

SOLVING GENERIC NONARCHIMEDEAN SEMIDEFINITE PROGRAMS USING STOCHASTIC GAME ALGORITHMS

XAVIER ALLAMIGEON, STÉPHANE GAUBERT, AND MATEUSZ SKOMRA

ABSTRACT. A general issue in computational optimization is to develop combinatorial algorithms for semidefinite programming. We address this issue when the base field is *nonarchimedean*. We provide a solution for a class of semidefinite feasibility problems given by generic matrices. Our approach is based on tropical geometry. It relies on tropical spectrahedra, which are defined as the images by the valuation of nonarchimedean spectrahedra. We establish a correspondence between generic tropical spectrahedra and zero-sum stochastic games with perfect information. The latter have been well studied in algorithmic game theory. This allows us to solve nonarchimedean semidefinite feasibility problems using algorithms for stochastic games. These algorithms are of a combinatorial nature and work for large instances.

1. INTRODUCTION

Semidefinite programming consists in optimizing a linear function over a spectrahedron. The latter is a subset of \mathbb{R}^n defined by linear matrix inequalities, i.e., a set of the form

$$\mathcal{S} = \{x \in \mathbb{R}^n : Q^{(0)} + x_1 Q^{(1)} + \cdots + x_n Q^{(n)} \succcurlyeq 0\}$$

where the $Q^{(k)}$ are symmetric matrices of order m , and \succcurlyeq denotes the Loewner order on the space of symmetric matrices. By definition, $X \succcurlyeq Y$ if and only if $X - Y$ is positive semidefinite.

Semidefinite programming is a fundamental tool in convex optimization. It is used to solve various applications from engineering sciences, and also to obtain approximate solutions or bounds for hard problems arising in combinatorial optimization and semialgebraic optimization. We refer the reader to [BPT13] and [GM12] for information.

Semidefinite programs are usually solved via interior point methods. The latter provide an approximate solution in a polynomial number of iterations, provided that a strictly feasible initial solution, i.e., a point belonging to the interior of the set \mathcal{S} , is known. We refer the reader to [dKV16] for a detailed analysis of the complexity in the bit model of interior point methods for semidefinite programming, and for a discussion of earlier complexity results based on the ellipsoid method.

Semidefinite programming becomes a much harder matter if one requires an exact solution. The feasibility problem (deciding the emptiness of the set \mathcal{S}) belongs to $\text{NP}_{\mathbb{R}} \cap \text{coNP}_{\mathbb{R}}$, where the subscript \mathbb{R} refers to the BSS model of computation [Ram97]. It is not known to be in NP in the bit model. A difficulty here is that all feasible points may have entries of absolute value doubly exponential in the size of the input. Also, there may be no rational solution [Sch16]. Beyond their theoretical interest, exact algorithms for semidefinite programming may be useful to address problems of formal proofs, which sometimes lead to challenging (degenerate) instances [MC11]. Known exact methods rely either on general purpose semialgebraic techniques

Date: October 15, 2018.

2010 Mathematics Subject Classification. 90C22, 91A15, 12J25, 14T05.

Key words and phrases. Semidefinite programming, stochastic mean payoff games, nonarchimedean fields, tropical geometry.

This manuscript has been accepted for publication in Journal of Symbolic Computation, <https://doi.org/10.1016/j.jsc.2017.07.002>. ©2018. The present file is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>. The three authors were partially supported by the ANR projects CAFEIN (ANR-12-INSE-0007) and MALTHY (ANR-13-INSE-0003), by the PGMO program of EDF and Fondation Mathématique Jacques Hadamard, and by the “Investissement d’avenir”, référence ANR-11-LABX-0056-LMH, LabEx LMH. M. Skomra is supported by a grant from Région Ile-de-France.

(like cylindrical decomposition) or on dedicated methods based on the computation of critical points, see the recent work [HNSD16] and the references therein.

Semidefinite programming is meaningful in any real closed field, and in particular in *non-archimedean* (real closed) fields like the field of Puiseux series with real coefficients. Such nonarchimedean semidefinite programming problems arise when considering *parametric* semidefinite programming problems over the reals, or structured problems in which the entries of the matrices have different orders of magnitudes. They are also of an intrinsic interest, since, by analogy with the situation in linear programming [Meg89], shifting to the nonarchimedean case is expected to shed light on the complexity of the classical problem over the reals. We note that the fields of Puiseux series are representative of the general nonarchimedean situation, since any nonarchimedean ordered field can be embedded in a field of generalized (Hahn) formal series [Ste10, Th. 5.2.20].

Description of main results. We address semidefinite programming in the nonarchimedean case, to which methods from tropical geometry can be applied. These methods are expected to allow one, in generic situations, to reduce semialgebraic problems to combinatorial problems, involving only the nonarchimedean valuations (leading exponents) of the coefficients of the input.

We exploit the tropical approach by considering tropical spectrahedra, defined as the images by the valuation of nonarchimedean spectrahedra. Yu [Yu15] showed that the image of the cone of semidefinite matrices is obtained by tropicalizing the nonnegativity conditions of minors of order 1 and 2. More generally, we studied the tropicalization of spectrahedra in the companion paper [AGS16b]. We showed that, under a genericity condition, tropical spectrahedra are described by explicit inequalities still arising from the nonnegativity of tropical minors of order 1 and 2. We recall these results in Section 3. Moreover, we show that we can reduce the tropical semidefinite feasibility problem to the subproblem in which tropical spectrahedra are defined by matrices with a Metzler sign pattern.

In this paper, we show (Theorem 26) that the feasibility problem for a generic tropical spectrahedron is equivalent to solving a stochastic mean payoff game (with perfect information). The complexity of these games is a long-standing open problem. They are not known to be polynomial, however they belong to the class $\text{NP} \cap \text{coNP}$, and they can be solved efficiently in practice.

This allows us to apply stochastic game algorithms to solve nonarchimedean semidefinite feasibility problems (Section 6). We obtain in this way both theoretical bounds and a practicable method which solves some large scale instances.

Related work. This work originates from the equivalence between tropical polyhedral feasibility problems and deterministic zero-sum games with mean payoff, established by Akian, Gaubert, and Guterman [AGG12]. The novelty here is the handling of nonlinear semialgebraic convex problems, and the proof that the well-known class of *stochastic* mean payoff games correspond precisely to semidefinite feasibility problems with a Metzler structure, therefore relating two classes of problems which both have an unsettled complexity.

Moreover, as mentioned above, the computation of exact solutions of semidefinite programming problems is of current interest. In particular, Nie, Ranestad, and Sturmfels [NRS10] provided complexity measures based on the notion of algebraic degree, and dedicated algorithms have been developed by Henrion, Naldi, and Safey El Din [HNSD16, Nal18].

2. STATEMENT OF THE PROBLEM AND ILLUSTRATION OF OUR APPROACH

A convenient choice of nonarchimedean structure, which we make in this paper, is the *field* \mathbb{K} of (*absolutely convergent generalized real*) *Puiseux series*, which is composed of functions of a real positive parameter t of the form

$$(1) \quad \mathbf{x} = \sum_{i=1}^{\infty} c_{\lambda_i} t^{\lambda_i},$$




FIGURE 1. Stochastic game associated with the spectrahedron considered in Section 2.

where $(\lambda_i)_{i \geq 1}$ is a strictly decreasing sequence of real numbers that is either finite or unbounded, $c_{\lambda_i} \in \mathbb{R} \setminus \{0\}$ for all λ_i , and the latter series is required to be absolutely convergent for t large enough. There is also a special, empty series, which is denoted by 0. We denote by $\text{lc}(\mathbf{x})$ the coefficient c_{λ_1} of the leading term in the series \mathbf{x} , with the convention that $\text{lc}(0) = 0$. The addition and multiplication in \mathbb{K} are defined in a natural way. Moreover, \mathbb{K} can be endowed with a linear order \geq , which is defined as $\mathbf{x} \geq \mathbf{y}$ if $\text{lc}(\mathbf{x} - \mathbf{y}) \geq 0$. We denote $\mathbb{K}_{\geq 0}$ the set of nonnegative series \mathbf{x} , i.e., satisfying $\mathbf{x} \geq 0$. The *valuation* of an element $\mathbf{x} \in \mathbb{K}$ as in (1) is defined as the greatest exponent λ_1 occurring in the series. Equivalently, the valuation is given by

$$\lim_{t \rightarrow +\infty} \log_t |\mathbf{x}(t)|,$$

where $\log_t(z) := \log(z)/\log(t)$. Up to a change of variables, \mathbb{K} coincides with the field of *generalized Dirichlet series* introduced by Hardy and Riesz [HR15]. Indeed, ordinary Dirichlet series are obtained by setting $t := \exp(s)$ and $\lambda_i := -\log i$. Van den Dries and Speissegger showed, among other results, that \mathbb{K} is real closed [vdDS98, Corollary 9.2 and Section 10.2]. The familiar field of ordinary, convergent Puiseux series, is obtained by requiring the sequence λ_i to consist of rational numbers included in an arithmetic progression. Working with the larger field \mathbb{K} , whose value group is \mathbb{R} instead of \mathbb{Q} , leads to more transparent results.

In this paper, we consider the semidefinite feasibility problem over the field \mathbb{K} . It is more convenient to start with the homogeneous case. More precisely, given symmetric matrices $\mathbf{Q}^{(1)}, \dots, \mathbf{Q}^{(n)} \in \mathbb{K}^{m \times m}$, we focus on the problem of determining whether or not the following spectrahedral cone

$$\{\mathbf{x} \in \mathbb{K}_{\geq 0}^n : \mathbf{x}_1 \mathbf{Q}^{(1)} + \dots + \mathbf{x}_n \mathbf{Q}^{(n)} \succcurlyeq 0\}$$

is *trivial*, meaning that it is reduced to the identically null point of $\mathbb{K}_{\geq 0}^n$. The linear map $\mathbf{x} \mapsto \mathbf{x}_1 \mathbf{Q}^{(1)} + \dots + \mathbf{x}_n \mathbf{Q}^{(n)}$ is a *matrix pencil*, which we denote by $\mathbf{Q}(\mathbf{x})$ for more brevity. The aforementioned decision problem is in fact equivalent to the nontriviality problem of general spectrahedral cones in which the nonnegativity condition $\mathbf{x} \in \mathbb{K}_{\geq 0}^n$ is relaxed. Indeed, given a pencil $\mathbf{Q}(\mathbf{x})$ of symmetric matrices, the spectrahedron $\{\mathbf{x} \in \mathbb{K}^n : \mathbf{Q}(\mathbf{x}) \succcurlyeq 0\}$ is trivial if and only if the spectrahedron $\{(\mathbf{y}, \mathbf{z}) \in \mathbb{K}_{\geq 0}^{2n} : (\mathbf{y}_1 - \mathbf{z}_1) \mathbf{Q}^{(1)} + \dots + (\mathbf{y}_n - \mathbf{z}_n) \mathbf{Q}^{(n)} \succcurlyeq 0\}$ is trivial. Even if the instances arising in this way are unlikely to be generic in the sense we discuss later in Section 3, we consider that the decision problem which we focus on already retains much of the complexity of the semidefinite feasibility problem over the field of Puiseux series. As we shall see in Section 5.2, our method can be extended to handle the feasibility problem for affine spectrahedra given by the relation $\{\mathbf{x} \in \mathbb{K}_{\geq 0}^n : \mathbf{Q}^{(0)} + \mathbf{x}_1 \mathbf{Q}^{(1)} + \dots + \mathbf{x}_n \mathbf{Q}^{(n)} \succcurlyeq 0\}$.

Our approach is best explained when the off-diagonal entries of the matrices $\mathbf{Q}^{(1)}, \dots, \mathbf{Q}^{(n)}$ are nonpositive. We associate to these matrices the following zero-sum game. There are two players, Player Min and Player Max, who control disjoint sets of states. The states of Player Min can be identified to the variables $\mathbf{x}_1, \dots, \mathbf{x}_n$. The states of Player Max can be identified to the row (or column) indices of the matrices, i.e., to the elements of $\{1, \dots, m\}$. In state \mathbf{x}_k , Player Min chooses a subset $\{i, j\} \subset \{1, \dots, m\}$ such that the entry $\mathbf{Q}_{ij}^{(k)}$ is negative. Next, Player Min pays to Player Max the opposite of the valuation of $\mathbf{Q}_{ij}^{(k)}$. Then, Nature selects one

element among i or j at random, with uniform probabilities (i.e., i or j is drawn with probability $1/2$).¹ If i is drawn, meaning that the current state is now i , Player Max chooses a variable \mathbf{x}_i such that $\mathbf{Q}_{ii}^{(l)}$ has a positive sign. He receives the valuation of $\mathbf{Q}_{ii}^{(l)}$ from Player Min, and the next state becomes \mathbf{x}_i . If j is drawn, the rule of move and the payment are identical, up to the replacement of i by j . A (stationary) policy of one player is a map which associates to a state of the player an admissible move. We are interested in infinite plays, with an infinite number of turns. Player Max looks for a policy which maximizes the mean payment received from Player Min per time unit, while Player Min looks for a policy which minimizes it.

We can think informally of this construction as follows: Player Min wishes to show that the semidefinite programming problem is infeasible, whereas Player Max wishes to show that it is feasible.

Let us now illustrate our approach on an example in dimension 3. We consider the following pencil of symmetric matrices

$$\mathbf{Q}(\mathbf{x}) := \begin{bmatrix} t\mathbf{x}_3 & -\mathbf{x}_1 & -t^{3/4}\mathbf{x}_3 \\ -\mathbf{x}_1 & t^{-1}\mathbf{x}_1 + t^{-5/4}\mathbf{x}_3 - \mathbf{x}_2 & -\mathbf{x}_3 \\ -t^{3/4}\mathbf{x}_3 & -\mathbf{x}_3 & t^{9/4}\mathbf{x}_2 \end{bmatrix},$$

and we aim at checking whether or not the spectrahedron $\{\mathbf{x} \in \mathbb{K}_{\geq 0}^3 : \mathbf{Q}(\mathbf{x}) \succcurlyeq 0\}$ is trivial. The associated game is depicted in Figure 1. The states of Player Min are depicted by circles. The states of Player Max are depicted by squares. The states in which Nature plays are depicted by full dots. The admissible moves of the game are represented by the edges between the states. The corresponding payments received by Player Max are indicated on these edges.

Observe that both players in this game have only two policies: at state $\textcircled{2}$ Player Max can choose the move that goes to $\textcircled{1}$ or the one that goes to $\textcircled{3}$, whereas at state $\textcircled{3}$ Player Min can choose the move $\{\textcircled{1}, \textcircled{3}\}$ or the move $\{\textcircled{2}, \textcircled{3}\}$.

Suppose that Player Max chooses the policy which goes to state $\textcircled{1}$ from state $\textcircled{2}$. If Player Min plays using the move $\{\textcircled{1}, \textcircled{3}\}$ at state $\textcircled{3}$, then a standard computation on Markov chain (which we present in Example 9) shows that the long-term average payoff of Player Max is equal to $3/40$. Similarly, if Player Min chooses the move $\{\textcircled{2}, \textcircled{3}\}$ instead, then the payoff of Player Max is equal to $1/56$. Therefore, Player Max has a policy which guarantees him to win the game, i.e., to obtain a nonnegative payoff. The main theorem of this paper states that this fact is equivalent to the nontriviality of the spectrahedron $\{\mathbf{x} \in \mathbb{K}_{\geq 0}^3 : \mathbf{Q}(\mathbf{x}) \succcurlyeq 0\}$. As we shall see, a winning policy of Player Max (resp. Min) provides a feasibility (resp. infeasibility) certificate. The mean payoff represents a feasibility/infeasibility margin.

3. PRELIMINARY RESULTS ON TROPICAL SPECTRAHEDRA

3.1. Tropical algebra. In this section, we recall some basic concepts of tropical algebra.

The *tropical semifield* \mathbb{T} is the set $\mathbb{R} \cup \{-\infty\}$, endowed with the addition $x \oplus y := \max(x, y)$ and the multiplication $x \odot y := x + y$. The term semifield refers to the fact that the addition does not have an opposite law. We use the notation $\bigoplus_{i=1}^n a_i = a_1 \oplus \dots \oplus a_n$ and $a^{\odot n} = a \odot \dots \odot a$ (n times). We also endow \mathbb{T} with the standard order \leq . The reader may consult [But10, MS15] for more information on the tropical semifield.

We denote by $\text{val} : \mathbb{K} \rightarrow \mathbb{R} \cup \{-\infty\}$ the map which associates a Puiseux series $\mathbf{x} \in \mathbb{K}$ to its valuation. We use the convention $\text{val}(0) = -\infty$. It is immediate to see that the map val satisfies the following properties

$$\begin{aligned} (2) \quad & \text{val}(\mathbf{x} + \mathbf{y}) \leq \max(\text{val}(\mathbf{x}), \text{val}(\mathbf{y})) \\ (3) \quad & \text{val}(\mathbf{x}\mathbf{y}) = \text{val}(\mathbf{x}) + \text{val}(\mathbf{y}) \end{aligned}$$

¹We allow the case $i = j$. In other words, Player Min can choose a subset $\{i\} \subset \{1, \dots, m\}$ such that $\mathbf{Q}_{ii}^{(k)}$ is negative. In this case, Nature selects i .

meaning that val is a *nonarchimedean valuation*. Moreover, the equality holds in (2) if the leading terms of \mathbf{x} and \mathbf{y} do not cancel, which is the case if $\text{val}(\mathbf{x}) \neq \text{val}(\mathbf{y})$ or if $\mathbf{x}, \mathbf{y} \geq 0$. In particular, the map val yields an order-preserving morphism of semifields from $\mathbb{K}_{\geq 0}$ to \mathbb{T} .

The *sign* of a series $\mathbf{x} \in \mathbb{K}$ is equal to $+1$ if $\mathbf{x} > 0$, -1 if $\mathbf{x} < 0$, and 0 otherwise. The *signed valuation*, denoted by sval , associates with a series $\mathbf{x} \in \mathbb{K}$ the couple $(\text{sign}(\mathbf{x}), \text{val}(\mathbf{x}))$. We denote by \mathbb{T}_{\pm} the image of \mathbb{K} by sval . We call it the set of *signed tropical numbers*. For brevity, we denote an element of the form (ϵ, a) by a if $\epsilon = 1$, $\ominus a$ if $\epsilon = -1$, and $-\infty$ if $\epsilon = 0$. Here, \ominus is a formal symbol. We call the elements of the first and second kind the *positive* and *negative* tropical numbers, respectively. We denote by \mathbb{T}_+ and \mathbb{T}_- the corresponding sets. In this way, (-2) is tropically positive, but $\ominus(-2)$ is tropically negative. Also, \mathbb{T} is embedded in \mathbb{T}_{\pm} , i.e., $\mathbb{T} = \mathbb{T}_+ \cup \{-\infty\}$. In \mathbb{T}_{\pm} , we define a *modulus* function, $|\cdot|: \mathbb{T}_{\pm} \rightarrow \mathbb{T}$, as $|\ominus\infty| = -\infty$ and $|a| = |\ominus a| = a$ for all $a \in \mathbb{T}_+$. We point out that \odot straightforwardly extends to \mathbb{T}_{\pm} using the standard rules for the sign, for instance $2 \odot (\ominus 3) = \ominus 5$. In contrast, we only partially extend the tropical addition \oplus to elements of \mathbb{T}_{\pm} of identical sign, e.g., $2 \oplus 3 = 3$ and $(\ominus 2) \oplus (\ominus 3) = \ominus 3$. It is possible to embed \mathbb{T}_{\pm} in an idempotent semiring, called the symmetrized tropical semiring [AGG09], so that the tropical addition becomes defined for all elements of \mathbb{T}_{\pm} . Alternatively, this addition may be defined by working in the setting of hyperfields [Vir10, CC11, BB16]. Here, we shall only use \mathbb{T}_{\pm} as a convenient notation, and we do not rely on an algebraic structure on \mathbb{T}_{\pm} .

We shall extend the valuation maps val and sval to vectors and matrices in a coordinate-wise manner.

Finally, we use the notion of tropical polynomials. A *tropical (signed) polynomial* over the variables X_1, \dots, X_n is a formal expression of the form

$$(4) \quad P(X) = \bigoplus_{\alpha \in \Lambda} a_{\alpha} \odot X_1^{\odot \alpha_1} \odot \dots \odot X_n^{\odot \alpha_n},$$

where $\Lambda \subset \{0, 1, 2, \dots\}^n$, and $a_{\alpha} \in \mathbb{T}_{\pm} \setminus \{-\infty\}$ for all $\alpha \in \Lambda$. We say that the tropical polynomial P *vanishes* on the point $x \in \mathbb{T}_{\pm}^n$ if the terms $a_{\alpha} \odot x_1^{\odot \alpha_1} \odot \dots \odot x_n^{\odot \alpha_n}$ which have the greatest modulus do not have the same sign. If P does not vanish on x , we define $P(x)$ as the tropical sum of the terms which have the greatest modulus. As an example, if $P(X) = 2 \odot X_1^{\odot 3} \odot X_2^{\odot 4} \oplus (\ominus 0) \odot X_2$, then $P(1, \ominus 5) = 25$, $P(1, -5) = \ominus(-5)$, whereas P vanishes on $(1, -5/3)$. These definitions are motivated by the following immediate observation. Suppose that

$$(5) \quad P(X) = \sum_{\alpha \in \Lambda} \mathbf{a}_{\alpha} X_1^{\alpha_1} \dots X_n^{\alpha_n} \in \mathbb{K}[X_1, \dots, X_n]$$

and let P be defined as in (4) with $a_{\alpha} := \text{sval}(\mathbf{a}_{\alpha})$. Then, for all $\mathbf{x} \in \mathbb{K}^n$,

$$\text{sval}(P(\mathbf{x})) = P(\text{sval}(\mathbf{x})),$$

provided that P does not vanish on $\text{sval}(\mathbf{x})$.

Given a polynomial \mathbf{P} as in (5), we denote by \mathbf{P}^+ the polynomial formed by the terms $\mathbf{a}_{\alpha} X_1^{\alpha_1} \dots X_n^{\alpha_n}$ such that $\mathbf{a}_{\alpha} > 0$. Similarly, \mathbf{P}^- refers to the polynomial consisting of the terms $-\mathbf{a}_{\alpha} X_1^{\alpha_1} \dots X_n^{\alpha_n}$ verifying $\mathbf{a}_{\alpha} < 0$. In this way, $\mathbf{P} = \mathbf{P}^+ - \mathbf{P}^-$. We also use the analogues of these polynomials in the tropical setting. If P is the tropical polynomial given in (4), we define P^+ (resp. P^-) as the tropical polynomial generated by the terms $|a_{\alpha}| \odot X_1^{\odot \alpha_1} \odot \dots \odot X_n^{\odot \alpha_n}$ where $a_{\alpha} \in \mathbb{T}_+$ (resp. \mathbb{T}_-). Observe that the quantities $P^+(x)$ and $P^-(x)$ are well defined for all $x \in \mathbb{T}^n$, since the tropical polynomials P^+ and P^- only involve tropically positive coefficients.

Throughout the paper, we denote the set $\{1, \dots, k\}$ by $[k]$.

3.2. Tropicalization of nonarchimedean spectrahedra. We now discuss the class of tropical spectrahedra. These objects were introduced in our companion paper [AGS16b]. This paper also contains the detailed proofs of all the results presented in this section.

Definition 1. A set $\mathcal{S} \subset \mathbb{T}^n$ is said to be a *tropical spectrahedron* if there exists a spectrahedron $\mathcal{S} \subset \mathbb{K}_{\geq 0}^n$ such that $\mathcal{S} = \text{val}(\mathcal{S})$.

If $\mathcal{S} = \text{val}(\mathcal{S})$, then we refer to \mathcal{S} as the *tropicalization* of the spectrahedron \mathcal{S} , and \mathcal{S} is said to be a *lift* (over the field \mathbb{K}) of \mathcal{S} . Checking the triviality of a spectrahedral cone $\mathcal{S} \subset \mathbb{K}_{\geq 0}^n$ is equivalent to determining whether or not the corresponding tropical spectrahedron $\mathcal{S} = \text{val}(\mathcal{S})$ is *trivial*, i.e., is reduced to the singleton $\{(-\infty, \dots, -\infty)\}$. Therefore, it is convenient to exploit some explicit description of the set \mathcal{S} . As stated in Theorem 3, such a description can be obtained from the tropical minors of order 1 and 2, provided that the valuation of the coefficients of the matrices $\mathbf{Q}^{(k)}$ is generic. To this purpose, given $i, j \in [m]$, we denote by $Q_{ij}(X)$ the tropical polynomial:

$$Q_{ij}(X) := Q_{ij}^{(1)} \odot X_1 \oplus \dots \oplus Q_{ij}^{(n)} \odot X_n.$$

Definition 2. Let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ be symmetric tropical matrices. We introduce the set $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ (or simply \mathcal{S}) of points $x \in \mathbb{T}^n$ that fulfill the following two conditions:

- for all $i \in [m]$, $Q_{ii}^+(x) \geq Q_{ii}^-(x)$;
- for all $i, j \in [m]$, $i < j$, we have $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq (Q_{ij}^+(x) \oplus Q_{ij}^-(x))^{\odot 2}$ or $Q_{ij}^+(x) = Q_{ij}^-(x)$.

Given $d \geq 1$, the *support* of a point $y \in \mathbb{T}^d$ is defined as the set of indices $k \in [d]$ such that $y_k \neq -\infty$. Given a nonempty subset $K \subset [d]$, and a set $Y \subset \mathbb{T}^d$, we define the *stratum* of Y associated with K as the subset of \mathbb{R}^K formed by the projection $(y_k)_{k \in K}$ of the points $y \in Y$ with support K . We say that a set $Y \subset \mathbb{T}^d$ is *negligible* if every stratum of Y has Lebesgue measure zero.

Theorem 3 ([AGS16b, Theorems 32 and 37]). *There exists a negligible set $Y \subset \mathbb{T}^d$ with $d = nm(m+1)/2$ such that the following property holds. Suppose that $\mathbf{Q}^{(1)}, \dots, \mathbf{Q}^{(n)} \in \mathbb{K}^{m \times m}$ is a sequence of symmetric matrices and let $\mathcal{S} = \{x \in \mathbb{K}_{\geq 0}^n : \mathbf{Q}(x) \succcurlyeq 0\}$ be the associated spectrahedron. Denote $Q^{(k)} = \text{sval}(\mathbf{Q}^{(k)})$ for all k . If the vector with entries $|Q_{ij}^{(k)}|$ (for $i \leq j$) does not belong to Y , then we have*

$$\text{val}(\mathcal{S}) = \mathcal{S}(Q^{(1)}, \dots, Q^{(n)}).$$

The set Y from Theorem 3 can be constructed explicitly, as a finite union of hyperplanes. Every hyperplane arises from a condition expressing the absence of a flow in a certain directed hypergraph associated with $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$. This construction is detailed in [AGS16b, Section 5.4]. Nevertheless, the number of directed hypergraphs (and subsequently, of hyperplanes) to be considered is exponential. For this reason, we do not give an explicit description of Y .

One important special case is when the matrices $\mathbf{Q}^{(k)}$ are Metzler matrices. Recall that a matrix is called (*negated*) *Metzler matrix* if its off-diagonal coefficients are nonpositive. Similarly, we say that a matrix $M \in \mathbb{T}_{\pm}^{m \times m}$ is a *tropical Metzler matrix* if $M_{ij} \in \mathbb{T}_- \cup \{-\infty\}$ for all $i \neq j$. Under the assumption that the matrices $Q^{(k)}$ are tropical Metzler matrices, Definition 2 gets slightly simpler:

Lemma 4. *Suppose that the matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ are symmetric tropical Metzler matrices. Then the set $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ consists of the points $x \in \mathbb{T}^n$ such that:*

- for all $i \in [m]$, $Q_{ii}^+(x) \geq Q_{ii}^-(x)$;
- for all $i, j \in [m]$, $i < j$, $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq (Q_{ij}(x))^{\odot 2}$.

Observe that in the lemma above, the term $Q_{ij}(x)$ ($i \neq j$) is well defined for any $x \in \mathbb{T}^n$ thanks to the Metzler property of the matrices $Q^{(k)}$. We can equivalently rewrite the constraints defining \mathcal{S} using classical notation as follows: for all $i \in [m]$,

$$(6) \quad \max_{Q_{ii}^{(k)} \in \mathbb{T}_+} (Q_{ii}^{(k)} + x_k) \geq \max_{Q_{ii}^{(l)} \in \mathbb{T}_-} (|Q_{ii}^{(l)}| + x_l),$$

and for all $i, j \in [m]$ such that $i < j$,

$$(7) \quad \max_{Q_{ii}^{(k)} \in \mathbb{T}_+} (Q_{ii}^{(k)} + x_k) + \max_{Q_{jj}^{(k')} \in \mathbb{T}_+} (Q_{jj}^{(k')} + x_{k'}) \geq 2 \max_{l \in [n]} (|Q_{ij}^{(l)}| + x_l).$$




FIGURE 2. The tropicalization of the spectrahedron of Section 2.

When the matrices $Q^{(k)}$ are Metzler, we refer to the set \mathcal{S} as a *tropical Metzler spectrahedron*. This terminology relies on the fact that in this case we can build a spectrahedron in $\mathbb{K}_{\geq 0}^n$ which is a lift of \mathcal{S} [AGS16b, Proposition 23]. We point out that if we restrict our considerations to the case of Metzler matrices, then the construction of the set Y from Theorem 3 can be simplified, but it still requires an exponential number of hyperplanes.

Example 5. Let us illustrate these results on the spectrahedron defined in Section 2. The corresponding tropical matrices $Q^{(1)}, Q^{(2)}, Q^{(3)} \in \mathbb{T}^{3 \times 3}$ are of Metzler type, and the associated tropical Metzler spectrahedron \mathcal{S} is defined by the constraints:

$$\begin{aligned} \max(-1 + x_1, -5/4 + x_3) &\geq x_2 \\ \max(x_1 + x_3, -1/4 + 2x_3) &\geq 2x_1 \\ x_2 &\geq -7/4 + x_3 \\ \max(5/4 + x_1 + x_2, 1 + x_2 + x_3) &\geq 2x_3 \end{aligned}$$

The first inequality comes from (6) with $i = 2$, and the last three constraints from (7).² The intersection of \mathcal{S} with the hyperplane $x_3 = 0$ is depicted in Figure 2. The matrices $Q^{(k)}$ fulfill the genericity conditions mentioned in Theorem 3 so that the tropicalization of the spectrahedron of Section 2 is precisely described by the four inequalities given above.

Even though the Metzler case may look special, we can reduce the problem of deciding whether the set $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ is trivial to the subproblem in which the matrices are Metzler. To show this, we prove that every set $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ is a projection of a tropical Metzler spectrahedron. Furthermore, this Metzler spectrahedron can be constructed in poly-time.

Proposition 6. *Let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ be symmetric tropical matrices. Then the set $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ is a projection of a tropical Metzler spectrahedron.*

Proof. Let us denote $\mathcal{S}(Q^{(1)}, \dots, Q^{(n)})$ by \mathcal{S} . For every pair $(i, j) \in [m] \times [m]$, $i < j$ we introduce a variable y_{ij} and we consider the set $\tilde{\mathcal{S}}$ defined as the set of the points $(x, y) \in \mathbb{T}^n \times \mathbb{T}^{m(m-1)/2}$ that fulfill the following conditions:

- for all $i \in [m]$, $Q_{ii}^+(x) \geq Q_{ii}^-(x)$;
- for all $i, j \in [m]$, $i < j$, $y_{ij} \oplus Q_{ij}^+(x) \geq Q_{ij}^-(x)$;
- for all $i, j \in [m]$, $i < j$, $y_{ij} \oplus Q_{ij}^-(x) \geq Q_{ij}^+(x)$;
- for all $i, j \in [m]$, $i < j$, $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq y_{ij}^{\odot 2}$.

The set $\tilde{\mathcal{S}}$ is a tropical Metzler spectrahedron defined by matrices of size $m^2 \times m^2$. We claim that \mathcal{S} is a projection of $\tilde{\mathcal{S}}$. First, let us take a point $x \in \mathcal{S}$. For every (i, j) with $i < j$ such that $Q_{ij}^-(x) = Q_{ij}^+(x)$ we put $y_{ij} = -\infty$. For every (i, j) with $i < j$ such that $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq (Q_{ij}^+(x) \oplus Q_{ij}^-(x))^{\odot 2}$ we put $y_{ij} = Q_{ij}^+(x) \oplus Q_{ij}^-(x)$. It is clear that we have $(x, y) \in \tilde{\mathcal{S}}$. Conversely, let $(x, y) \in \tilde{\mathcal{S}}$. For every (i, j) with $i < j$ we consider two cases. If $y_{ij} \geq Q_{ij}^-(x)$, then we have $y_{ij} \geq Q_{ij}^+(x)$ and hence $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq y_{ij} \geq (Q_{ij}^+(x) \oplus Q_{ij}^-(x))^{\odot 2}$. If $y_{ij} < Q_{ij}^-(x)$, then we have $Q_{ii}^+(x) \geq Q_{ij}^-(x)$ and $Q_{ij}^-(x) \geq Q_{ij}^+(x)$. Hence $Q_{ii}^+(x) = Q_{ij}^-(x)$. Therefore $x \in \mathcal{S}$. \square

²The constraints of the form (6) with $i = 1, 3$ are trivial.

In the light of Proposition 6, we restrict in the rest of the paper to the problem of deciding whether a tropical Metzler spectrahedron is trivial.

4. TROPICAL SPECTRAHEDRA AND STOCHASTIC GAMES

4.1. Stochastic mean payoff games. In this section, we present the class of games which is related to nonarchimedean semidefinite feasibility problems. For simplicity, we refer to them as stochastic mean payoff games, although as we shall see, this terminology usually corresponds to a larger class of games. This abuse of language is justified by the fact that the associated decision and computational problems are poly-time equivalent, as discussed below.

In our setting, a *stochastic mean payoff game* involves two players, Max and Min, who control disjoint sets of states. The states owned by Max and Min are respectively indexed by elements of $[m]$ and $[n]$. We will use the symbols i, j to refer to states of Player Max, and k, l to states of Player Min. Both players alternatively move a pawn over these states as follows. When the pawn is on a state $k \in [n]$, Player Min chooses an action $a \in A^{(k)}$, which is defined as a subset of states of Max of cardinality 1 or 2: if $a = \{i\}$, then the pawn is moved to the state i , while if $a = \{i, j\}$ with $i \neq j$, it is moved to the state i (resp. j) with probability $1/2$. In both cases, Player Max receives from Player Min a reward denoted by r_k^a . Once the pawn is on a state $i \in [m]$, Player Max picks an action $b \in B^{(i)}$, where b is a subset of states of Min of cardinality 1. Then, Player Max moves the pawn to the state l such that $b = \{l\}$, and Player Min pays him a payment denoted by r_i^b .

We suppose that Player Min starts the game and that players can always make the next move, i.e., that $A^{(k)} \neq \emptyset$ and $B^{(i)} \neq \emptyset$ for all $i \in [m]$ and $k \in [n]$.

A *policy for Player Min* is a function mapping every state $k \in [n]$ to an action $\sigma(k)$ in $A^{(k)}$. Analogously, a *policy for Player Max* is a function τ such that $\tau(i) \in B^{(i)}$ for all $i \in [m]$. Suppose that the game starts from a state k^* of Player Min. When players play according to a couple (σ, τ) of policies, the movement of the pawn is described by a Markov chain on the space $E = [m] \uplus [n]$. The average payoff of Player Max in the long-term is then defined as the average payoff of the controller in this Markov chain. In other words, the payoff of Player Max is given by

$$(8) \quad g_{k^*}(\sigma, \tau) = \lim_{N \rightarrow \infty} \mathbb{E}_{\sigma, \tau} \left(\frac{1}{2N} \sum_{p=1}^N (r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\tau(i_p)}) \right),$$

where the expectation $\mathbb{E}_{\sigma, \tau}$ is taken over all the trajectories $k_1, i_1, k_2, \dots, i_p$ starting from $k_1 = k^*$ in the Markov chain. The goal of Player Max is to find a policy which maximizes his average payoff, while Player Min aims at minimizing this quantity. The basic theorem of stochastic mean payoff games is the existence of “optimal” policies for both players. This was proven by Liggett and Lippman [LL69].

Theorem 7. *There exists a pair of policies $(\bar{\sigma}, \bar{\tau})$ and a unique vector $\chi \in \mathbb{R}^n$ such that for all initial states $k \in [n]$, the following two conditions are satisfied:*

- *for each policy σ of Player Min, $\chi_k \leq g_k(\sigma, \bar{\tau})$;*
- *for each policy τ of Player Max, $\chi_k \geq g_k(\bar{\sigma}, \tau)$.*

The vector χ in Theorem 7 is referred to as the *value* of the game. Note that for every k , the quantity χ_k coincides with average payoff $g_k(\bar{\sigma}, \bar{\tau})$ associated with the couple of policies $(\bar{\sigma}, \bar{\tau})$. The first condition in Theorem 7 states that, by playing according to the policy $\bar{\tau}$, Player Max is certain to get an average payoff greater than or equal to the value χ_k associated with the initial state. Symmetrically, Player Min is ensured to limit her average loss to the quantity χ_k by following the policy $\bar{\sigma}$.

A state $k \in [n]$ is said to be *winning (for Player Max)* if the value χ_k of the game starting from the initial state k is nonnegative. We denote by SMPG the following decision problem: “given a stochastic mean payoff game, does there exist an initial state which is winning for Player Max?” It can be shown that if one of the players has no choice (i.e., has only one possible policy), then

the problem of finding values and optimal policies can be solved in polynomial time by linear programming (see, e.g., [FV07, Section 2.9]). This readily implies that SMPG is in $\text{NP} \cap \text{coNP}$ (one guesses an optimal policy for one player, fixes it and then finds an optimal policy and the value of a 1-player game). However, as discussed in the introduction, the question whether there exists a poly-time algorithm to solve this problem is open.

Remark 8. In the literature, stochastic mean payoff games correspond to a larger class of problems [AM09]. These more general games admit a value and optimal policies as well. It turns out that the associated decision problem is poly-time equivalent to the simpler problem SMPG defined above. Moreover, these decision problems are poly-time equivalent to the problem of computing the value and a pair of optimal policies. We refer to Section 7 for details.

Example 9. Let us revisit the example presented in Section 2, see Figure 1. As noted previously, both players in this example game have only two policies: at state $\boxed{2}$ Player Max can choose the action that goes to $\textcircled{1}$ or the action that goes to $\textcircled{3}$, whereas at state $\textcircled{3}$ Player Min can choose the action that goes to $\{\boxed{1}, \boxed{3}\}$ or the action that goes to $\{\boxed{2}, \boxed{3}\}$. Suppose that Player Max chooses the action that goes to $\textcircled{1}$ and that Player Min chooses the action that goes to $\{\boxed{1}, \boxed{3}\}$. The Markov chain obtained in this way has the transition matrix of form

$$P = \begin{bmatrix} 0 & U \\ V & 0 \end{bmatrix},$$

where U describes the probabilities of transition from circle states to square states, and V describes the probabilities of transition from square states to circle states, i.e.,

$$U = \begin{matrix} & \boxed{1} & \boxed{2} & \boxed{3} \\ \begin{matrix} \textcircled{1} \\ \textcircled{2} \\ \textcircled{3} \end{matrix} & \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1 & 0 \\ 1/2 & 0 & 1/2 \end{bmatrix} \end{matrix}, \quad V = \begin{matrix} & \textcircled{1} & \textcircled{2} & \textcircled{3} \\ \begin{matrix} \boxed{1} \\ \boxed{2} \\ \boxed{3} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \end{matrix}.$$

This chain has only one recurrent class and all states belong to this class. Moreover, it is easy to verify that

$$\pi = \frac{1}{10}(2, 1, 2, 2, 2, 1)$$

is the stationary distribution of this chain. Therefore, by Theorem 53, the payoff of Player Max is equal to

$$g(\{\boxed{1}, \boxed{3}\}, \{\textcircled{1}\}) = \frac{1}{10}(-\frac{6}{4} + 2 - 2 + \frac{9}{4}) = \frac{3}{40}.$$

Similarly, if Player Max chooses the action that goes to $\textcircled{1}$ and Player Min chooses the action that goes to $\{\boxed{2}, \boxed{3}\}$, then the transition matrix of the generated Markov chain has the form $P' = \begin{bmatrix} 0 & U' \\ V & 0 \end{bmatrix}$, where

$$U' = \begin{matrix} & \boxed{1} & \boxed{2} & \boxed{3} \\ \begin{matrix} \textcircled{1} \\ \textcircled{2} \\ \textcircled{3} \end{matrix} & \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1 & 0 \\ 0 & 1/2 & 1/2 \end{bmatrix} \end{matrix}$$

and V is the same as previously. In this case the chain also has only one recurrent class and every state belongs to this class. Furthermore,

$$\pi' = \frac{1}{14}(4, 1, 2, 2, 4, 1)$$

is the stationary distribution of this chain. Hence the payoff of Player Max is equal to

$$g(\{\boxed{2}, \boxed{3}\}, \{\textcircled{1}\}) = \frac{1}{14}(2 - 4 + \frac{9}{4}) = \frac{1}{56}.$$

In both cases, the payoff of Player Max is positive. Therefore, if Player Max chooses the action that goes to $\textcircled{1}$, then he is guaranteed to obtain a nonnegative payoff. One can check, by doing

the same calculations for the remaining policies, that $(\{\underline{2}, \underline{3}\}, \{\underline{1}\})$ is the unique couple of optimal policies in this game.

4.2. Shapley operators. One of the possible approaches to analyze stochastic mean payoff games is to introduce the associated *Shapley operator*, which is a map F from \mathbb{T}^n to itself defined as:

$$(9) \quad (F(x))_k = \min_{\substack{a \in A^{(k)} \\ a = \{i, j\}}} \left(r_k^a + \frac{1}{2} \left(\max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right) \right),$$

where we use the convention that $i = j$ when $a \in A^{(i)}$ is the singleton $\{i\}$ and we set $\frac{1}{2}(-\infty) = -\infty$.

Let us point out that \mathbb{T} can be endowed with a natural metric $d(x, y) = |\exp(x) - \exp(y)|$. This metric generates the product topology on \mathbb{T}^n in which the functions

$$\begin{aligned} \mathbb{T}^2 \ni (a, b) &\rightarrow a + b \in \mathbb{T}, \quad \mathbb{T}^2 \ni (a, b) \rightarrow \max\{a, b\} \in \mathbb{T}, \\ \mathbb{T}^2 \ni (a, b) &\rightarrow \min\{a, b\} \in \mathbb{T}, \quad \mathbb{T} \ni a \rightarrow \frac{1}{2}a \in \mathbb{T} \end{aligned}$$

are continuous. If $x, y \in \mathbb{T}^n$ are vectors, then we denote $x \leq y$ if the inequality $x_k \leq y_k$ is fulfilled for every $k \in [n]$. We say that a map $f: \mathbb{R}^n \rightarrow \mathbb{R}^d$ is *piecewise affine* if \mathbb{R}^n can be partitioned into finitely many polyhedra such that f restricted to any of these polyhedra is affine. The basic properties of Shapley operators are summarized in the next lemma. Hereafter, if $\lambda \in \mathbb{R}$ and $x \in \mathbb{T}^n$, then we use the notation $\lambda + x$ to denote the vector $(\lambda + x_1, \dots, \lambda + x_n)$.

Lemma 10. *The Shapley operator F has the following properties:*

- (i) *it is order preserving, i.e., for all $x, y \in \mathbb{T}^n$ such that $x \leq y$ we have $F(x) \leq F(y)$;*
- (ii) *it is additively homogeneous, i.e., if $\lambda \in \mathbb{R}$, then $F(\lambda + x) = \lambda + F(x)$ for all $x \in \mathbb{T}^n$;*
- (iii) *it is continuous in the topology of \mathbb{T}^n ;*
- (iv) *$F|_{\mathbb{R}^n}$ is nonexpansive in the supremum norm, i.e., for all $x, y \in \mathbb{R}^n$ we have $\|F(x) - F(y)\| \leq \|x - y\|$, where $\|x\| = \max_{k \in [n]} |x_k|$;*
- (v) *$F|_{\mathbb{R}^n}$ is piecewise affine.*

Proof. The first two properties follow trivially from the definition of F . The third follows from the remarks about topology of \mathbb{T}^n stated above. The fourth follows from the first two and the fact that F preserves \mathbb{R}^n . Indeed, for all $x, y \in \mathbb{R}^n$ we have $x \leq \|x - y\| + y$. Therefore $F(x) \leq F(\|x - y\| + y) = \|x - y\| + F(y)$. Analogously, $F(y) \leq \|x - y\| + F(x)$ and hence $\|F(x) - F(y)\| \leq \|x - y\|$. The last point is immediate from the definition of F . \square

A central question, given $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$, is to decide whether the limit $\lim_{N \rightarrow \infty} f^N(x)/N$ exists. In order to prove the existence of this limit, Kohlberg introduced the notion of an invariant half-line and proved the next theorem [Koh80].

Theorem 11. *Suppose that function $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is piecewise affine and nonexpansive in any norm. Then there exists a triple $(v, \eta, \gamma_0) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}$, $\gamma_0 \geq 0$, such that $f(v + \gamma\eta) = v + (\gamma + 1)\eta$ for all $\gamma \geq \gamma_0$. A pair (v, η) is called an invariant half-line. Furthermore, if (v_1, η_1) and (v_2, η_2) are invariant half-lines, then $\eta_1 = \eta_2$.*

In particular, we obtain $\lim_{N \rightarrow \infty} f^N(v)/N = \eta$. Moreover, for any $x \in \mathbb{R}^n$ and $N \geq 1$ we have $\|f^N(x) - f^N(v)\| \leq \|x - v\|$. Thus, we get the following corollary.

Corollary 12. *Suppose that $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ fulfills the assumptions of Theorem 11. Then, for every $x \in \mathbb{R}^n$ we have*

$$\lim_{N \rightarrow \infty} \frac{1}{N} f^N(x) = \eta \quad \text{and} \quad \lim_{N \rightarrow \infty} \max_{k \in [n]} \frac{1}{N} f_k^N(x) = \max_{k \in [n]} \eta_k.$$

Lemma 10 and Theorem 11 show that every Shapley operator F has an invariant half-line (v, η) . The vectors x such that $x \leq F(x)$ may be thought of as nonlinear analogues of subharmonic functions. Akian, Gaubert, and Guterman showed that the existence of such a nontrivial vector x is equivalent to the property that the mean payoff game has at least one winning state, in the case of deterministic games [AGG12]. This is based in particular on a nonlinear fixed point theorem (Collatz–Wielandt theorem), building on earlier work of Nussabum [Nus86, Theorem 3.1].

Theorem 13 (Collatz–Wielandt properties, [AGG12, Lemma 2.8]). *Let $f: \mathbb{T}^n \rightarrow \mathbb{T}^n$ be an order-preserving, additively homogeneous, and continuous map. Let $x \in \mathbb{R}^n$ be any point. Then*

$$(10) \quad \lim_{N \rightarrow \infty} \max_{k \in [n]} \frac{1}{N} f_k^N(x) = \max\{\lambda \in \mathbb{T} : \exists u \in \mathbb{T}^n, u \neq -\infty, f(u) \geq \lambda + u\},$$

$$(11) \quad = \inf\{\lambda \in \mathbb{R} : \exists u \in \mathbb{R}^n, f(u) \leq \lambda + u\}$$

where we use the notation $\mathbb{T}^n \ni -\infty = (-\infty, -\infty, \dots, -\infty)$.

Note that this theorem is stated in [AGG12] with “sup” instead of “max” on the right hand side of (10), but their proof shows that this supremum is attained. The infimum in (11) is not attained in general.

The limit η from Corollary 12 is closely related to the value χ of the corresponding mean payoff game. More precisely, we have $\eta = 2\chi$. In order to prove this statement, we need the following notation: for every fixed policy σ of Player Min, we denote by F^σ the Shapley operator of a 1-player game in which Player Min can only use σ . Analogously, for every fixed policy τ of Player Max, we denote by F^τ the Shapley operator of a 1-player game in which Player Min can only use τ . By $F^{\sigma, \tau}$ we denote the Shapley operator of a 0-player game in which Player Min can only use σ and Player Max can only use τ . In other words, these operators are defined as follows:

$$\begin{aligned} F^\sigma(x)_k &= r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{j \in \sigma(k)} \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l), \\ F^\tau(x)_k &= \min_{a \in A^{(k)}} \left(r_k^a + \frac{1}{|a|} \sum_{\substack{j \in a \\ \tau(j) = \{l\}}} (r_j^{\tau(j)} + x_l) \right), \\ F^{\sigma, \tau}(x)_k &= r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{\substack{j \in \sigma(k) \\ \tau(j) = \{l\}}} (r_j^{\tau(j)} + x_l). \end{aligned}$$

It is easy to verify that we have the following selection lemma.

Lemma 14. *For every $x \in \mathbb{R}^n$ we have*

$$F(x) = \min_{\sigma} F^\sigma(x) = \max_{\tau} F^\tau(x),$$

where min and max denote the minimum and maximum in the partial order \leq on \mathbb{R}^n , the minimum is taken over the set of policies of Player Min, and the maximum is taken over the set of policies of Player Max. Similarly, $F^\sigma(x) = \max_{\tau} F^{\sigma, \tau}(x)$ and $F^\tau(x) = \min_{\sigma} F^{\sigma, \tau}(x)$.

Moreover, an easy induction shows the following result:

Lemma 15. *Suppose that Player Min uses a policy σ and that Player Max uses a policy τ . Then, for every initial state $k = k_1 \in [n]$ the expected total payoff of Player Max after the N th turn is equal to $((F^{\sigma, \tau})^N(0))_k$. In other words, we have*

$$\forall N, \mathbb{E}_{\sigma, \tau} \left(\sum_{p=1}^N (r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\tau(i_p)}) \right) = ((F^{\sigma, \tau} \circ F^{\sigma, \tau} \circ \dots \circ F^{\sigma, \tau})(0))_k = ((F^{\sigma, \tau})^N(0))_k.$$

Proof. Let us denote the expected total payoff of Player Max after the N th turn by $f(N, k)$, where $k \in [n]$ denotes the initial state. Observe that we have the relations

$$f(1, k) = r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{\substack{j \in \sigma(k) \\ \tau(j) = \{l\}}} r_j^{\tau(j)} = F^{\sigma, \tau}(0)_k$$

and

$$f(N+1, k) = r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{\substack{j \in \sigma(k) \\ \tau(j) = \{l\}}} (r_j^{\tau(j)} + f(N, l)).$$

This proves by induction that $f(N, k) = ((F^{\sigma, \tau})^N(0))_k$. \square

Lemma 16. *There exists a pair of policies $(\bar{\sigma}, \bar{\tau})$ such that F , $F^{\bar{\sigma}}$, and $F^{\bar{\tau}}$ have the same invariant half-line (v, η) .*

Proof. Let (v, η) denote any invariant half-line of F . We start by showing the existence of $\bar{\sigma}$. By Lemma 14 we have $F(v + \gamma\eta) = \min_{\sigma} F^{\sigma}(v + \gamma\eta)$. Therefore, there exists a strategy $\bar{\sigma}$ such that the equality $v + (N+1)\eta = F^{\bar{\sigma}}(v + N\eta)$ is verified for infinitely many values of $N \in \mathbb{N}^*$. Since the operator $F^{\bar{\sigma}}$ is piecewise-affine, there exists an affine function $Ax + b$ such that $F^{\bar{\sigma}}(v + \gamma\eta) = A(v + \gamma\eta) + b$ for all $\gamma \geq 0$ large enough. In particular, $v + \eta + N\eta = Av + b + NA\eta$ for infinitely many values of $N \in \mathbb{N}^*$. Hence $A\eta = \eta$ and $v + \eta = Av + b$. Therefore $F^{\bar{\sigma}}(v + \gamma\eta) = A(v + \gamma\eta) + b = v + (\gamma+1)\eta$ for all $\gamma \geq 0$ large enough. The construction of $\bar{\tau}$ is analogous. \square

Theorem 17. *Policies $(\bar{\sigma}, \bar{\tau})$ are optimal. Furthermore, the value of the game is equal to $\eta/2$.*

Proof. Suppose that Player Min uses $\bar{\sigma}$. Let τ be any policy of Player Min. By Lemma 15 we have

$$\forall \tau, g_k(\bar{\sigma}, \tau) = \lim_{N \rightarrow \infty} \mathbb{E}_{\bar{\sigma}, \tau} \left(\frac{1}{2N} \sum_{p=1}^N (r_{k_p}^{\bar{\sigma}(k_p)} + r_{i_p}^{\tau(i_p)}) \right) = \lim_{N \rightarrow \infty} \frac{(F^{\bar{\sigma}, \tau})^N(0)}{2N}.$$

Therefore, Lemma 14, Lemma 16, and Corollary 12 show that

$$\forall \tau, g_k(\bar{\sigma}, \tau) \leq \lim_{N \rightarrow \infty} \frac{(F^{\bar{\sigma}})^N(0)}{2N} = \frac{1}{2} \eta_k.$$

Analogously, if Player Max uses $\bar{\tau}$, we have

$$\forall \sigma, g_k(\sigma, \bar{\tau}) = \lim_{N \rightarrow \infty} \mathbb{E}_{\sigma, \bar{\tau}} \left(\frac{1}{2N} \sum_{p=1}^N (r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\bar{\tau}(i_p)}) \right) \geq \frac{1}{2} \eta_k.$$

Hence $(\bar{\sigma}, \bar{\tau})$ are optimal and $g_k(\bar{\sigma}, \bar{\tau}) = \eta_k/2$ for every $k \in [n]$. \square

Theorem 18. *Player Max has at least one winning initial state if and only if the set $\{x \in \mathbb{T}^n : x \leq F(x)\}$ is nontrivial (i.e., contains a point different than $(-\infty, \dots, -\infty)$).*

Proof. By Theorem 17 and Corollary 12 we have

$$\max_{k \in [n]} 2\chi_k = \lim_{N \rightarrow \infty} \max_{k \in [n]} \frac{1}{N} F_k^N(x).$$

Hence, by (10), we have

$$\max_{k \in [n]} 2\chi_k = \max\{\lambda \in \mathbb{T} : \exists x \in \mathbb{T}^n, x \neq -\infty, \lambda + x \leq F(x)\}.$$

In particular, the inequality $\max_{k \in [n]} \chi_k \geq 0$ holds if and only if the set $\{x \in \mathbb{T}^n : x \leq F(x)\}$ is nontrivial. \square

The Shapley operator F may be thought of as a non-linear Markov operator, and so, the vectors x such that $x \leq F(x)$ may be thought of as non-linear *sub-harmonic* vectors. Theorem 18 shows that the existence of a winning state is characterized by the existence of a non-trivial sub-harmonic vector.

5. EQUIVALENCE BETWEEN STOCHASTIC GAMES AND TROPICAL SPECTRAHEDRA

5.1. The stochastic game associated with a Metzler tropical spectrahedron. We now describe the connection between tropical spectrahedra and stochastic mean payoff games. Let us consider a tropical Metzler spectrahedron \mathcal{S} associated with the matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$. We construct a stochastic mean payoff game Γ consisting of m states of Player Max and n states of Player Min. For each state $k \in [n]$ of Player Min, the set $A^{(k)}$ of actions available to Player Min at this state consists of:

- the actions $\{i\}$ with payment $-|Q_{ii}^{(k)}|$ for all $i \in [m]$ such that $Q_{ii}^{(k)} \in \mathbb{T}_-$;
- the actions $\{i, j\}$ with payment $-|Q_{ij}^{(k)}|$ for all $i < j$ such that $Q_{ij}^{(k)} \in \mathbb{T}_-$.

For every state $i \in [m]$ of Player Max, the set $B^{(i)}$ of actions available to Player Max at this state is formed by the actions $\{k\}$ with payment $Q_{ii}^{(k)}$ for all states $k \in [n]$ satisfying $Q_{ii}^{(k)} \in \mathbb{T}_+$.

Recall that in the games which we consider, every state has to be equipped with at least one action, i.e., the sets $A^{(k)}$ and $B^{(i)}$ must be nonempty. In consequence, our construction is valid provided that the following assumption on the matrices $Q^{(k)}$ is satisfied:

Assumption 19. (a) For all $k \in [n]$, the matrix $Q^{(k)}$ has at least one coefficient in \mathbb{T}_- .
(b) For all $i \in [m]$, there exists $k \in [n]$ such that the diagonal coefficient $Q_{ii}^{(k)}$ belongs to \mathbb{T}_+ .

Concerning the nontriviality of tropical Metzler spectrahedra, Assumption 19 can be made without loss of generality, up to extracting submatrices from $Q^{(k)}$ or eliminating some of them, as shown by the next lemma.

Lemma 20. Let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ be symmetric Metzler matrices, and \mathcal{S} be the associated tropical spectrahedron. We can build (in poly-time) symmetric Metzler matrices $R^{(1)}, \dots, R^{(q)} \in \mathbb{T}_{\pm}^{p \times p}$ ($p \leq m$, $q \leq n$) satisfying Assumption 19 such that the associated tropical spectrahedron is nontrivial if and only if \mathcal{S} is nontrivial.

Proof. We first examine Assumption 19(a). Suppose that the matrix $Q^{(k)}$ has no coefficient in \mathbb{T}_- . In this case, the spectrahedron \mathcal{S} is nontrivial, since it contains the vector $x \in \mathbb{T}^n$ such that $x_k = 0$ and $x_l = -\infty$ for all $l \neq k$.

Now, let us look at Assumption 19(b). Consider $i \in [m]$ and suppose that no diagonal coefficient $Q_{ii}^{(k)}$ is in \mathbb{T}_+ . We distinguish three cases:

- if the set K of indices $k \in [n]$ such that $Q_{ii}^{(k)} \in \mathbb{T}_-$ is nonempty, the relation $Q_{ii}^+(x) \geq Q_{ii}^-(x)$ enforces to have $x_k = -\infty$ for all $x \in \mathcal{S}$ and $k \in K$. This means that we can reduce the nontriviality of \mathcal{S} to the nontriviality of the spectrahedron associated with the matrices $Q^{(k)}$ with $k \notin K$;
- if the aforementioned set K is empty, and the i -th row and column of the matrices $Q^{(k)}$ are all identically equal to $-\infty$, then we can remove all these rows and columns, and reduce to a problem with matrices of order $m - 1$ over the variables x_1, \dots, x_n ;
- if the set K is empty and some matrix $Q^{(k)}$ contains an entry different than $-\infty$ on its i -th row, namely $Q_{ij}^{(k)} \in \mathbb{T}_-$ with $i < j$, then the relation $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq (Q_{ij}(x))^{\odot 2}$ enforces $x_k = -\infty$ for any $x \in \mathcal{S}$. We consequently reduce the problem to the nontriviality of the spectrahedron associated with the matrices $Q^{(l)}$ with $l \neq k$. \square

We are going to show that the value of the game Γ is related with the feasibility of sets \mathcal{S}_λ obtained from \mathcal{S} by adding a reinforcing coefficient to the inequalities that correspond to minors of order 1 and 2.

Definition 21. For any $\lambda \in \mathbb{R}$ we denote by \mathcal{S}_λ the set of all points $x \in \mathbb{T}^n$ verifying

- for all $i \in [m]$, $Q_{ii}^+(x) \geq \lambda \odot Q_{ii}^-(x)$;
- for all $i, j \in [m]$, $i < j$, $Q_{ii}^+(x) \odot Q_{jj}^+(x) \geq (\lambda \odot Q_{ij}(x))^{\odot 2}$.

Observe that we have $\mathcal{S}_0 = \mathcal{S}$.

Lemma 22. *Let $F: \mathbb{T}^n \rightarrow \mathbb{T}^n$ be the Shapley operator associated with Γ . Then, for any $\lambda \in \mathbb{R}$ we have $\mathcal{S}_\lambda = \{x \in \mathbb{T}^n : \lambda + x \leq F(x)\}$.*

Proof. By construction of the game Γ , a vector $x \in \mathbb{T}^n$ satisfies $\lambda + x \leq F(x)$ if and only if for all $k \in [n]$,

$$\lambda + x_k \leq \min_{Q_{ij}^{(k)} \in \mathbb{T}_-} \left(-|Q_{ij}^{(k)}| + \frac{1}{2} \left(\max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l) + \max_{Q_{jj}^{(l)} \in \mathbb{T}_+} (Q_{jj}^{(l)} + x_l) \right) \right),$$

or, equivalently, for all $i, j \in [m]$,

$$2 \left(\lambda + \max_{Q_{ij}^{(k)} \in \mathbb{T}_-} (|Q_{ij}^{(k)}| + x_k) \right) \leq \max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l) + \max_{Q_{jj}^{(l)} \in \mathbb{T}_+} (Q_{jj}^{(l)} + x_l).$$

By distinguishing whether i and j are equal in these inequalities, and recalling that $Q_{ij}^{(k)}$ is in $\mathbb{T}_- \cup \{-\infty\}$ for all $k \in [n]$ when $i \neq j$, we recover the constraints that describe \mathcal{S}_λ . \square

From (10) and Lemma 22, we obtain:

Theorem 23. *The set \mathcal{S}_λ is nontrivial if and only if $\lambda \geq 2 \max_{k \in [n]} \chi_k$, where χ is the value of the game Γ .*

In particular, the tropical spectrahedron \mathcal{S} is nontrivial if and only if the stochastic game Γ has at least one winning initial state.

In case at least one initial state of the game has a positive value, the latter statement can be refined in order to obtain an explicit point in the spectrahedron \mathcal{S} .

Lemma 24. *Suppose that $x \in \mathcal{S}_\lambda$ for some positive λ . Let $Q^{(1)}, \dots, Q^{(n)}$ denote any symmetric matrices in $\mathbb{K}^{m \times m}$ such that $\text{sval}(Q^{(k)}) = Q^{(k)}$ for all $k \in [n]$. Then, the point $x = (t^{x_1}, \dots, t^{x_n})$ belongs to $\mathcal{S} := \{x \in \mathbb{K}_{\geq 0}^n : x_1 Q^{(1)} + \dots + x_n Q^{(n)} \succcurlyeq 0\}$.*

Proof. The proof of [AGS16b, Lemma 26] shows that if $x \in \mathcal{S}_\lambda$ for some positive λ , then any point $x \in \mathbb{K}_{\geq 0}^n$ verifying $\text{val}(x) = x$ belongs to \mathcal{S} . \square

Remark 25. In case all the initial states of the game have a negative value, we know from the second Collatz–Wielandt identity (11) that there exists a vector $u \in \mathbb{R}^n$ and a scalar $\lambda < 0$ such that $F(u) \leq \lambda + u$. The pair (u, λ) yields an emptiness certificate for the nonarchimedean spectrahedron \mathcal{S} .

Along the same lines, we can reciprocally associate a tropical spectrahedron with any stochastic game. In more details, let Γ be a stochastic mean payoff game over the sets of states $[m]$ and $[n]$. We recall that $A^{(k)}$ and $B^{(i)}$ correspond to the sets of possible actions of Player Min at state $k \in [n]$ and Player Max at $i \in [m]$ respectively. Furthermore, r_k^a is the reward obtained by Player Max when Player Min chooses an action $a \in A^{(k)}$ at state $k \in [n]$ and r_i^b is the reward obtained by Player Max when he chooses an action $b \in B^{(i)}$ at state $i \in [m]$. Out of this data we construct symmetric Metzler matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$. We define $Q_{ij}^{(k)} = Q_{ji}^{(k)} := \ominus(-r_k^a)$ for all $k \in [n]$ and $i, j \in [m]$ such that $i \neq j$ and $a = \{i, j\}$ is an available action of Player Min at the state k . Diagonal coefficients of the matrices $Q^{(k)}$ are defined according to whether the actions $a := \{i\}$ and $b := \{k\}$ are available from the states k and i respectively:

- if $a \in A^{(k)}$ and $b \notin B^{(i)}$, we set $Q_{ii}^{(k)} := \ominus(-r_k^a)$, while if $b \in B^{(i)}$ and $a \notin A^{(k)}$, we define $Q_{ii}^{(k)} := r_i^b$;
- if both actions a and b occur simultaneously, we set $Q_{ii}^{(k)}$ to $\ominus(-r_k^a)$ if $(-r_k^a) > r_i^b$, and to r_i^b if $r_i^b \geq (-r_k^a)$.

Finally, all the other entries of the matrices $Q^{(k)}$ are set to $-\infty$. With this construction, it can be verified that Lemma 22 and, subsequently, Theorem 23 are still valid. For the sake of readability, we prove this in Appendix A.

We denote by TMSDFP the tropical Metzler semidefinite feasibility problem: “given symmetric tropical Metzler matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ satisfying Assumption 19, is the associated tropical Metzler spectrahedron trivial?” As a consequence of Theorem 23, we obtain the equivalence between the two decision problems SMPG and TMSDFP. This equivalence can be refined thanks to the fact that SMPG restricted to games with payments equal to 0 or ± 1 is poly-time equivalent to SMPG for arbitrary payments. We show the proof of this reduction in Section 7. This yields the following result:

Theorem 26. *The problems SMPG and TMSDFP are poly-time equivalent. Furthermore, if either of these problems can be solved in pseudopolynomial time, then both of them can be solved in polynomial time.*

Proof. The fact that TMSDFP is poly-time equivalent to SMPG follows directly from Theorem 23. Moreover, the same statement shows that SMPG restricted to games with payoffs in $\{-1, 0, 1\}$ is poly-time equivalent to TMSDFP restricted to matrices with entries in

$$\{0, \ominus 0, \pm 1, \ominus(\pm 1), -\infty\}.$$

By using Corollary 51 (Section 7), we deduce that if either of these problems can be solved in pseudopolynomial time, then both are solvable in polynomial time. \square

We now extend Theorem 26 in two different directions. In Section 5.2, we deal with the problem of checking whether a certain stratum of a tropical Metzler spectrahedron is empty. This allows us to handle the case of affine spectrahedra. Section 5.3 provides a result on the asymptotic feasibility of real spectrahedra arising from a spectrahedron over Puiseux series.

5.2. Strata of tropical spectrahedra and dominions. We want to characterize the possible supports of points belonging to $\{x \in \mathbb{T}^n : x \leq F(x)\}$. These supports can be interpreted in terms of the associated stochastic mean payoff game, using the following notion. A dominion of Player Max is a subset of initial states $D \subset [n]$ such that, if the game starts in D , then Player Max can ensure that the game never reaches any state from $[n] \setminus D$. Formally, we make the following definition.

Definition 27. We say that a subset $D \subset [n]$ is a *dominion (of Player Max)* if for every state $k \in D$ and every action $\{i, j\} \in A^{(k)}$, there exists a pair of states $l_1 \in D, l_2 \in D$ such that $\{l_1\} \in B^{(i)}$ and $\{l_2\} \in B^{(j)}$.

We refer to Figure 3 for an illustration. Given a dominion $D \subset [n]$, we can construct a *subgame induced by D* . This game is created as follows: the set of states controlled by Player Min is equal to D . Each of these states is equipped with the same set of actions as in the original game. The set of states controlled by Player Max consists of all states $i \in [m]$ such that there exists a state $k \in D$ and a state $j \in [m]$ verifying $\{i, j\} \in A^{(k)}$. Every such state i is equipped with all the actions of the form $\{l\} \in B^{(i)}$ with $l \in D$. The payoffs associated with the actions in the induced subgame are the same as payoffs in the original game.

Definition 28. We say that a dominion $D \subset [n]$ is a *winning dominion* if all states of D are winning in the subgame induced by D .

Remark 29. It follows from the definition that if a state $k \in [n]$ belongs to a winning dominion, then it is a winning state. The converse is not true, even if we suppose that k belongs to a dominion that contains only winning states. An example of such situation is presented in Figure 3.

We can now give a characterization of all possible supports of points $\{x \in \mathbb{T}^n : x \leq F(x)\}$.

Theorem 30. *Take a Shapley operator $F : \mathbb{T}^n \rightarrow \mathbb{T}^n$ and a nonempty set $K \subset [n]$. Then, the set $\{x \in \mathbb{T}^n : x \leq F(x)\}$ contains a point with support K if and only if K is a winning dominion in the stochastic mean payoff game associated with F .*




FIGURE 3. A mean payoff games illustrating the notion of winning dominions. The states ①, ②, and ③ are winning. The minimal (inclusion-wise) dominions are given by $\{\textcircled{1}\}$, $\{\textcircled{3}\}$, $\{\textcircled{4}\}$, and $\{\textcircled{2}, \textcircled{3}, \textcircled{4}\}$. The state ① is winning and constitutes a dominion, because Player Max can force to return to ① by playing $\overline{1} \rightarrow \textcircled{1}$. However, this state does not belong to any winning dominion, because the induced subgame, reduced to the states ① and $\overline{1}$, is not winning for Player Max. The state ③ is the only winning dominion.

Proof. Suppose that we are given a point $x \in \mathbb{T}^n$ with support K . Then, for every $i \in [m]$ we have

$$(12) \quad \max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) = \max_{\substack{b \in B^{(i)} \\ b = \{l\}, l \in K}} (r_i^b + x_l).$$

Suppose that $x \leq F(x)$. If K is not a dominion, then there exists a state $k \in K$ and an action $\{i, j\} \in A^{(k)}$ such that we have $l \notin K$ for all $\{l\} \in B^{(i)}$. Hence, we have

$$x_k \leq (F(x))_k = \min_{\substack{a \in A^{(k)} \\ a = \{i, j\}}} \left(r_k^a + \frac{1}{2} \left(-\infty + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right) \right) = -\infty,$$

what gives a contradiction. Therefore, K is a dominion. Let $\tilde{x} \in \mathbb{R}^{|K|}$ denote the stratum of x associated with K . Observe that (12) shows that if $F^K: \mathbb{T}^{|K|} \rightarrow \mathbb{T}^{|K|}$ denotes the Shapley operator of the subgame induced by K , then we have $(F(x))_k = (F^K(\tilde{x}))_k$ for every $k \in K$. Let $(v, \eta) \in \mathbb{R}^{|K|} \times \mathbb{R}^{|K|}$ denote an invariant half-line of $F_{|\mathbb{R}^{|K|}}^K$. Suppose that $\eta_k < 0$ for some $k \in K$. Fix a natural number $N \geq 1$. We have

$$x_k = \tilde{x}_k \leq (F^K(\tilde{x}))_k \leq ((F^K \circ F^K)(\tilde{x}))_k \leq \dots \leq ((F^K \circ \dots \circ F^K)(\tilde{x}))_k,$$

where F^K is composed N times. Hence $x_n/N \leq ((F^K \circ \dots \circ F^K)(\tilde{x}))_k/N$. By Corollary 12, the right-hand side of this inequality converges to $\eta_k < 0$. This implies that $x_k = -\infty$, what gives a contradiction. Hence $\eta_k \geq 0$ for every $k \in K$ and K is a winning dominion by Theorem 17.

Conversely, suppose that K is a winning dominion. Let F^K denote the Shapley operator of the subgame induced by K and let (v, η) be the invariant half-line of $F_{|\mathbb{R}^{|K|}}^K$. Theorem 17 shows that $\eta_k \geq 0$ for all $k \in K$. By the definition, for $\gamma \geq 0$ large enough we have $F^K(v + \gamma\eta) = v + (\gamma + 1)\eta \geq v + \gamma\eta$. Take any such γ and denote $\tilde{x} = v + \gamma\eta$. We extend the point $\tilde{x} \in \mathbb{R}^{|K|}$ to $x \in \mathbb{T}^n$ by setting $x_k = \tilde{x}_k$ for $k \in K$ and $x_k = -\infty$ otherwise. As previously, (12) shows that we have $(F(x))_k = (F^K(\tilde{x}))_k$ for every $k \in K$. In particular, $x_k = \tilde{x}_k \leq (F^K(\tilde{x}))_k = (F(x))_k$ for all $k \in K$. Finally, for every $k \notin K$ we have $x_k = -\infty \leq (F(x))_k$. \square

The following is an immediate corollary.

Corollary 31. *Take a Shapley operator $F: \mathbb{T}^n \rightarrow \mathbb{T}^n$ such that all the initial states $i \in [n]$ of the associated stochastic mean payoff game have the same mean payoff. Then, the set $\{x \in \mathbb{T}^n: x \leq F(x)\}$ is nontrivial if and only if it contains a finite vector.*

One may reinforce the condition of this corollary by requiring that the mean payoff of the game remains independent of the initial state for all numerical values of the payments of the game, the transitions being unchanged. The latter property admits a convenient combinatorial characterization, which relies on the *recession function* \hat{F} of the Shapley operator:

$$\hat{F}(x) := \lim_{\gamma \rightarrow \infty} \gamma^{-1} F(\gamma x) .$$

Since F commutes with the addition of a constant vector, it is immediate that $\hat{F}(\alpha, \dots, \alpha) = (\alpha, \dots, \alpha)$ holds for all $\alpha \in \mathbb{R}$. We call such fixed point of \hat{F} *uniform*.

Theorem 32 ([AGH15, Th. 3.1]). *Suppose that the recession function \hat{F} has only uniform fixed points. Then, there exist $u \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$ such that $F(u) = \lambda + u$. In particular, the mean payoff of the game is equal to $\lambda/2$ for all initial states.*

The condition that F has only uniform fixed points can be checked by finding invariant sets in directed hypergraphs [AGH15]. For concrete examples of games arising from spectrahedra, we shall see that this leads to explicit conditions on the zero/non-zero pattern of the matrices defining the spectrahedron (Remark 44).

Let us now suppose that $Q^{(0)}, \dots, Q^{(n)} \in \mathbb{K}^{m \times m}$ are Metzler matrices, and let $Q^{(k)} := \text{sval}(Q^{(k)})$. By [AGS16b, Lemma 20] and Theorem 3, provided that the matrices $Q^{(k)}$ are generic, the affine spectrahedron

$$(13) \quad \mathcal{S} := \{x \in \mathbb{K}_{\geq 0}^n : Q^{(0)} + x_1 Q^{(1)} + \dots + x_n Q^{(n)} \succcurlyeq 0\}$$

is nonempty if and only if the tropical Metzler spectrahedron $\mathcal{S}(Q^{(0)}, \dots, Q^{(n)})$ contains a point x such that $x_0 \neq -\infty$. Let Γ be the stochastic mean payoff game associated arising from the matrices $Q^{(0)}, \dots, Q^{(n)}$. The latter property can be checked using dominions as follows:

Corollary 33. *The set $\{x \in \mathcal{S}(Q^{(0)}, \dots, Q^{(n)}) : x_0 \neq -\infty\}$ is nonempty if and only if there is a winning dominion D in Γ such that $0 \in D$.*

5.3. The archimedean feasibility problem. We next relate the tropical feasibility problem with the archimedean feasibility problem. For simplicity of exposition, we consider the case of conic spectrahedra.

We suppose that $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{K}^{m \times m}$ are symmetric Metzler matrices, set $Q(x) := x_1 Q^{(1)} + \dots + x_n Q^{(n)}$, and consider \mathcal{S} as in (13) with $Q^{(0)} = 0$. For any fixed value of the parameter $t \in \mathbb{R}$, we also consider the real spectrahedron $\mathcal{S}(t) \subset \mathbb{R}_{\geq 0}^n$ described by $Q^{(1)}(t), \dots, Q^{(n)}(t)$ in a similar manner. We want to study the feasibility problem of $\mathcal{S}(t)$ as t goes to infinity. We denote $\text{sval}(Q^{(k)}) = Q^{(k)}$ for all k and we further suppose that the matrices $Q^{(1)}, \dots, Q^{(n)}$ satisfy Assumption 19. The proof is based on the following definition and lemma.

Definition 34. For any $\alpha \geq 1$ we define the set $\mathcal{S}_{2,\alpha} \subset \mathbb{K}_{\geq 0}^n$ as

$$\mathcal{S}_{2,\alpha} := \left\{ x \in \mathbb{K}_{\geq 0}^n : \forall i, Q_{ii}(x) \geq 0, \forall i \neq j, Q_{ii}(x)Q_{jj}(x) \geq \alpha(Q_{ij}(x))^2 \right\}.$$

Lemma 35. *We have the inclusion $\mathcal{S}_{2,(m-1)^2} \subset \mathcal{S} \subset \mathcal{S}_{2,1}$.*

We refer to [AGS16b, Section 5.1] for the proof of Lemma 35. We also introduce a threshold $T > 1$ such that for all $t \geq T$, every series $Q_{ij}^{(k)}(t)$ converges, and the signs of $Q_{ij}^{(k)}(t)$ and $Q_{ij}^{(k)}$ are the same. For any $t \geq T$, we define

$$\delta(t) := \max_{Q_{ij}^{(k)} \neq -\infty} \left| |Q_{ij}^{(k)}| - \log_t |Q_{ij}^{(k)}(t)| \right|.$$

By the definition of valuation and order in \mathbb{K} we have $\lim_{t \rightarrow \infty} \delta(t) = 0$. Furthermore, for all k, i, j we have $t^{|Q_{ij}^{(k)}| - \delta(t)} \leq |Q_{ij}^{(k)}(t)| \leq t^{|Q_{ij}^{(k)}| + \delta(t)}$, even if $Q_{ij}^{(k)} = -\infty$ (we set $t^{-\infty} = 0$). The next theorem relates the feasibility of the spectrahedron $\mathcal{S}(t)$ with the value of the mean payoff game associated with $Q^{(1)}, \dots, Q^{(n)}$ when t is sufficiently large. We point out that this result does not require the genericity assumptions of Theorem 3.

Theorem 36. Let $m \geq 2$, and χ be the value of the stochastic mean payoff game associated with $Q^{(1)}, \dots, Q^{(n)}$. Let $\lambda := \max_k \chi_k$, and suppose that $\lambda \neq 0$. Take any $t \geq T$ such that $\delta(t) < |\lambda|$ and

$$t > (2(m-1)n)^{1/(2|\lambda|-2\delta(t))}.$$

Then, the spectrahedron $\mathcal{S}(t)$ is nontrivial if and only if λ is positive.

Proof. Suppose that λ is positive. Theorem 23 shows that there is a point $x \in \mathbb{T}^n$, $x \neq -\infty$ such that

$$\forall (i, j) \in [m]^2, \quad 2\lambda + Q_{ij}^-(x) \leq \frac{1}{2}Q_{ii}^+(x) + \frac{1}{2}Q_{jj}^+(x).$$

Thus, we have $t^{4\lambda}t^{2Q_{ij}^-(x)} \leq t^{Q_{ii}^+(x)}t^{Q_{jj}^+(x)}$. Take the point $\mathbf{x} = (t^{x_1}, \dots, t^{x_n})$, where $t^{-\infty} = 0$. Observe that we have

$$\forall (i, j) \in [m]^2, \quad Q_{ij}^-(t)(\mathbf{x}) \leq \sum_{Q_{ij}^{(k)} \in \mathbb{T}_-} t^{|Q_{ij}^{(k)}| + \delta(t) + x_k} \leq nt^{Q_{ij}^-(x) + \delta(t)}$$

and

$$\forall i \in [m], \quad Q_{ii}^+(t)(\mathbf{x}) \geq \sum_{Q_{ii}^{(k)} \in \mathbb{T}_+} t^{Q_{ii}^{(k)} - \delta(t) + x_k} \geq t^{Q_{ii}^+(x) - \delta(t)}.$$

Therefore, for all $i \in [m]$ we have

$$Q_{ii}^+(t)(\mathbf{x}) \geq \frac{t^{2\lambda - 2\delta(t)}}{n} Q_{ii}^-(t)(\mathbf{x}).$$

Since $t \geq (2n)^{1/(2\lambda - 2\delta(t))}$, we have $Q_{ii}^+(t)(\mathbf{x}) \geq 2Q_{ii}^-(t)(\mathbf{x})$ and hence $Q_{ii}(\mathbf{x}) \geq \frac{1}{2}Q_{ii}^+(\mathbf{x})$. Thus, for any $i < j$ we have

$$\begin{aligned} Q_{ii}(\mathbf{x})Q_{jj}(\mathbf{x}) &\geq \frac{1}{4}Q_{ii}^+(\mathbf{x})Q_{jj}^+(\mathbf{x}) \geq \frac{1}{4}t^{Q_{ii}^+(x) + Q_{jj}^+(x) - 2\delta(t)} \\ &\geq \frac{1}{4}t^{2|Q_{ij}^-(x)| + 4\lambda - 2\delta(t)} \geq \frac{t^{4\lambda - 4\delta(t)}}{4n^2} (Q_{ij}(\mathbf{x}))^2. \end{aligned}$$

Hence, since $t \geq (2(m-1)n)^{1/(2\lambda - 2\delta(t))}$, we get $(Q_{ii}(t)(\mathbf{x}))(Q_{jj}(t)(\mathbf{x})) \geq (m-1)^2(Q_{ij}(t)(\mathbf{x}))^2$ for all $i < j$ and the point \mathbf{x} belongs to $\mathcal{S}(t)$ by Lemma 35.

Conversely, suppose that λ is negative but $\mathcal{S}(t)$ is nontrivial. Take any nonzero point $\mathbf{x} \in \mathcal{S}(t)$ and a point $x \in \mathbb{T}^n$ defined as $x_k = \log_t(\mathbf{x}_k)$ for all $k \in [n]$ (where $\log_t(0) = -\infty$). Since λ is negative, Theorem 23 shows that for every fixed $\varepsilon > 0$ there is a pair $(i, j) \in [m]^2$ such that

$$2\lambda + \varepsilon + Q_{ij}^-(x) > \frac{1}{2}Q_{ii}^+(x) + \frac{1}{2}Q_{jj}^+(x).$$

Hence $t^{4\lambda + 2\varepsilon}t^{2Q_{ij}^-(x)} > t^{Q_{ii}^+(x)}t^{Q_{jj}^+(x)}$. Similarly to the previous case, observe that we have

$$Q_{ij}^-(t)(\mathbf{x}) \geq \sum_{Q_{ij}^{(k)} \in \mathbb{T}_-} t^{Q_{ij}^{(k)} - \delta(t) + x_k} \geq t^{Q_{ij}^-(x) - \delta(t)}$$

and

$$Q_{ii}^+(t)(\mathbf{x}) \leq \sum_{Q_{ii}^{(k)} \in \mathbb{T}_+} t^{Q_{ii}^{(k)} + \delta(t) + x_k} \leq nt^{Q_{ii}^+(x) + \delta(t)}.$$

Therefore, we have

$$t^{4\lambda + 2\varepsilon + 2\delta(t)}(Q_{ij}^-(t)(\mathbf{x}))^2 > t^{Q_{ii}^+(x) + Q_{jj}^+(x)} \geq \frac{t^{-2\delta(t)}}{n^2} (Q_{ii}^+(t)(\mathbf{x}))(Q_{jj}^+(t)(\mathbf{x})).$$

Thus, if we take ε such that $2\lambda + \varepsilon + 2\delta(t) < 0$ and $t \geq n^{1/|2\lambda + \varepsilon + 2\delta(t)|}$, we have $(Q_{ij}^-(t)(\mathbf{x}))^2 > (Q_{ii}^+(t)(\mathbf{x}))(Q_{jj}^+(t)(\mathbf{x}))$, what gives a contradiction. \square

Remark 37. It is easy to see from the proof that if the matrices $Q^{(k)}$ are diagonal (which holds, in particular, if $m = 1$), then the term $2(m - 1)$ is not needed, and the bound for t takes the form $t > n^{1/(2|\lambda| - 2\delta(t))}$.

Example 38. We can apply this bound to the example presented in Section 2. Let us take the simplest lift of matrices $Q^{(1)}, Q^{(2)}, Q^{(3)}$, namely $Q_{ij}^{(k)} = \text{sign}(Q_{ij}^{(k)})t^{Q_{ij}^{(k)}}$. In this case we have $\delta(t) = 0$ for all $t > 1$. Moreover, the computations presented in Example 9 show that $\lambda = 1/56$. Thus, if we take $t > 12^{28}$, then the spectrahedron $\mathcal{S}(t)$ is nontrivial. In this special instance, the lower bound provided by Theorem 36 is large because we chose a game that is nearly singular: the mean payoff λ is close to 0. There are, however, other instances of nonarchimedean spectrahedra for which λ is of order one, leading to a $(mn)^{O(1)}$ bound for t .

6. ALGORITHMS

In this section, we discuss algorithms that can solve TMSDFP thanks to the equivalence with SMPG established in Theorem 26.

6.1. Complexity bounds. We first derive complexity bounds for TMSDFP. Let L denote the maximal number of bits needed to encode an entry from matrices $Q^{(1)}, \dots, Q^{(n)}$. An exponential bound for TMSDFP is achieved by a naive algorithm that enumerates all policies of one player and uses linear programming to solve the remaining 1-player game.

Theorem 39. *There is an algorithm that solves TMSDFP in*

$$\min\{n^m, m^{2n}\}\text{Poly}(m, n, L)$$

arithmetic operations.

Proof. Take the matrices $Q^{(1)}, \dots, Q^{(n)}$ and consider the associated stochastic game Γ . If $m^{2n} \leq n^m$, then we do the following: for each policy σ of Player Min, we consider a 1-player game Γ^σ induced by fixing σ . The value of Γ^σ , denoted χ^σ , can then be found in polynomial time by linear programming (see [FV07, Section 2.9]). Moreover, for every fixed state $k \in [n]$, the value χ_k of Γ can be found by taking the minimum over σ of values χ_k^σ . Now, by Theorem 23, the tropical Metzler spectrahedron \mathcal{S} associated with $Q^{(1)}, \dots, Q^{(n)}$ is nontrivial if and only if χ has at least one nonnegative entry. Since Player Min has at most m^{2n} policies in Γ , this procedure takes at most $m^{2n}\text{Poly}(m, n, L)$ arithmetic operations. If $m^{2n} > n^m$, then we do the analogous operation for Player Max. \square

The best currently known bounds for TMSDFP can be derived from the interpretation of stochastic games as LP-type problems. These bounds are randomized subexponential, [Hal07], and also [HZ15] for a recent improvement. A distinctive feature of this approach is that it works in *strongly* subexponential time, i.e., its arithmetic complexity does not depend on L .

Theorem 40. *There is a randomized algorithm that solves TMSDFP in expected*

$$\exp(O(\sqrt{(m+n)\log(m+n)}))$$

arithmetic operations.

Proof. As in the proof of Theorem 39, take the matrices $Q^{(1)}, \dots, Q^{(n)}$ and consider the associated game Γ . Now, we apply to Γ the first step of reduction presented in [AM09, Lemma 1 and Lemma 3] (i.e., the step presented in Fig. 2 of the latter reference). This converts Γ in a strongly polynomial complexity to a game Γ_1 , which has the same optimal policies as Γ and has the form considered by Halman [Hal07]. Moreover, Γ_1 has at most n min vertices and at most m max vertices (in Halman's terminology). Furthermore, Halman [Hal07, Theorem 4.1] gives an algorithm that finds optimal policies for games belonging to his class in the expected $\exp(O(\sqrt{p\log p}))$ arithmetic operations, where p is the total number of min and max vertices. Hence, a pair of optimal policies in Γ_1 (which is also optimal in Γ) can be found in the expected

$$\exp(O(\sqrt{(m+n)\log(m+n)}))$$

```

1: procedure CHECKFEASIBILITY( $F, \varepsilon$ )
2:    $\triangleright F$  a Shapley operator from  $\mathbb{T}^n$  to  $\mathbb{T}^n$ ,  $\varepsilon > 0$  a numerical precision.
3:    $u := 0 \in \mathbb{R}^n$ ,  $v := 0 \in \mathbb{R}^n$ 
4:   while  $\max_k u_k > -\varepsilon$  and  $\min_k u_k < \varepsilon$  do
5:      $v := \max(v, u)$ ,  $u := F(u)$ 
6:      $\triangleright$  The operation  $\max$  on vectors is understood entrywise.
7:   done
8:   if  $\max_k u_k \leq -\varepsilon$  then
9:     There is no vector  $x \in \mathbb{T}^n$ ,  $x \neq -\infty$ , such that  $x \leq F(x)$ .
10:  else
11:    The vector  $v$  satisfies  $v \leq F(v)$ .
12:  end

```

FIGURE 4. Checking the feasibility of a tropical semidefinite problem by value iteration.

arithmetic operations. Once a pair of optimal policies is known, the value of Γ can be found in strongly polynomial complexity by Remark 54. As in the proof of Theorem 39, we use the value of Γ to decide whether the tropical Metzler spectrahedron \mathcal{S} is trivial. \square

6.2. Value iteration. We now present an algorithm with poorer theoretical bounds, but which is well adapted practically to some large scale instances: value iteration. This algorithm may be thought of as a nonlinear analogue of the power algorithm to compute the dominant eigenvalue of a matrix. Its advantage lies in scalability. This leads to a procedure called CHECKFEASIBILITY, provided in Figure 4 and which checks the existence of a sub-harmonic vector of a Shapley operator of the form (9). By Lemma 22, this is equivalent to finding a point in a tropical spectrahedron.

We next establish the correctness of the CHECKFEASIBILITY algorithm, under the assumption of Corollary 31, that all the initial states in the mean payoff game associated with the Shapley operator F have the same value λ . It follows from Section 7 (Corollary 49) that every stochastic mean payoff game reduces polynomially to a game with this property.

Theorem 41. *Suppose that all the initial states of the game with Shapley operator F have the same value. Suppose in addition that this value, λ , differs from 0. Then, for all choices of $\varepsilon > 0$, the procedure CHECKFEASIBILITY terminates and is correct.*

Proof. It will be convenient to use the following notation, for a vector $z \in \mathbb{R}^n$:

$$\mathbf{t}(z) := \max_i z_i, \quad \mathbf{b}(z) := \min_i z_i,$$

so that the halting condition reads

$$\mathbf{t}(u) \leq -\varepsilon \text{ or } \mathbf{b}(u) \geq \varepsilon.$$

The sequences $u^{(0)}, u^{(1)}, \dots$ and $v^{(0)}, v^{(1)}, \dots$ generated by the algorithm implemented in exact arithmetic satisfy, for $\ell \geq 1$,

$$u^{(\ell)} = F^\ell(0), \quad v^{(\ell)} = \max(0, F(0), \dots, F^{\ell-1}(0)).$$

We know from Corollary 12 and Theorem 17 that the limit $\lim_\ell u^{(\ell)}/2\ell$ coincides with the mean payoff vector χ . The assumptions of the present theorem imply that χ is a constant vector with nonzero entries, i.e., $\chi = (\lambda, \dots, \lambda)^\top$ where $\lambda \neq 0$. Therefore, $\lim_\ell \mathbf{b}(u^{(\ell)})/\ell = \lim_\ell \mathbf{t}(u^{(\ell)})/\ell = \lambda$. In particular, for ℓ sufficiently large, we have either $\mathbf{b}(u^{(\ell)}) > \varepsilon$, or $\mathbf{t}(u^{(\ell)}) < -\varepsilon$, depending on the sign of λ , and so the halting condition is ultimately satisfied. If $\mathbf{t}(u^{(\ell)}) < -\varepsilon$, then $F^\ell(0) \leq -\varepsilon e$, where $e = (1, 1, \dots, 1)$. Since F is order preserving and commutes with the addition of a constant vector, we get, after an immediate induction, that $F^{\ell p}(0) \leq -p\varepsilon e$ holds for all $p \geq 1$, and therefore, the mean payoff vector of the game, which coincides with $\lim_p F^{\ell p}/(2\ell p)$, has negative entries. By the Collatz–Wielandt property (10), this implies that the tropical

spectrahedron $\{x : x \leq F(x)\}$ is reduced to the trivial vector with $-\infty$ entries. If $\mathbf{b}(u^{(\ell)}) > \varepsilon$, then the dual argument shows that the mean payoff of the game $\lim_p F^{\ell p}/(2\ell p)$ has positive entries. Moreover, in this case we get

$$\begin{aligned} F(v) &= F(\max(0, F(0), \dots, F^{\ell-2}(0), F^{\ell-1}(0))) \\ &\geq \max(F(0), F^2(0), \dots, F^{\ell-1}(0), F^{\ell}(0)) \\ &\geq \max(F(0), F^2(0), \dots, F^{\ell-1}(0), \varepsilon e) \\ &\geq \max(F(0), F^2(0), \dots, F^{\ell-1}(0), 0) = v, \end{aligned}$$

which implies that v belongs to the tropical spectrahedron $\{x : x \leq F(x)\}$.

We note that the algorithm still terminates, and provides a correct yes/no answer, if the operator F is evaluated in fixed precision arithmetic, in such a way that u never differs from its true value by more than ε in the sup norm $\|\cdot\|$. Indeed, let $\tilde{u}^{(0)}, \tilde{u}^{(1)}, \dots$ denote the successive approximate values of the variables u which are computed. Our assumption entails that $\|\tilde{u}^{(\ell)} - u^{(\ell)}\| \leq \varepsilon$ at every step ℓ of the algorithm. The termination proof relies on the fact that either $\mathbf{b}(u^{(\ell)})$ tends to infinity, or $\mathbf{t}(u^{(\ell)})$ tends to $-\infty$. Since $\|\tilde{u}^{(\ell)} - u^{(\ell)}\| \leq \varepsilon$, the analogous property is still satisfied by $\tilde{u}^{(\ell)}$, and so, the algorithm implemented with finite precision does terminate.

Now, if ℓ^* is the step at which the procedure terminates, we have either $\mathbf{b}(\tilde{u}^{(\ell^*)}) > \varepsilon$ or $\mathbf{t}(\tilde{u}^{(\ell^*)}) < -\varepsilon$, which entails that either $\mathbf{b}(u^{(\ell^*)}) > 0$ or $\mathbf{t}(u^{(\ell^*)}) < 0$. Reasoning as above, we deduce that the mean payoff of the game is positive in the first situation and negative in the second one. Hence, the algorithm still decides correctly the triviality of the spectrahedron. \square

Remark 42. The situation in which the mean payoff is zero is degenerate: then, an infinitesimal perturbation of the entries of the matrices can make the spectrahedron trivial or nontrivial. Value iteration cannot naturally handle such degenerate situations, which can be solved by different methods, like the policy iteration algorithm presented in [ACTDG13], which is based on the solution of a finite sequence of linear systems, and can be implemented in exact arithmetic.

Remark 43. Every iteration (while loop) of the procedure CHECKFEASIBILITY takes a time $O(nm^2)$, which is *linear* in the size of the input. The number of iterations can only be bounded by an exponential in the size of the input. However, our benchmarks indicate that the algorithm is fast when the instance is “far from being degenerate”.

Remark 44. The condition that all the initial states have the same mean payoff can be checked by appealing to Theorem 32 involving the recession function \hat{F} of the Shapley operator. For instance, suppose that F arises from a Metzler tropical spectrahedron given by matrices whose entries $Q_{ij}^{(k)}$ are all finite. Then, one can verify that the k -th coordinate map of \hat{F} is given by

$$\hat{F}_k(x) = \min_{i < j} \frac{1}{2} \left(\max_{Q_{ii}^{(i)} \in \mathbb{T}_+} x_i + \max_{Q_{jj}^{(i)} \in \mathbb{T}_+} x_j \right),$$

and is therefore independent of the choice of $k \in [n]$. Then, it follows from Theorem 32 that \hat{F} has only fixed points of the form (α, \dots, α) , which implies that the first condition of Theorem 41 is satisfied.

Remark 45. By Theorem 3, the procedure CHECKFEASIBILITY allows us to verify the feasibility of a Metzler nonarchimedean spectrahedron if the entries of the matrices $Q_{ij}^{(k)}$ are finite and generic.

Example 46. Procedure CHECKFEASIBILITY applied to the nonarchimedean spectrahedron of Section 2, with $\varepsilon = 10^{-8}$, terminates in 20 iterations. It returns a vector having floating point entries

$$(1.05612 \dots, 0.0204082 \dots, 1.12755 \dots).$$




FIGURE 5. Phase transition between feasibility and infeasibility of spectrahedra. For each (n, m) , the color scheme reports the ratio of feasible instances among 10 samples.

(n, m)	(50, 10)	(50, 40)	(50, 50)	(50, 100)	(50, 1000)
time	0.000065	0.000049	0.000077	0.000279	0.026802
(n, m)	(100, 10)	(100, 15)	(100, 80)	(100, 100)	(100, 1000)
time	0.000025	0.000270	0.000366	0.000656	0.053944
(n, m)	(1000, 10)	(1000, 50)	(1000, 100)	(1000, 200)	(1000, 500)
time	0.000233	0.073544	0.015305	0.027762	0.148714
(n, m)	(2000, 10)	(2000, 70)	(2000, 100)	(10000, 150)	(10000, 400)
time	0.000487	1.852221	0.087536	19.919844	2.309174

TABLE 1. Execution time (in sec.) of Procedure CHECKFEASIBILITY on random instances.

When converted into a vector with rational entries, it reads

$$v := \left(1107425 \times 2^{-20}, 42799 \times 2^{-21}, 4729289 \times 2^{-22} \right).$$

We have checked using exact precision arithmetic over rationals (provided by the GNU multiple precision arithmetic library, <https://gmplib.org/>) that this vector lies in the interior of the tropical spectrahedron shown in Figure 2, i.e., $v_i < F(v_i)$ for $i = 1, 2, 3$. Following Lemma 24, the vector $\mathbf{x} := (t^{v_1}, t^{v_2}, t^{v_3})$ fulfills $\mathbf{Q}(\mathbf{x}) \succ 0$.

We report in Table 1 experimental results for different values of (n, m) . We chose all the $|Q_{ij}^{(k)}|$, for $i \leq j$, to be independent random variables uniformly distributed on $[0, 1]$. Moreover, the diagonal coefficients $Q_{ii}^{(k)}$ were chosen to have a positive tropical sign (they belong to \mathbb{T}_+). We took $\varepsilon = 10^{-8}$ (the performance was similar for $\varepsilon = 10^{-6}$ or $\varepsilon = 10^{-10}$). Our experiments were obtained using a C program, distributed as an ancillary file attached to this arXiv manuscript for reproducibility purposes.³ This program was compiled under Linux with gcc -O3, and executed on a single core of an Intel(R) i7-4600U CPU at 2.10 GHz with 16 GB RAM. We report the average execution time over 10 samples for every value of (n, m) . The number of iterations did not exceed 731 on this benchmark, and, for most (n, m) , it was limited to a few units. Indeed, random instances exhibit experimentally a phase transition, as shown in Figure 5: for a given (n, m) , the system is either feasible with overwhelming probability, or infeasible with overwhelming probability, unless (n, m) lies in a tiny region of the parameter space. Value iteration quickly decides feasibility, except in regions close to the phase transition. This explains why the execution time does not increase monotonically with (n, m) in our experiments (we included both easy and hard values of (n, m)).

³The ancillary file can be downloaded from <http://arxiv.org/src/1603.06916/anc>.

7. EQUIVALENT FORMS OF STOCHASTIC MEAN PAYOFF GAMES PROBLEM

In this section we present the (algorithmically) equivalent forms of stochastic mean payoff games, as mentioned in Remark 8. In order to do that, we need to introduce the notions of *simple stochastic games* and *stopping games*.

We start by defining the class of stopping games. We say that a pair of states $(i, k) \in [m] \times [n]$ is a *sink* if, when i or k is reached, the game loops forever between these two states. More formally, we have $A^{(k)} = \{\{i\}\}$, $B^{(i)} = \{\{k\}\}$. We say that the game is *stopping* if it has at least one sink and the probability that the game will reach a sink is equal one for every choice of policies (σ, τ) and every initial state. Note that if a game is stopping, then only the payoffs in sinks are important to determine its solution. Indeed, if we denote the sinks by $(i_1, k_1), \dots, (i_p, k_p)$ and the game is stopping, then the payoff of Player Max is given by

$$g_k(\sigma, \tau) = \sum_{s=1}^p (r_{k_s}^{\{i_s\}} + r_{i_i}^{\{k_s\}}) \mu_s(k, \sigma, \tau),$$

where $\mu_s(k, \sigma, \tau)$ is the probability that the game starting from k reaches the sink (i_s, k_s) if the players use the policies (σ, τ) . This expression depends only on payoffs in sinks.

Now, we introduce the class of simple stochastic games. We say that the stochastic game is *simple* if its set of states can be divided into three classes: states controlled by Player Min, Player Max, and Nature. Players Min and Max have only deterministic choices and Nature chooses the next state by tossing a coin. (To be coherent with the previous definition of stochastic game, we assume that Player Min controls the states of Nature — but she has no other choice than to toss a coin.) Formally, we suppose that for every $k \in [n]$ and every $a \in A^{(k)}$ we have the implication $|a| = 2 \implies |A^{(k)}| = 1$. Moreover, a simple game has two sinks: one with payoff 1 and the other with payoff 0. All other payoffs are equal to 0.

It may seem that solving simple games is indeed simpler than solving games in their full generality. Andersson and Miltersen [AM09] have shown that this is not the case. Let SMPG-COMP denote the problem of finding the value and a pair of optimal policies in a stochastic mean payoff game.

Theorem 47 ([AM09]). *SMPG-COMP is poly-time equivalent to the problem of finding the values of stopping simple stochastic games.*

Note that this theorem was originally without the word “stopping”, but this is what the authors actually showed in the latter reference. As already mentioned in Remark 8, this result is valid for a much wider class of games than those considered in this work. Andersson and Miltersen defined stopping simple stochastic games in a slightly more general way — in their version, the players can make multiple moves in a row. Nevertheless, observe that we can always add dummy states to a stopping simple stochastic game (i.e., states in which player has only one action) and obtain an equivalent game that belongs to the class considered here. This shows that, from the algorithmic point of view, these classes are equivalent.

Moreover, observe that if the game is both simple and stopping, then we can change its payoffs in sinks — instead of payoffs equal to 0 and 1, we can demand them to be equal to -1 and 1 . This does not change the optimal policies of the game and acts as an affine transformation on the value vector. Henceforth, we assume that payoffs in sinks of stopping simple stochastic games are equal to -1 and 1 . We now show that the computational problem of finding values of simple games can be reduced to the decision problem. By $\text{SMPG}(k)$ we will denote the problem of deciding if a given state $k \in [n]$ is winning in the stochastic mean payoff game.

Lemma 48. *$\text{SMPG}(k)$ restricted to stopping simple stochastic games is poly-time equivalent to SMPG-COMP.*

Proof. It is obvious that $\text{SMPG}(k)$ can be reduced to SMPG-COMP. We will show the opposite reduction. By Theorem 47, SMPG-COMP is poly-time reducible to the problem of finding values of stopping simple stochastic games. Fix such a game and let $\chi \in [-1, 1]^n$ denote its value.




FIGURE 6. Transformation of stopping simple stochastic games. Circle states are controlled by Player Min. Numbers indicate payoffs received by Player Max after each move.

First, we show an auxiliary reduction. Fix a rational number $\alpha \in [-1, 1] \cap \mathbb{Q}$ and an initial state $k \in [n]$. Suppose that we want to decide if $\chi_k \geq \alpha$. We will show that this is poly-time reducible (poly-time in the size of the game and the number of bits needed to encode α) to $\text{SMPG}(k)$. If $\alpha = -1$, then the answer is “yes”. If $\alpha > -1$, then we may modify the game as follows: we suppose that when the sink with payoff 1 is reached, the game does not start to loop, but instead moves with probability $1 - \frac{1}{\alpha+1}$ to the sink with payoff -1 and with probability $\frac{1}{\alpha+1}$ to the (newly created) sink with payoff 1. Denote the value of the modified game by $\tilde{\chi}$. We have $\tilde{\chi}_k \geq 0 \iff \chi_k \geq \alpha$. The modified game is stopping but not simple. Nevertheless, we may apply the construction of Zwick and Paterson [ZP96, remarks preceding Theorem 6.1] and obtain (in poly-time) a new game, which is stopping, simple, and has $\tilde{\chi}_k$ as the value at state k . This gives the auxiliary reduction.

Second, [ACS14, Lemma 1] (which improves [Con92, Lemma 2]) shows that χ_k is a rational number of form $a/b \in [-1, 1]$, where a, b are integers and $0 \leq b \leq 3^n$. Thus, χ_k can be found by the Kwek–Mehlhorn algorithm [KM03], using polynomially many queries to the oracle given by our auxiliary reduction. \square

Finally, we want to show that $\text{SMPG}(k)$ is poly-time reducible to SMPG . This requires an auxiliary construction which is presented in Figure 6. We take a stopping simple stochastic game Γ , fix an initial state $k_0 \in [n]$ and suppose that the sinks of Γ are indexed as $(n-1, m-1)$ (sink with payoff -1) and (n, m) (sink with payoff 1). Now, we modify the game as follows: we add two states, $n+1$ (controlled by Player Min) and $m+1$ (controlled by Player Max). At $n+1$, Player Min has only one possible action: to go to $m+1$; after this action Player Min pays 0 to Player Max. Moreover, at $m+1$ Player Max also has only one action: to go to k_0 ; after this action Player Max receives 0 from Player Min. Finally, we modify the sinks of Γ as follows: at $m-1$ (resp. m) Player Max has only one possible action: to go to $n+1$; after this action he receives -1 (resp. 1) from Player Min. Denote the modified game by $\bar{\Gamma}$. By construction, it is quite intuitive that the value of $\bar{\Gamma}$ does not depend on the initial state and that the state k_0 is winning in $\bar{\Gamma}$ if and only if it is winning in Γ . To prove this formally, we use Theorem 53. Henceforth, by $\Gamma^{\sigma, \tau}$ we denote the 0-player game obtained from Γ by fixing a pair of policies (σ, τ) .

Corollary 49. *The state k_0 is winning in Γ if and only if it is winning in $\bar{\Gamma}$. Moreover, the value of $\bar{\Gamma}$ is independent of the initial state.*

Proof. First, observe that there exists a natural bijection between policies of Γ and $\bar{\Gamma}$. Hence, we will use the same letters to denote policies in both games. Fix a pair of policies (σ, τ) . Let $g(\sigma, \tau)$ (resp. $\bar{g}(\sigma, \tau)$) denote the payoff of Player Max in $\Gamma^{\sigma, \tau}$ (resp. in $\bar{\Gamma}^{\sigma, \tau}$). Since Γ was stopping, the 0-player game $\bar{\Gamma}^{\sigma, \tau}$ obtained from $\bar{\Gamma}$ by fixing the policies (σ, τ) has only one recurrent class and k_0 belongs to this class. By Theorem 53 we see that $\bar{g}(\sigma, \tau)$ is constant for all initial states. Since (σ, τ) are arbitrary, this shows that the value of $\bar{\Gamma}$ does not depend on the choice of initial state. Furthermore, Theorem 53 shows that $\bar{g}_{k_0}(\sigma, \tau) = g_{k_0}(\sigma, \tau)/\theta_{k_0}$, where $\theta_{k_0} = \theta_{k_0}(\sigma, \tau)$ is the expected time of first return to k_0 .

Now, suppose that the value of Γ starting from k_0 is higher or equal than 0. Let τ^* denote the optimal policy of Player Max in Γ . For any policy σ of Player Min we have $g_{k_0}(\sigma, \tau^*) \geq 0$ and hence $\bar{g}_{k_0}(\sigma, \tau^*) \geq 0$. Therefore, τ^* is a winning (but not necessarily optimal) policy for Player Max in $\bar{\Gamma}$ starting from k_0 . Thus, the value of $\bar{\Gamma}$ starting from k_0 is higher or equal than 0. The opposite implication is analogous. \square

Corollary 50. *SMPG-COMP is poly-time equivalent to SMPG restricted to games with payoffs in $\{-1, 0, 1\}$.*

Proof. By Lemma 48, SMPG-COMP can be reduced to SMPG(k) restricted to stopping simple stochastic games. By the construction described above and Corollary 49, SMPG(k) for stopping simple stochastic games can be reduced to SMPG restricted to games with payoffs in $\{-1, 0, 1\}$. The opposite reduction is trivial. \square

Corollary 51. *SMPG restricted to games with payoffs in $\{-1, 0, 1\}$ is poly-time equivalent to SMPG for general games.*

Proof. SMPG is trivially reducible to SMPG-COMP. Therefore, the claim follows from Corollary 50. \square

8. CONCLUDING REMARKS

In this paper, we have shown that under a genericity condition on the valuations, solving feasibility semidefinite problems over the field of Puiseux series reduces to a well studied class of zero-sum stochastic games. This leads both to complexity bounds and to algorithms capable experimentally to solve large scale nonarchimedean instances. The interest is also to relate two different problems which both have unsettled complexities. This is the first exposition of this approach.

It would be interesting to relax the current genericity conditions. We believe that finer genericity conditions could involve both the valuations and leading coefficients of the series.

Another interesting question is to use the present approach to deal with the real case. We already showed in Section 5.3 that the nonarchimedean feasibility problem is equivalent to the archimedean one for large values of t , with an exponential number of bits. We may ask however whether it is possible to use combinatorial methods to work for smaller values of t .

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their numerous remarks which helped to improve the presentation of the paper. An abridged version of the present work appeared initially in the ISSAC paper [AGS16a]. We also thank the referees of ISSAC for their detailed comments.

REFERENCES

- [ACS14] D. Auger, P. Couheney, and Y. Strobecki. Finding optimal strategies of almost acyclic simple stochastic games. In *Proceedings of the 11th Annual Conference on Theory and Applications of Models of Computation (TAMC)*, volume 8402 of *Lecture Notes in Comput. Sci.*, pages 67–85. Springer, 2014.

- [ACTDG13] M. Akian, J. Cochet-Terrasson, S. Detournay, and S. Gaubert. Solving multichain stochastic games with mean payoff by policy iteration. In *52nd IEEE Annual Conference on Decision and Control (CDC)*, pages 1834–1841. IEEE, 2013.
- [AGG09] M. Akian, S. Gaubert, and A. Guterman. Linear independence over tropical semirings and beyond. In *Proceedings of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemp. Math.*, pages 1–38. AMS, 2009.
- [AGG12] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *Int. J. Algebra Comput.*, 22(1):125001 (43 pages), 2012.
- [AGH15] M. Akian, S. Gaubert, and A. Hochart. Ergodicity conditions for zero-sum games. *Discrete Contin. Dyn. Syst.*, 35(9):3901–3931, 2015.
- [AGS16a] X. Allamigeon, S. Gaubert, and M. Skomra. Solving generic nonarchimedean semidefinite programs using stochastic game algorithms. In *Proceedings of the 41st International Symposium on Symbolic and Algebraic Computation (ISSAC)*, pages 31–38. ACM, 2016.
- [AGS16b] X. Allamigeon, S. Gaubert, and M. Skomra. Tropical spectrahedra. arXiv:1610.06746v2, 2016.
- [AM09] D. Andersson and P. B. Miltersen. The complexity of solving stochastic games on graphs. In *Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC)*, volume 5878 of *Lecture Notes in Comput. Sci.*, pages 112–121. Springer, 2009.
- [BB16] M. Baker and N. Bowler. Matroids over hyperfields. arXiv:1601.01204, 2016.
- [BPT13] G. Blekherman, P. A. Parrilo, and R. R. Thomas. *Semidefinite Optimization and Convex Algebraic Geometry*, volume 13 of *MOS-SIAM Ser. Optim.* SIAM, Philadelphia, PA, 2013.
- [But10] P. Butkovič. *Max-linear Systems: Theory and Algorithms*. Springer Monogr. Math. Springer, London, 2010.
- [CC11] A. Connes and C. Consani. The hyperring of adèle classes. *J. Number Theory*, 131(2):159–194, 2011.
- [Chu67] K. L. Chung. *Markov Chains With Stationary Transition Probabilities*, volume 104 of *Grundlehren Math. Wiss.* Springer, Heidelberg, 1967.
- [Con92] A. Condon. The complexity of stochastic games. *Inform. and Comput.*, 96(2):203–224, 1992.
- [dKV16] E. de Klerk and F. Vallentin. On the Turing model complexity of interior point methods for semidefinite programming. *SIAM J. Optim.*, 26(3):1944–1961, 2016.
- [FV07] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, New York, 2007.
- [GLS93] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms Combin.* Springer, Berlin, 1993.
- [GM12] B. Gärtner and J. Matoušek. *Approximation Algorithms and Semidefinite Programming*. Springer, Heidelberg, 2012.
- [Hal07] N. Halman. Simple stochastic games, parity games, mean payoff games and discounted payoff games are all LP-type problems. *Algorithmica*, 49(1):37–50, 2007.
- [HNSED16] D. Henrion, S. Naldi, and M. Safey El Din. Exact algorithms for linear matrix inequalities. *SIAM J. Optim.*, 26(4):2512–2539, 2016.
- [HR15] G. H. Hardy and M. Riesz. *The general theory of Dirichlet’s series*. Cambridge University Press, Cambridge, 1915.
- [HZ15] T. D. Hansen and U. Zwick. An improved version of the Random-Facet pivoting rule for the simplex algorithm. In *Proceedings of the 47th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 209–218. ACM, 2015.
- [KM03] S. Kwek and K. Mehlhorn. Optimal search for rationals. *Inform. Process. Lett.*, 86(1):23–26, 2003.
- [Koh80] E. Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Math. Oper. Res.*, 5(3):366–372, 1980.
- [LL69] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Rev.*, 11(4):604–607, 1969.
- [MC11] D. Monniaux and P. Corbineau. On the generation of Positivstellensatz witnesses in degenerate cases. In *Proceedings of the Second international conference on Interactive theorem proving (ITP)*, pages 249–264. ACM, 2011.
- [Meg89] N. Megiddo. On the complexity of linear programming. In T. F. Bewley, editor, *Advances in economic theory*, volume 12 of *Econom. Soc. Monogr.*, pages 225–268. Cambridge University Press, Cambridge, 1989.
- [MS15] D. Maclagan and B. Sturmfels. *Introduction to Tropical Geometry*, volume 161 of *Grad. Stud. Math.* AMS, Providence, RI, 2015.
- [Nal18] S. Naldi. Solving rank-constrained semidefinite programs in exact arithmetic. *J. Symbolic Comput.*, 85:206–223, 2018.
- [NRS10] J. Nie, K. Ranestad, and B. Sturmfels. The algebraic degree of semidefinite programming. *Math. Program.*, 122(2):379–405, 2010.
- [Nus86] R. D. Nussbaum. Convexity and log convexity for the spectral radius. *Linear Algebra Appl.*, 73:59–122, 1986.

- [Put05] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Ser. Probab. Stat. Wiley, Hoboken, NJ, 2005.
- [Ram97] M. V. Ramana. An exact duality theory for semidefinite programming and its complexity implications. *Math. Program.*, 77(1):129–162, 1997.
- [Sch16] C. Scheiderer. Sums of squares of polynomials with rational coefficients. *J. Eur. Math. Soc.*, 18(7):1495–1513, 2016.
- [Ste10] S. A. Steinberg. *Lattice-ordered Rings and Modules*. Springer, New York, 2010.
- [vdDS98] L. van den Dries and P. Speissegger. The real field with convergent generalized power series. *Trans. Amer. Math. Soc.*, 350(11):4377–4421, 1998.
- [Vir10] O. Viro. Hyperfields for tropical geometry I. Hyperfields and dequantization. arXiv:1006.3034, 2010.
- [Yu15] J. Yu. Tropicalizing the positive semidefinite cone. *Proc. Amer. Math. Soc.*, 143(5):1891–1895, 2015.
- [ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoret. Comput. Sci.*, 158(1–2):343–359, 1996.

APPENDIX A. CONSTRUCTING SPECTRAHEDRA FROM MEAN PAYOFF GAMES

Let Γ be a stochastic mean payoff game, and let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ be the matrices as constructed in the paragraph following Theorem 23. From these matrices, we can define the set \mathcal{S}_{λ} as in Definition 21. Denoting by F the Shapley operator of Γ , we still have:

Lemma 52. *For any $\lambda \in \mathbb{R}$ we have $\mathcal{S}_{\lambda} = \{x \in \mathbb{T}^n : \lambda + x \leq F(x)\}$.*

Proof. Let

$$(F(x))_k = \min_{\substack{a \in A^{(k)} \\ a = \{i, j\}}} \left(r_k^a + \frac{1}{2} \left(\max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right) \right)$$

denote the Shapley operator of Γ . As in the proof of Lemma 22, we have the equivalence

$$\begin{aligned} \forall k \in [n], \lambda + x_k &\leq (F(x))_k \iff \\ \forall k \forall a \in A^{(k)}, a = \{i, j\}, \lambda + x_k - r_k^a &\leq \frac{1}{2} \left(\max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right) \iff \\ \forall (i, j), \lambda + \max_{\{a \in \bigcup_k A^{(k)} : a = \{i, j\}\}} (x_k - r_k^a) &\leq \frac{1}{2} \left(\max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right). \end{aligned}$$

We want to show that the last set of constraints describes \mathcal{S}_{λ} . To do this, recall that an inequality of the form $\max(x, \alpha + y) \geq \max(x', \beta + y)$ is equivalent to $\max(x, \alpha + y) \geq x'$ if $\alpha \geq \beta$, and to $x \geq \max(x', \beta + y)$ if $\beta > \alpha$. Therefore, for every $i \in [m]$ we have the equivalence

$$\begin{aligned} \lambda + \max_{\{a \in \bigcup_k A^{(k)} : a = \{i, i\}\}} (x_k - r_k^a) &\leq \max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l) \iff \\ \lambda + \max_{Q_{ii}^{(k)} \in \mathbb{T}_-} (x_k + |Q_{ii}^{(k)}|) &\leq \max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l). \end{aligned} \tag{14}$$

Moreover, note that if $x \in \mathbb{T}^n$ verifies (14), then we have the equality

$$\max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l) = \max_{\substack{b \in B^{(i)} \\ b = \{l\}}} (r_i^b + x_l). \tag{15}$$

Indeed, if we have $\max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l) < r_i^{\{l^*\}} + x_{l^*}$ for some $l^* \in [n]$, then by the construction we get $(-r_{l^*}^{\{i\}}) > r_i^{\{l^*\}}$ and $Q_{ii}^{(l^*)} = \ominus(-r_{l^*}^{\{i\}})$. In particular, $\max_{Q_{ii}^{(k)} \in \mathbb{T}_-} (x_k + |Q_{ii}^{(k)}|) > r_i^{\{l^*\}} + x_{l^*}$, what gives a contradiction with (14). Furthermore, observe that for any $i \neq j$ we have the equality

$$\lambda + \max_{\{a \in \bigcup_k A^{(k)} : a = \{i, j\}\}} (x_k - r_k^a) = \lambda + \max_{Q_{ij}^{(k)} \in \mathbb{T}_-} (x_k + |Q_{ij}^{(k)}|). \tag{16}$$

Suppose that $x \in \mathcal{S}_\lambda$. Then x verifies (14) for all $i \in [m]$. Hence, by (15) and (16), for any $i \neq j$ we have

$$\begin{aligned}
(17) \quad \lambda + \max_{\{a \in \bigcup_k A^{(k)} : a = \{i, j\}\}} (x_k - r_k^a) &= \lambda + \max_{Q_{ij}^{(k)} \in \mathbb{T}_-} (x_k + |Q_{ij}^{(k)}|) \\
&\leq \frac{1}{2} \max_{Q_{ii}^{(l)} \in \mathbb{T}_+} (Q_{ii}^{(l)} + x_l) + \frac{1}{2} \max_{Q_{jj}^{(l)} \in \mathbb{T}_+} (Q_{jj}^{(l)} + x_l) \\
&= \frac{1}{2} \left(\max_{\substack{b \in B^{(i)} \\ b = \{i\}}} (r_i^b + x_l) + \max_{\substack{b \in B^{(j)} \\ b = \{l\}}} (r_j^b + x_l) \right).
\end{aligned}$$

Thus $\lambda + x \leq F(x)$. Conversely, if $\lambda + x \leq F(x)$, then x also verifies (14) for all $i \in [m]$, and the same argument as in (17) shows that $x \in \mathcal{S}_\lambda$. \square

APPENDIX B. MARKOV CHAINS

Let us recall some facts about Markov chains with rewards. We only consider Markov chains on finite spaces. Suppose that we are given a Markov chain (X_0, X_1, \dots) defined on a finite space E . Recall that such a chain can be described by a *transition matrix* $P \in [0, 1]^{|E| \times |E|}$, where P_{uw} denotes the probability that chain moves from state u to state w in one step.

If $C \subset E$, then we say that $\pi \in \mathbb{R}_{\geq 0}^{|C|}$ is a *stationary distribution* on the set C if $\pi_u = \sum_{w \in C} \pi_w P_{wu}$ for all $u \in C$ and $\sum_{u \in C} \pi_u = 1$. Furthermore, a set $C \subset E$ is called a *recurrent class* if it has the following two properties: (a) if C is reached, then the chain will never leave it; (b) if C is reached, then every state in C will be visited infinitely many times (with probability one). Every Markov chain has at least one recurrent class, and every recurrent class has a unique stationary distribution. A state $u \in E$ is called *recurrent* if it belongs to a recurrent class. Otherwise, it is called *transient*.

We now introduce Markov chains with payoffs. To this end, with every state $u \in E$ we associate a payoff $r_u \in \mathbb{R}$. This quantity is interpreted as follows: there is a controller of the chain, who receives a payoff r_u as soon as the chain leaves the state u . A *(long-term) average payoff* of the controller is defined as

$$\forall u \in E, g_u = \lim_{N \rightarrow \infty} \mathbb{E} \frac{1}{N} \sum_{p=1}^N r_{u_p},$$

where the expectation is taken over all trajectories u_1, \dots, u_N starting from $u_1 = u$ in the Markov chain. The next theorem characterizes the average payoff. Before that, let us introduce some additional notation.

For any state $u \in E$, let the random variable $T_u = \inf\{s \geq 1 : X_s = u\}$ denote the time of first return to u . By θ_u we denote the expected time of first return to u ,

$$\theta_u = \mathbb{E}(T_u | X_0 = u).$$

Furthermore, let ξ_u be the expected payoff the controlled obtained before returning to u , i.e.,

$$\xi_u = \mathbb{E} \left(\sum_{s=0}^{T_u-1} r_{X_s} \mid X_0 = u \right).$$

Theorem 53. *If $u \in E$ is a fixed initial state, then the average payoff g_u is well defined and characterized as follows:*

- (i) *Suppose that u is a recurrent state belonging to the recurrent class C . Let $(\pi_w)_{w \in C}$ be the stationary distribution on C . Then $\pi_u = 1/\theta_u$. Furthermore, we have*

$$g_u = \frac{\xi_u}{\theta_u} = \sum_{w \in C} r_w \pi_w.$$

In particular, g_u is constant for all states u belonging to C .

- (ii) If u is transient and C_1, \dots, C_p denote all the recurrent classes of the Markov chain, then $g_u = \sum_{s=1}^p g_{u_s} \psi_s$, where, for all s , ψ_s denotes the probability that the chain starting from u reaches the recurrent class C_s , and $u_s \in C_s$ is an arbitrary state of C_s .

Remark 54. We point out that given the transition matrix P , the average payoff g can be computed using the algorithm presented in [Put05, Appendix A.3 and A.4]. This algorithm can be implemented to run in strongly polynomial complexity using the strongly polynomial version of gaussian elimination (presented, for example, in [GLS93, Section 1.4]).

Theorem 53 is well known, and can be easily derived from the analysis of Markov chains presented in the textbook of Chung [Chu67, Part I, §6, §7, and §9]. We give the details for the sake of completeness. Let μ_{uw} denote the probability that the Markov chain starting from $u \in E$ will reach $w \in E$ at least once, $\mu_{uw} = \mathbb{P}(\exists s \geq 1, X_s = w | X_0 = u)$. By definition, the state u is recurrent if $\mu_{uu} = 1$, and it is transient otherwise. By ζ_{uw} we denote the expected number of visits in w before returning to u , i.e.,

$$\zeta_{uw} = \mathbb{E} \left(\sum_{s=0}^{T_u-1} \mathbf{1}_{\{X_s=w\}} \mid X_0 = u \right).$$

The following theorem describes the ergodic behavior of any finite (or countable) Markov chain.

Theorem 55. *The Cesaro limit*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{s=0}^N P^s$$

is well defined (we will denote it by M). Moreover, the entries of M are given as follows: if w is a transient state, then $M_{uw} = 0$ for all u . If w is recurrent, then $M_{uw} = \frac{\mu_{uw}}{\theta_w}$ for all u .

Proof. See [Chu67, Part I, §6, Theorem 4 and its Corollary]. \square

Remark 56. Note that the theorem above does not state that $\theta_w < \infty$ if the state w is recurrent. (We work under the convention that $\frac{a}{+\infty} = 0$ for all finite a .) Nevertheless, if the chain is finite, then we have $\theta_w < \infty$ for all recurrent states w , and this can be deduced as a corollary of the theorem above, as discussed below.

Observe that M is a stochastic matrix (as a limit of stochastic matrices). Moreover, we have $MP = PM = M$. This leads to the following corollary.

Corollary 57. *If $C \subset E$ is a recurrent class, then $M_{uw} = \frac{1}{\theta_w}$ for all $u, w \in C$. Furthermore, $(\pi_u)_{u \in C}$ defined as $\pi_u = \frac{1}{\theta_u}$ is the unique stationary distribution on C . In particular, if $u \in E$ is a recurrent state, then $\theta_u < \infty$.*

Proof. Let first claim follows immediately from Theorem 55. We will prove that $(\pi_u)_{u \in C}$ is a stationary distribution on C . Let P_C denote the square submatrix of P formed by the rows and columns of P with indices in C . We define M_C analogously. The first claim implies that M_C has identical rows. Since C is a recurrent class, we have $P_{uw'} = 0$ for all $u \in C, w' \notin C$. Hence, for all s we have $[P^s]_C = [P_C]^s$. Therefore $M_C = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{s=0}^N P_C^s$. Hence, M_C is stochastic. In other words, every row of M_C is a probability distribution on C and, since $M_C P_C = M_C$, this distribution is a stationary distribution on C . Since C is a recurrent class, stationary distribution on C has only strictly positive values. Hence we have $\theta_u < \infty$. The fact that the stationary distribution is unique follows from [Chu67, Part I, §7, Theorem 1]. \square

The next theorem characterizes the relationship between entries of M and the values ζ_{uw} .

Theorem 58. *If (u, w) belong to the same recurrent class, then $0 < \zeta_{uw} < \infty$. Moreover, if (f, g, h, u, w) are (not necessarily distinct) states belonging to the same recurrent class, then $M_{fg}/M_{hu} = \zeta_{wg}/\zeta_{wu}$.*

Proof. The fact that $0 < \zeta_{uw} < \infty$ follows from [Chu67, Part I, §9, Theorem 2 and 3]. Moreover, by Corollary 57 we have $M_{uw} = \frac{1}{\theta_w} > 0$. Hence the claim follows from [Chu67, Part I, §9, Theorem 5 and remarks that precede it]. \square

Corollary 59. *If (u, w) belong to the same recurrent class, then*

$$M_{uw} = \frac{1}{\theta_w} = \frac{\zeta_{uw}}{\theta_u}.$$

Proof. By Theorem 55 and Theorem 58 we have $\theta_u/\theta_w = M_{ww}/M_{uu} = \zeta_{uw}/\zeta_{uu}$. By definition $\zeta_{uu} = 1$ and hence $\theta_u/\theta_w = \zeta_{uw}$. \square

Proof of Theorem 53. Fix $u \in E$. Observe that for all $N \geq 1$ we have

$$\mathbb{E}\left(\sum_{s=0}^N r_{X_s} \middle| X_0 = u\right) = [r + Pr + \dots + P^N r]_u.$$

Therefore $g_u = [Mr]_u$. In particular, g_u is well defined.

Let us suppose that the initial state u is recurrent and denote its recurrent class by C . In this case, Corollary 57 and Theorem 55 imply that $g_u = \sum_{w \in C} r_w \pi_w$, where $(\pi_w)_{w \in C}$ is the stationary distribution on C . Moreover, Corollary 59 gives the identity

$$\begin{aligned} g_u &= \frac{1}{\theta_u} \sum_{w \in C} \zeta_{uw} r_w \\ &= \frac{1}{\theta_u} \mathbb{E}\left(\sum_{w \in C} \sum_{s=0}^{T_u-1} r_w \mathbf{1}_{\{X_s=w\}} \middle| X_0 = u\right) \\ &= \frac{1}{\theta_u} \mathbb{E}\left(\sum_{s=0}^{T_u-1} \sum_{w \in E} r_w \mathbf{1}_{\{X_s=w\}} \middle| X_0 = u\right) \\ &= \frac{1}{\theta_u} \mathbb{E}\left(\sum_{s=0}^{T_u-1} r_{X_s} \middle| X_0 = u\right) = \frac{\xi_u}{\theta_u}. \end{aligned}$$

Now, suppose that the initial state u is transient. Let C_1, \dots, C_p denote the recurrent classes in our Markov chain. In this case, Theorem 55 gives the identity

$$g_u = \sum_{s=1}^p \sum_{w \in C_s} \frac{\psi_s}{\theta_w} r_w = \sum_{s=1}^p \psi_s g_{u_s}. \quad \square$$

INRIA AND CMAP, ÉCOLE POLYTECHNIQUE, CNRS, UNIVERSITÉ PARIS-SACLAY, 91128 PALAISEAU CEDEX FRANCE

E-mail address: `firstname.lastname@inria.fr`