

Condition numbers of stochastic mean payoff games and what they say about nonarchimedean semidefinite programming*

X. Allamigeon¹

S. Gaubert¹

R. Katz²

M. Skomra¹

Abstract—Semidefinite programming can be considered over any real closed field, including fields of Puiseux series equipped with their nonarchimedean valuation. Nonarchimedean semidefinite programs encode parametric families of classical semidefinite programs, for sufficiently large values of the parameter. Recently, a correspondence has been established between nonarchimedean semidefinite programs and stochastic mean payoff games with perfect information. This correspondence relies on tropical geometry. It allows one to solve generic nonarchimedean semidefinite feasibility problems, of large scale, by means of stochastic game algorithms. In this paper, we show that the mean payoff of these games can be interpreted as a condition number for the corresponding nonarchimedean feasibility problems. This number measures how close a feasible instance is from being infeasible, and vice versa. We show that it coincides with the maximal radius of a ball in Hilbert’s projective metric, that is included in the feasible set. The geometric interpretation of the condition number relies in particular on a duality theorem for tropical semidefinite feasibility programs. Then, we bound the complexity of the feasibility problem in terms of the condition number. We finally give explicit bounds for this condition number, in terms of the characteristics of the stochastic game. As a consequence, we show that the simplest algorithm to decide whether a stochastic mean payoff game is winning, namely value iteration, has a pseudopolynomial complexity when the number of random positions is fixed.

I. INTRODUCTION

A. Motivation

Semidefinite programming (SDP) consists in optimizing a linear function over a spectrahedron, the latter being the intersection of a cone of positive semidefinite matrices with an affine space. Semidefinite programs arise in a number of applications from engineering sciences and combinatorial optimization. We refer the reader to [1], [2] for more background on the theory and applications of semidefinite programming.

Spectrahedra form a class of convex semialgebraic sets. Even though these sets are usually defined over the field

of real numbers, their definition is meaningful over any real closed field. In particular, the complexity of SDP and related questions can be investigated over real closed nonarchimedean fields, like fields of Puiseux series. Such nonarchimedean SDP instances, which arise in perturbation theory, encode parametric families of classical SDP instances (over the reals), for large enough (or small enough) values of the parameter. The study of the nonarchimedean case is also motivated by unsettled questions concerning the complexity of ordinary SDP. Indeed, the latter are solvable in “polynomial time” only in a restricted sense. More precisely, complexity bounds for SDP, obtained by interior point or by the ellipsoid methods, are only polynomial in the log of certain metric estimates whose bit-size can be doubly exponential in the size of the input [3]. It is unknown whether the SDP feasibility problem belongs to NP.

Semidefinite feasibility problems over the nonarchimedean valued field of Puiseux series have been studied in [4]. It is shown there that, under a genericity condition, these problems are equivalent to solving stochastic mean payoff games with perfect information and finite state and action spaces. Stochastic mean payoff games have an unsettled complexity: they belong to $NP \cap coNP$ but no polynomial time algorithm is currently known [5], [6]. However, several practically efficient algorithms to solve stochastic mean payoff games have been developed. In this way, one can solve nonarchimedean semidefinite instances of a scale probably unreachable by interior point methods. For instance, the benchmarks presented in [4] show that random instances of these problems with as many as 10000 variables could be solved by value iteration in a few seconds. In addition, the benchmarks there indicate a phase transition, value iteration easily decides most feasibility instances. The hard random instances are concentrated in a small region of the instance space.

B. Main results

In order to explain why value iteration is so efficient on many nonarchimedean SDP instances, we introduce here a notion of *condition number* for stochastic mean payoff games. Essentially, for a feasible instance, the condition number is the inverse of the distance of the data to an infeasible instance, and vice versa. We show that this condition number coincides with the absolute value of the mean payoff. We establish a universal bound for the time of convergence of value iteration, involving the condition number and an auxiliary metric estimate, the distance of point 0 to the set of “bias vectors” (Theorem 16). Then, we effectively bound

*X. Allamigeon, S. Gaubert, and M. Skomra 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. The first two authors also gratefully acknowledge the support of Mittag-Leffler institute.

¹INRIA and CMAP, École Polytechnique, CNRS, 91128 Palaiseau Cedex, France `firstname.lastname@inria.fr`

²CONICET-CIFASIS, Bv. 27 de Febrero 210 bis, 2000 Rosario, Argentina `katz@cifasis-conicet.gov.ar`

AMS Classification – Primary: 14T05; Secondary: 91A15, 90C22

Key words: Semidefinite programming, stochastic games, tropical geometry, nonarchimedean fields

the condition number and the latter distance, for stochastic mean payoff games with perfect information (Theorems 18 and 19). We arrive, in particular, at a bound that becomes pseudopolynomial when the number of “random” positions of the game is fixed.

To arrive at these results, we develop a metric geometry approach of the condition number. We use Hilbert’s projective metric, which arises in Perron–Frobenius theory [7]. The same metric, up to a logarithmic change of variable, arises in tropical geometry [8], [9]. We also prove duality results for stochastic mean payoff games, showing, essentially, that the condition number of the primal and dual problems coincide. In summary, our main results show that the complexity of value iteration is governed by metric geometry properties: this leads to a general method to derive complexity bounds, which can be applied to various classes of Shapley operators.

C. Related works

When specialized to stochastic mean payoff games with perfect information, our bounds should be compared with the one of Boros, Elbassioni, Gurvich, and Makino [10]. The authors of [10] generalize the “pumping” algorithm, developed for deterministic games by Gurvich, Karzanov, and Khachiyan [11], to the case of stochastic games. The resulting algorithm is also pseudopolynomial if the number of random positions is fixed. The algorithm of Ibsen-Jensen and Miltersen [12] yields a stronger bound in the case of simple stochastic games, still assuming that the number of random positions is fixed.

The duality results in Section IV-A extend to stochastic games some duality results for deterministic games by Grigoriev and Podolskii [13]. In contrast, our approach builds on [9], deriving duality results from a minimax Collatz–Wielandt type theorem of Nussbaum [14]. Other duality results, by Bodirsky and Mamino, in the context of satisfiability problems, have appeared in [15].

D. Organization of the paper

Earlier results on the relation between nonarchimedean semidefinite programming and stochastic mean payoff games are presented in Section II, leading to the introduction of the notion of condition number. Some background on nonlinear Perron–Frobenius theory is presented in Section III. The new results are included in Section IV, in which we characterize the condition number, and in Section V, in which we derive complexity estimates for value iteration in terms of the condition number.

II. MOTIVATION: THE CORRESPONDENCE BETWEEN NONARCHIMEDEAN SEMIDEFINITE PROGRAMMING AND STOCHASTIC MEAN PAYOFF GAMES

In this section, we summarize some of the main results of [4], which motivate the present work. Throughout this paper, given $k \in \mathbb{N}$, we denote the set $\{1, \dots, k\}$ by $[k]$.

A. Nonarchimedean semidefinite programs

We start by introducing semidefinite programming over nonarchimedean fields. More specifically, the model of nonarchimedean field used in this paper is the field \mathbb{K} of (*absolutely convergent generalized real*) *Puiseux series*, which are series in the parameter t of the form

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

where (i) $(\lambda_i)_{i \geq 1}$ is a strictly decreasing sequence of real numbers that is either finite or unbounded, (ii) $c_{\lambda_i} \in \mathbb{R} \setminus \{0\}$ for all λ_i , (iii) and the series (1) is absolutely convergent for $t \in \mathbb{R}$ sufficiently large. There is also a special, empty series, which is denoted by $\mathbf{0}$. The field \mathbb{K} is ordered, with a total order defined by $\mathbf{x} > \mathbf{0} \iff c_{\lambda_1} > 0$. In addition, it is known that \mathbb{K} is a real closed field [16]. Actually, our approach applies to other nonarchimedean fields with a real value group [17], it is helpful, however to have a concrete field in mind, like \mathbb{K} . Henceforth, we denote by $\mathbb{K}_{\geq 0}$ the set of nonnegative series, i.e., the series \mathbf{x} that satisfy $\mathbf{x} = \mathbf{0}$ or $\mathbf{x} > \mathbf{0}$.

Given symmetric matrices $\mathbf{Q}^{(0)}, \mathbf{Q}^{(1)}, \dots, \mathbf{Q}^{(n)} \in \mathbb{K}^{m \times m}$, we define the associated *spectrahedron* (over Puiseux series) as the set

$$\{\mathbf{x} \in \mathbb{K}^n : \mathbf{Q}^{(0)} + \mathbf{x}_1 \mathbf{Q}^{(1)} + \dots + \mathbf{x}_n \mathbf{Q}^{(n)} \text{ is PSD}\}, \quad (2)$$

where “PSD” stands for *positive semidefinite*. (We point out that the definition of positive semidefinite matrices makes sense over any real closed field.) The problem which we are interested in is to determine whether a spectrahedron over Puiseux series is empty or not. This corresponds to the analog of the semidefinite feasibility problem over the field \mathbb{K} . This problem is also related to the standard semidefinite feasibility problem over the field of real numbers associated with the spectrahedra

$$\{\mathbf{x} \in \mathbb{R}^n : \mathbf{Q}^{(0)}(t) + \mathbf{x}_1 \mathbf{Q}^{(1)}(t) + \dots + \mathbf{x}_n \mathbf{Q}^{(n)}(t) \text{ is PSD}\}, \quad (3)$$

for t large enough. Here, $\mathbf{Q}^{(i)}(t)$ stands for the real symmetric matrix obtained by evaluating the entries of $\mathbf{Q}^{(i)}$ at the value t . The relation between the problem over Puiseux series and the one over real numbers is described in the following lemma, and is a consequence of quantifier elimination over real closed fields:

Lemma 1. *The spectrahedron (2) over the field \mathbb{K} is empty if and only if, for t sufficiently large, the spectrahedron (3) over \mathbb{R} is empty.*

In this paper, we consider a slightly different problem which already retains much of the difficulty of the semidefinite feasibility problem over the field \mathbb{K} : given symmetric matrices $\mathbf{Q}^{(1)}, \dots, \mathbf{Q}^{(n)} \in \mathbb{K}^{m \times m}$, determine whether the following *spectrahedral cone*

$$\{\mathbf{x} \in \mathbb{K}_{\geq 0}^n : \mathbf{x}_1 \mathbf{Q}^{(1)} + \dots + \mathbf{x}_n \mathbf{Q}^{(n)} \text{ is PSD}\} \quad (4)$$

is *trivial*, meaning that it is reduced to the zero point. We

refer to [4] for further details on the relation between the original semidefinite feasibility problems and the problems above for spectrahedral cones.

B. Valuation map and tropical semifield

As a nonarchimedean field, \mathbb{K} is equipped with a *valuation* map $\text{val}: \mathbb{K} \rightarrow \mathbb{R} \cup \{-\infty\}$ defined by $\text{val}(x) := \lambda_1$ for $x \neq \mathbf{0}$ as in (1), and $\text{val}(\mathbf{0}) := -\infty$. This valuation map has the following properties:

$$\text{val}(x + y) \leq \max(\text{val}(x), \text{val}(y)) \quad (5)$$

$$\text{val}(xy) = \text{val}(x) + \text{val}(y). \quad (6)$$

We point out that equality holds in (5) as soon as the leading terms of x and y do not cancel. In particular, this condition is satisfied when $x, y \geq \mathbf{0}$.

The *tropical* (or *max-plus*) semifield \mathbb{T}_{\max} can be thought of as the image of $\mathbb{K}_{\geq 0}$ by the valuation map. More precisely, this semifield is defined as the set $\mathbb{T}_{\max} := \mathbb{R} \cup \{-\infty\}$ 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. The reader may consult [18], [19], [20] for more information on the tropical semifield.

The operations above are extended in the usual way to matrices with entries in \mathbb{T}_{\max} . The resulting matrix product is also denoted by \odot . Henceforth, for any $z \in \mathbb{T}_{\max}^n$ and $\beta \in \mathbb{T}_{\max}$, we denote by $\beta + z$ the vector of \mathbb{T}_{\max}^n with entries $\beta + z_i$. Finally, we denote by $\mathbf{0}$ the neutral element for addition in \mathbb{T}_{\max} (i.e., $\mathbf{0} := -\infty$), as well as any vector that has all components equal to $\mathbf{0}$.

We consider \mathbb{T}_{\max} equipped with the topology defined by the distance $(a, b) \mapsto |\exp(a) - \exp(b)|$, and \mathbb{T}_{\max}^n equipped with the product topology. On \mathbb{R}^n we also use *Hilbert’s seminorm* [8], defined by $\|x\|_H := \mathbf{t}(x) - \mathbf{b}(x)$, where $\mathbf{t}(x) := \max_{i \in [n]} x_i$ and $\mathbf{b}(x) := \min_{i \in [n]} x_i$. This seminorm induces a norm on the quotient space of \mathbb{R}^n by the tropical parallelism relation, which is defined by: $x \parallel y$ if, and only if, there exists $\alpha \in \mathbb{R}$ such that $x = \alpha + y$. We denote by $B_H(z, r)$ the Hilbert ball of center $z \in \mathbb{R}^n$ and radius $r \in \mathbb{R}_+$, i.e., $B_H(z, r) := \{x \in \mathbb{R}^n \mid \|x - z\|_H \leq r\}$. We also endow \mathbb{T}_{\max} with the standard order \leq , which is extended to vectors entrywise.

Another algebraic structure that we will use in this paper is the *completed min-plus semiring* $\overline{\mathbb{T}}_{\min}$, which is the set $\mathbb{R} \cup \{+\infty\} \cup \{-\infty\}$ equipped with $(a, b) \mapsto \min\{a, b\}$ as addition and $(a, b) \mapsto a + b$ as multiplication (with the convention $(-\infty) + (+\infty) = (+\infty) + (-\infty) = (+\infty)$). The corresponding matrix product for matrices with entries in $\overline{\mathbb{T}}_{\min}$ will be denoted by \odot' . Given $A \in \mathbb{T}_{\max}^{m \times n}$, the operator $A^\sharp: \overline{\mathbb{T}}_{\min}^m \mapsto \overline{\mathbb{T}}_{\min}^n$ is defined by:

$$A^\sharp(y) := (-A^\top) \odot' y,$$

where A^\top denotes the transpose of A . The operator A^\sharp will be called the *adjoint* of A , being an adjoint in a categorical sense as it satisfies the following property:

$$A \odot x \leq y \text{ if and only if } x \leq A^\sharp(y), \quad (7)$$

for any $y \in \overline{\mathbb{T}}_{\min}^m$ and $x \in \overline{\mathbb{T}}_{\max}^n$.

C. Stochastic zero-sum games with mean payoff

A *stochastic mean payoff game* can be specified by two matrices $A \in \mathbb{T}_{\max}^{m \times n}$ and $B \in \mathbb{T}_{\max}^{m \times q}$, and a row-stochastic matrix $P \in [0, 1]^{q \times n}$, where m, n and q are integers. We assume that every row of B has at least one finite entry, and that the same is true for every column of A .

The rules of the game are as follows. Two players, called Max and Min, control disjoint sets of states, respectively indexed by $[m]$ and $[n]$, and alternatively move a pawn over these states. When the pawn is located on a state $j \in [n]$ of Player Min, she selects a state $i \in [m]$ such that $A_{ij} \neq -\infty$, moves the pawn to state i and pays to Player Max the amount $-A_{ij}$. When the pawn is on a state $i \in [m]$ of Player Max, he selects $k \in [q]$ such that $B_{ik} \neq -\infty$, moves the pawn to state k and receives from Player Min the payment B_{ik} . Finally, at state k the pawn is moved by *nature* to state $l \in [n]$ with probability P_{kl} .

A (*positional*) *strategy* for Player Min is a function $\sigma: [n] \rightarrow [m]$ such that $A_{\sigma(j)j} \neq -\infty$ for all j . Similarly, a (*positional*) *strategy* for Player Max is a function $\tau: [m] \rightarrow [q]$ such that $B_{i\tau(i)} \neq -\infty$ for all i . If Min and Max play according to the strategies σ and τ , and start from state $j_0 \in [n]$, the movement of the pawn is described by a Markov chain over the disjoint union $[m] \uplus [n] \uplus [q]$. Then, the *payoff* (of Player Max) is defined as the average payoff

$$g_{j_0}(\sigma, \tau) := \lim_{N \rightarrow +\infty} \mathbb{E}_{\sigma\tau} \left(\frac{1}{N} \sum_{p=1}^N (-A_{i_p j_p} + B_{i_p k_p}) \right),$$

where $\mathbb{E}_{\sigma\tau}$ refers to the expectation over the trajectories $j_0, i_0, k_0, j_1, i_1, k_1, \dots$, with respect to the probability measure determined by these strategies. The objective of Players Min and Max is to find a strategy which respectively minimizes and maximizes the payoff. Liggett and Lippman [21] showed that there exists a pair of optimal strategies (σ^*, τ^*) , which satisfies

$$g_j(\sigma^*, \tau) \leq g_j(\sigma^*, \tau^*) \leq g_j(\sigma, \tau^*),$$

for every initial state $j \in [n]$ and pair of Min/Max strategies (σ, τ) . In this case, the quantity $g_j(\sigma^*, \tau^*)$ is referred to as the *value* of the game when starting from state j . The state j is said to be *winning* (for Player Max) when the associated value is nonnegative. It is said to be *strictly winning* for the same player if the associated value is positive. A dual terminology applies to Player Min.

With any such a game is associated a *Shapley operator*. The latter is the map $F: \overline{\mathbb{T}}_{\max}^n \rightarrow \overline{\mathbb{T}}_{\max}^n$ defined by

$$F = A^\sharp \circ B \circ P, \quad (8)$$

i.e., $F(x) = A^\sharp(B \odot (Px))$, where Px denotes the usual matrix-vector product of P and x . The finiteness assumptions on the entries of the matrices A, B imply that F preserves both $\overline{\mathbb{T}}_{\max}^n$ and \mathbb{R}^n . It is convenient to consider the vector $v^k := F^k(0)$, for $k \in \mathbb{N}$, where $F^k = F \circ \dots \circ F$ denotes the k -th iterate of F . The j th entry v_j^k represents the value

of the game in finite horizon k with initial state j , associated with the same data. The vector

$$\chi(F) := \lim_{k \rightarrow \infty} v^k / k = \lim_{k \rightarrow \infty} F^k(0) / k$$

is known as the *escape rate vector* of F . We shall recall in Section III why this escape rate does exist. It is known that

$$g_j(\sigma^*, \tau^*) = \chi_j(F) \quad , \quad (9)$$

i.e., the value of the mean payoff game coincides with the limit of the mean value per time unit of the finite horizon game, as the horizon tends to infinity.

In this way, solving a mean payoff games reduces to a dynamical systems issue: computing the escape rate vector of a Shapley operator.

D. Zero-sum games associated with nonarchimedean semidefinite programs

The correspondence between semidefinite feasibility problems for spectrahedral cones and stochastic mean payoff games is given in the next theorem:

Theorem 2. *With every spectrahedral cone \mathcal{C} of the form (4) is associated a stochastic mean payoff game that satisfies the following property: if the valuation of the entries of the matrices $Q^{(i)}$ are chosen in a generic way, then \mathcal{C} is nontrivial if and only if at least one state in the associated game is winning.*

This correspondence is established in [4] by considering the following problem:

$$\mathcal{P}(F): \text{ does there exist } x \in \mathbb{T}_{\max}^n \text{ such that } \\ x \neq \mathbb{0} \text{ and } x \leq F(x)?$$

where $F: \mathbb{T}_{\max}^n \rightarrow \mathbb{T}_{\max}^n$ is the *Shapley operator* of the game associated with the spectrahedral cone \mathcal{C} . This problem is said to be *feasible* when it admits a solution, and *infeasible* otherwise. We point out that $\mathcal{P}(F)$ is feasible if, and only if, the associated stochastic mean payoff game has a winning state. Equivalently, this amounts to the fact that the set

$$\mathcal{S}(F) := \{x \in \mathbb{T}_{\max}^n : x \leq F(x)\} \quad (10)$$

is nontrivial, meaning that it is not reduced to the point $\mathbb{0}$. When the genericity assumption of Theorem 2 is satisfied, $\mathcal{S}(F)$ is precisely the image under the valuation map of \mathcal{C} . Similarly, we can consider the problem:

$$\mathcal{P}_{\mathbb{R}}(F): \text{ does there exist } x \in \mathbb{R}^n \text{ such that } x \ll F(x)?$$

where $y \ll z$ stands for the fact that $y_i < z_i$ for all i . This problem is feasible if, and only if, the set $\mathcal{S}(F)$ is *strictly nontrivial*, meaning that there exists $x \in \mathbb{R}^n$ such that $x \ll F(x)$. This corresponds to the property where every state of the game has a positive value.

The Shapley operator and the associated feasibility problems $\mathcal{P}(F)$ and $\mathcal{P}_{\mathbb{R}}(F)$ provide further conditions under which game algorithms are directly applicable to solve nonarchimedean feasibility problems, disregarding the genericity conditions of Theorem 2. This applies to the case where

the matrices $Q^{(1)}, \dots, Q^{(n)}$ are (*negated*) Metzler matrices, which means that off-diagonal entries are nonpositive:

Theorem 3. *For any Metzler matrices $Q^{(1)}, \dots, Q^{(n)}$, we have:*

- (i) *if $\mathcal{P}(F)$ is infeasible, or equivalently, $\mathcal{S}(F)$ is trivial, then \mathcal{C} is trivial.*
- (ii) *if $\mathcal{P}_{\mathbb{R}}(F)$ is feasible, or equivalently, $\mathcal{S}(F)$ is strictly nontrivial, then \mathcal{C} is strictly nontrivial, meaning that there exists $x \in \mathbb{K}_{>0}^n$ such that the matrix $x_1 Q^{(1)} + \dots + x_n Q^{(n)}$ is positive definite.*

Following the analogy with the classical condition number in linear programming (see, e.g., [22]), we are interested in finding a numerical quantity measuring (the inverse of) the distance to triviality when the instance is nontrivial or to nontriviality when it is trivial. In more details, we define the *condition number* $\text{cond}(F)$ of the above problem $\mathcal{P}(F)$ by:

$$(\inf\{\|u\|_{\infty} : u \in \mathbb{R}^n, \mathcal{P}(u + F) \text{ is infeasible}\})^{-1} \quad (11)$$

if $\mathcal{P}(F)$ is feasible, and

$$(\inf\{\|u\|_{\infty} : u \in \mathbb{R}^n, \mathcal{P}(u + F) \text{ is feasible}\})^{-1} \quad (12)$$

if $\mathcal{P}(F)$ is infeasible (with the convention $0^{-1} = +\infty$). Here, $u + F$ stands for the map $x \mapsto u + F(x)$, where the addition is understood entrywise, and $\|\cdot\|_{\infty}$ stands for the sup-norm, i.e., $\|u\|_{\infty} := \max_i |u_i|$. The condition number $\text{cond}_{\mathbb{R}}(F)$ of the problem $\mathcal{P}_{\mathbb{R}}(F)$ is defined in the same way as in (11) and (12) but replacing \mathcal{P} by $\mathcal{P}_{\mathbb{R}}$.

Remark 4. In [4], the Shapley operator associated with a nonarchimedean SDP feasibility problem is written as $A^{\sharp} \circ P \circ B$ instead of $A^{\sharp} \circ B \circ P$. As there are trivial reductions between the games correspondings to both forms, we consider here a Shapley operator in the latter form. This is more suitable to state the complexity estimates in Section V.

III. PRELIMINARY RESULTS OF NONLINEAR PERRON–FROBENIUS THEORY

In this section, we recall some elements of nonlinear Perron–Frobenius theory which will be used to study the condition numbers introduced above. To do so, we next axiomatize essential properties of the Shapley operators considered in Section II-C, following the “operator approach” of stochastic games [23], [24]

A self-map F of \mathbb{T}_{\max}^n is said to be *order-preserving* when

$$x \leq y \implies F(x) \leq F(y) \text{ for all } x, y \in \mathbb{T}_{\max}^n,$$

and *additively homogeneous* when

$$F(\lambda + x) = \lambda + F(x) \text{ for all } \lambda \in \mathbb{T}_{\max} \text{ and } x \in \mathbb{T}_{\max}^n.$$

We point out that any order-preserving and additively homogeneous self-map F of \mathbb{T}_{\max}^n which preserves \mathbb{R}^n is nonexpansive in the sup-norm, meaning that

$$\|F(x) - F(y)\|_{\infty} \leq \|x - y\|_{\infty} \text{ for all } x, y \in \mathbb{R}^n.$$

Given an order-preserving and additively homogeneous self-map F of \mathbb{T}_{\max}^n , the vectors $x \in \mathbb{T}_{\max}^n$ satisfying

$x \leq F(x)$ can be thought of as the nonlinear analogues of subharmonic functions. A central role in determining the existence of such vectors is played by the limit $\chi(F) = \lim_{k \rightarrow \infty} (F^k(x)/k)$, for $x \in \mathbb{R}^n$. When this limit exists, it can be shown to be independent of the choice of $x \in \mathbb{R}^n$, and so it coincides with the escape rate vector $\chi(F)$ of F . The following theorem of Kohlberg implies that the limit does exist when F preserves \mathbb{R}^n and its restriction to \mathbb{R}^n is *piecewise affine* (meaning that \mathbb{R}^n can be covered by finitely many polyhedra such that F restricted to any of them is affine).

Theorem 5. [25] *A piecewise affine self-map F of \mathbb{R}^n that is nonexpansive in any norm admits an invariant half-line, meaning that there exist $z, w \in \mathbb{R}^n$ such that*

$$F(z + \beta w) = z + (\beta + 1)w$$

for any $\beta \in \mathbb{R}$ large enough. In particular, the escape rate vector $\chi(F)$ exists, and is given by the vector w .

Kohlberg's theorem applies to Shapley operators of stochastic mean payoff games with finite state and action spaces and perfect information. Indeed, the Shapley operator (8) of the game described in Section II-C is order-preserving and additively homogeneous, and its restriction to \mathbb{R}^n is piecewise affine.

For a general order preserving and additively homogeneous self-map of \mathbb{T}_{\max}^n , the escape rate vector may not exist. We can still, however, recover information about the sequences $(F^k(x)/k)_k$ through the Collatz–Wielandt numbers of F . Assuming that F is a continuous, order-preserving, and additively homogeneous self-map F of \mathbb{T}_{\max}^n , we define the *upper Collatz–Wielandt number* of F by:

$$\overline{\text{cw}}(F) := \inf\{\mu \in \mathbb{R} : \exists z \in \mathbb{R}^n, F(z) \leq \mu + z\}, \quad (13)$$

and the *lower Collatz–Wielandt number* of F by:

$$\underline{\text{cw}}(F) := \sup\{\mu \in \mathbb{R} : \exists z \in \mathbb{R}^n, F(z) \geq \mu + z\}. \quad (14)$$

A relation between the escape rate vector and the upper Collatz–Wielandt number is given in the next theorem, which is derived in [9] from a minimax result of Nussbaum [14].

Theorem 6. [9, Lemma 2.8] *Let F be a continuous, order-preserving, and additively homogeneous self-map of \mathbb{T}_{\max}^n . Then,*

$$\lim_{k \rightarrow \infty} \mathbf{t}(F^k(x)/k) = \overline{\text{cw}}(F) = \sup\{\mu \in \mathbb{T}_{\max} : \exists z \in \mathbb{T}_{\max}^n, z \neq 0, F(z) \geq \mu + z\}$$

for any $x \in \mathbb{R}^n$.

As noted in [9, Remark 2.10], the previous result can be dualized when F preserves \mathbb{R}^n .

Corollary 7. *Let F be a continuous, order-preserving, and additively homogeneous self-map of \mathbb{T}_{\max}^n that preserves \mathbb{R}^n . Then,*

$$\lim_{k \rightarrow \infty} \mathbf{b}(F^k(x)/k) = \underline{\text{cw}}(F)$$

for any $x \in \mathbb{R}^n$.

As a consequence, when the escape rate vector exists, we simply have

$$\overline{\text{cw}}(F) = \mathbf{t}(\chi(F)) \quad \text{and} \quad \underline{\text{cw}}(F) = \mathbf{b}(\chi(F)).$$

Specializing this to the case where F is the Shapley operator of a game, the quantities $\overline{\text{cw}}(F)$ and $\underline{\text{cw}}(F)$ respectively correspond to the greatest and smallest values of the states for the mean payoff problem.

In the sequel, we will consider especially the situation in which there is a vector $v \in \mathbb{R}^n$ and a scalar $\lambda \in \mathbb{R}$ such that

$$F(v) = \lambda + v. \quad (15)$$

The scalar λ , which is unique, is known as the *ergodic constant*, and (15) is referred to as the *ergodic equation*. We will denote this scalar by $\rho(F)$ as it appears to be a nonlinear extension of the spectral radius. The vector v is known as a *bias*, or a *potential*. It is easily seen that if F admits such a bias vector, then

$$\underline{\text{cw}}(F) = \overline{\text{cw}}(F) = \rho(F),$$

and the condition that $\rho(F) \geq 0$ means that the game is winning for every initial state. The existence of a bias vector is guaranteed by certain ‘‘ergodicity’’ assumptions [26].

IV. METRIC GEOMETRY PROPERTIES OF CONDITION NUMBERS

A. Condition numbers vs Collatz–Wielandt numbers, and duality

We point out that the definitions given in Section II-D of the condition numbers $\text{cond}(F)$ and $\text{cond}_{\mathbb{R}}(F)$ can be generalized to any continuous, order-preserving, and additively homogeneous self-map F of \mathbb{T}_{\max}^n . The next proposition provides a characterization of these condition numbers in terms of the Collatz–Wielandt numbers of F .

Proposition 8. *Let F be a continuous, order-preserving, and additively homogeneous self-map of \mathbb{T}_{\max}^n . Then,*

$$\text{cond}_{\mathbb{R}}(F) = |\underline{\text{cw}}(F)|^{-1} \quad \text{and} \quad \text{cond}(F) = |\overline{\text{cw}}(F)|^{-1}.$$

We define the dual of the mean payoff game of Section II-D as the one whose Shapley operator is $F^* = (B^\top)^\sharp \circ P \circ A^\top$. The following theorem will allow us to relate $\mathcal{P}_{\mathbb{R}}(F)$ with $\mathcal{P}_{\mathbb{R}}(F^*)$ and $\mathcal{P}(F^*)$.

Theorem 9 (Duality theorem). *Let $F = A^\sharp \circ B \circ P$ and $F^* = (B^\top)^\sharp \circ P \circ A^\top$, where $A \in \mathbb{T}_{\max}^{m \times n}$ has at least one finite entry per column, $B \in \mathbb{T}_{\max}^{m \times q}$ has at least one finite entry per row, and $P \in \mathbb{R}^{q \times n}$ is a row-stochastic matrix. Then,*

$$\overline{\text{cw}}(F^*) = -\underline{\text{cw}}(F).$$

As a consequence of Theorem 9, we obtain:

Corollary 10. *Let $F = A^\sharp \circ B \circ P$ and $F^* = (B^\top)^\sharp \circ P \circ A^\top$, where $A \in \mathbb{T}_{\max}^{m \times n}$ has at least one finite entry per column, $B \in \mathbb{T}_{\max}^{m \times q}$ has at least one finite entry per row, and $P \in \mathbb{R}^{q \times n}$ is a row-stochastic matrix. Then,*

- (i) The condition number of $\mathcal{P}_{\mathbb{R}}(F)$ coincides with the condition number of $\mathcal{P}(F^*)$.
- (ii) Either $\mathcal{P}(F^*)$ is feasible or $\mathcal{P}_{\mathbb{R}}(F)$ is feasible.
- (iii) Only one of the problems $\mathcal{P}_{\mathbb{R}}(F)$ and $\mathcal{P}_{\mathbb{R}}(F^*)$ can be feasible.

B. A geometric characterization of condition numbers

In this section, we study the *inner radius* of the feasible sets of games, that is, given the Shapley operator $F: \mathbb{T}_{\max}^n \rightarrow \mathbb{T}_{\max}^n$ of a game, we study the maximal radius of a Hilbert ball contained in the set (10).

We start with the following simple lemma.

Lemma 11. *Let F be an order-preserving and additively homogeneous self-map of \mathbb{T}_{\max}^n . Assume $z \in \mathbb{R}^n$ and $r \in \mathbb{R}_+$ are such that $r \leq \mathbf{b}(F(z) - z)$. Then, the Hilbert ball $B_H(z, r)$ is contained in $\mathcal{S}(F)$.*

For the condition in the previous lemma to be also necessary for the inclusion to hold, we need an additional assumption on F .

Definition 1. An order-preserving and additively homogeneous self-map F of \mathbb{T}_{\max}^n is said to be *diagonal free* when $F_i(x)$ is independent of x_i for all $i \in [n]$.

Lemma 12. *When F is diagonal free, for any $z \in \mathbb{R}^n$ and $r \in \mathbb{R}_+$ the Hilbert ball $B_H(z, r)$ is contained in $\mathcal{S}(F)$ only if $r \leq \mathbf{b}(F(z) - z)$.*

If F is not diagonal free, the conclusion of Lemma 12 does not necessarily hold, as shown in the next example.

Example 13. Let us consider the order-preserving and additively homogeneous map $F = A^\# \circ B$, where $A = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} -1 & 0 & -1 \end{pmatrix}$. Then, for $z = \begin{pmatrix} 0 & 3 & 0 \end{pmatrix}^\top$, it can be verified that $B_H(z, 3) \subset \mathcal{S}(F) = \{x \in \mathbb{R}^3: x \leq A^\# \circ B(x)\} = \{x \in \mathbb{R}^3: A \odot x \leq B \odot x\}$. However, we have

$$F(x) = \begin{pmatrix} \max\{x_1 - 1, x_2, x_3 - 1\} \\ \max\{x_1 - 1, x_2, x_3 - 1\} \\ \max\{x_1 - 1, x_2, x_3 - 1\} \end{pmatrix},$$

and so $\mathbf{b}(F(z) - z) = 0$.

As a consequence of Lemmas 11 and 12, we obtain:

Theorem 14. *Let F be a diagonal free self-map of \mathbb{T}_{\max}^n . Then, $\mathcal{S}(F)$ contains a Hilbert ball of positive radius if and only if $\underline{\text{cW}}(F) > 0$. Moreover, when $\mathcal{S}(F)$ contains a Hilbert ball of positive radius, the supremum of the radii of the Hilbert balls contained in $\mathcal{S}(F)$ coincides with $\underline{\text{cW}}(F)$.*

Sergeev established in [27] a characterization of the inner radius of *polytropes*, which corresponds to the special case of Theorem 14 in which F is the Shapley operator of a game with only one player and deterministic transitions.

Remark 15. The condition in Theorem 14 is not too restrictive. Indeed, it can be shown that in most cases of interest, if the Shapley operator F of a mean payoff game is not diagonal free, one can construct another mean payoff game

such that its Shapley operator is diagonal free and the inner radius of its feasible set coincides with the one of $\mathcal{S}(F)$.

V. BOUNDING THE COMPLEXITY OF VALUE ITERATION BY THE CONDITION NUMBERS

In this section, F is an order-preserving and additively homogeneous self-map of \mathbb{T}_{\max}^n which preserves \mathbb{R}^n . We also assume that F admits a bias vector $v \in \mathbb{R}^n$, as in (15).

A. A universal complexity bound for value iteration

The most straightforward idea to solve a mean payoff game is probably value iteration: we infer whether or not the mean payoff game is winning by solving the finite horizon game, for a large enough horizon. This is formalized in Fig. 1.

```

1: procedure VALUEITERATION( $F$ )
2:    $\triangleright F$  a Shapley operator from  $\mathbb{R}^n$  to  $\mathbb{R}^n$ 
3:    $\triangleright$  The algorithm will report whether Player Max or Player
      Min wins the mean payoff game represented by  $F$ 
4:    $u := 0 \in \mathbb{R}^n$ 
5:   repeat  $u := F(u)$   $\triangleright$  At iteration  $\ell$ ,  $u = F^\ell(0)$  is the
      value vector of the game in finite horizon  $\ell$ 
6:   until  $\mathbf{t}(u) \leq 0$  or  $\mathbf{b}(u) \geq 0$ 
7:   if  $\mathbf{t}(u) \leq 0$  then return "Player Min wins"
8:   else return "Player Max wins"
9:   end
10: end

```

Fig. 1. Basic value iteration algorithm.

We next show, in Theorem 16, that this value iteration algorithm terminates and is correct, provided the mean payoff of the game is nonzero (i.e., $\rho(F) \neq 0$), and the operations are performed in exact arithmetic. We shall see in Corollaries 23 and 24 that these two restrictions can be eliminated, at the price of an increase of the complexity bound.

It is convenient to introduce the following metric estimate, which represents the minimal Hilbert's seminorm of a bias vector

$$R(F) := \inf \{\|u\|_H : u \in \mathbb{R}^n, F(u) = \rho(F) + u\}.$$

Since F is assumed to have a bias vector $v \in \mathbb{R}^n$, we have $R(F) \leq \|v\|_H < \infty$ and $\rho(F) = \underline{\text{cW}}(F) = \overline{\text{cW}}(F)$. Hence, by Proposition 8,

$$\begin{aligned} |\rho(F)|^{-1} &= |\underline{\text{cW}}(F)|^{-1} = |\overline{\text{cW}}(F)|^{-1} \\ &= \text{cond}_{\mathbb{R}}(F) = \text{cond}(F). \end{aligned}$$

We shall denote by $\text{cond}(F)$ this common quantity.

Note that $|\rho(F)|$ has a remarkable interpretation, as the value of an auxiliary game, in which there is an initial stage, at which Player Max can decide either to keep his role or to swap it with the role of Player Min. Then, the two players play the mean payoff game as usual. Swapping roles amounts to replacing F by the Shapley operator $\tilde{F}(x) := -F(-x)$. Observe also that $\rho(\tilde{F}) = -\rho(F)$. Hence, the value of this modified game is precisely $\max(\rho(F), \rho(\tilde{F})) = |\rho(F)|$.

The following result bounds the complexity of value iteration in terms of $R(F)$ and of the condition number $\text{cond}(F)$.

Theorem 16. *Suppose that the Shapley operator F has a bias vector and that the ergodic constant $\rho(F)$ is nonzero. Then, procedure VALUEITERATION terminates after*

$$N_{vi} \leq R(F) \text{cond}(F)$$

iterations and returns the correct answer.

B. Bounding the condition number and the bias vector of a stochastic mean payoff game

We next bound the condition number $|\rho(F)|^{-1}$, and the metric estimate $R(F)$, when F is a Shapley operator of a stochastic game with perfect information and finite action spaces. As in Section II-C, we assume that

$$F = A^\sharp \circ B \circ P \quad (16)$$

where $A \in \mathbb{T}_{\max}^{m \times n}$ has at least one finite entry per column, $B \in \mathbb{T}_{\max}^{m \times q}$ has at least one finite entry per row, and $P \in \mathbb{R}^{q \times n}$ is a row-stochastic matrix. To obtain explicit bounds, we will assume that the finite entries of the matrices A and B are integers, and we set

$$W := \max \{ |A_{ij} - B_{ih}| : A_{ij} \neq \emptyset, B_{ih} \neq \emptyset, \\ i \in [m], j \in [n]; h \in [q] \}. \quad (17)$$

This is not more special than assuming that the finite entries of A and B are rational numbers (we may always rescale rational payments so that they become integers). We also assume that the probabilities P_{il} are rational, and that they have a common denominator $M \in \mathbb{N}_{>0}$, $P_{il} = Q_{il}/M$, where $Q_{il} \in [M]$ for all $i \in [q]$ and $l \in [n]$.

We say that a state $i \in [q]$ is *nondeterministic* if there are at least two indices $l, l' \in [n]$ such that $P_{il} > 0$ and $P_{il'} > 0$.

The following lemma improves an estimate in [10].

Lemma 17. *Suppose that a Markov chain with n states is irreducible, and that the transition probabilities are rational numbers whose denominators divide an integer M . Let $k \leq n$ denote the number of states with at least 2 possible successors. Let $\pi \in (0, 1)^{n \times n}$ denote the invariant measure of the chain. Then, the least common denominator of the rational numbers $(\pi_i)_{i \in n}$ is not greater than $nM^{\min\{k, n-1\}}$.*

We deduce the following result.

Theorem 18. *Let F be a Shapley operator as above, still supposing that F has a bias vector and that $\rho(F)$ is nonzero. If k is the number of nondeterministic states of the game, then $\text{cond}(F) \leq nM^{\min\{k, n-1\}}$.*

To bound $R(F)$, we use the following idea. For $0 < \alpha < 1$, let v_α denote the value of the discounted game associated with F , meaning that $v_\alpha = F(\alpha v_\alpha)$. Since F represents a zero-sum game with perfect information and finite state and action spaces, it is known that v_α has a Laurent series expansion in powers of $(1-\alpha)$ with a pole of order at most 1 at $\alpha = 1$, see [25]. We can deduce from this that the limit

of $v_\alpha - \rho(F)/(1-\alpha)$ as $\alpha \rightarrow 1^-$ exists and that it is a bias, which we call the *Blackwell bias*. By working out the limit, we arrive at the following estimate.

Theorem 19. *Let F be the Shapley operator in (16), still supposing that it has a bias vector, and let v^* be its Blackwell bias. Then,*

$$R(F) \leq \|v^*\|_{\text{H}} \leq 10n^2WM^{\min\{k, n-1\}}.$$

By combining Theorems 18 and 19, we arrive at the following.

Corollary 20. *Let F be the Shapley operator in (16), still supposing that it has a bias vector and that $\rho(F)$ is nonzero. Then, procedure VALUEITERATION stops after*

$$N_{vi} \leq 10n^3WM^{2\min\{k, n-1\}} \quad (18)$$

iterations and correctly decides which of the two players is winning.

We next show that when specialized to deterministic games, the universal estimate of Theorem 16 gives precisely the complexity bound of Zwick–Paterson [28].

Lemma 21. *Let $F = A^\sharp \circ B$, where $A, B \in \mathbb{T}_{\max}^{m \times n}$, and suppose that there exists $v \in \mathbb{R}^n$ such that $F(v) = \rho(F) + v$. Then*

$$R(F) \leq (n-1)(|\rho(F)| + W),$$

where W is defined as in (17), setting $q = n$.

For deterministic games with integer payments, the mean payoff is given by the average weight of a circuit, which has length at most n . It follows that $|\rho(F)| \geq 1/n$, unless $\rho(F) = 0$. Note also that $\rho(F) \leq W$. By applying Theorem 18, we arrive at the following bound for the number of iterations N_{vi} of the algorithm in Fig. 1.

Corollary 22 (Compare with [28]). *Let $F = A^\sharp \circ B$ be the Shapley operator of a deterministic game, where the finite entries of $A, B \in \mathbb{T}_{\max}^{m \times n}$ are integers. If there exists $v \in \mathbb{R}^n$ such that $F(v) = \rho(F) + v$ with $\rho(F) \neq 0$, then*

$$N_{vi} \leq 2n^2W.$$

The assumption $\rho(F) \neq 0$ that is used in Theorem 16 can be relaxed, by appealing to the following perturbation and scaling argument. This leads to a bound in which the exponents of M and of n are increased.

Corollary 23. *Let $\mu := nM^{\min\{k, n-1\}}$. Then, procedure VALUEITERATION, applied to the perturbed and rescaled Shapley operator $1 + 2\mu F$, terminates unconditionally in at most $21n^4WM^{3\min\{k, n-1\}}$ iterations. If the algorithm reports that Max wins, then Max is winning in the original mean payoff game. If the algorithm reports that Min wins, then Min is strictly winning in the original mean payoff game.*

The algorithm in Fig. 1 can be adapted to work in finite precision arithmetic. Consider the variant of the main body of this algorithm, given in Fig. 2. We assume that each

evaluation of the Shapley operator F is performed with an error of at most $\epsilon > 0$ in the sup-norm.

```

 $u := 0 \in \mathbb{R}^n, \ell := 0 \in \mathbb{N}, \epsilon \in \mathbb{R}_{>0}$ 
repeat
     $u := F(u); \ell := \ell + 1$   $\triangleright$  The operator  $F$  is evaluated in
    approximate arithmetic, so that  $F(u)$  is at most at distance  $\epsilon$  in
    the sup-norm from its true value.
until  $\ell\epsilon + \mathbf{t}(u) \leq 0$  or  $-\ell\epsilon + \mathbf{b}(u) \geq 0$ 
if  $\ell\epsilon + \mathbf{t}(u) \leq 0$  then return "Player Min wins"
end
if  $-\ell\epsilon + \mathbf{b}(u) \geq 0$  then return "Player Max wins"
end
    
```

Fig. 2. Modification of the basic value iteration algorithm to work in finite precision arithmetic.

Corollary 24. *Let F be the Shapley operator in (16), still supposing that it has a bias vector and that $\rho(F)$ is nonzero. Let $\mu := nM^{\min\{k,n-1\}}$. Then, for any $0 < \epsilon \leq \mu^{-1}/3$, value iteration performed with a numerical precision of ϵ at each step (i.e., the algorithm in Fig. 2) stops after*

$$N_{vi} \leq 30n^3WM^2 \min\{k,n-1\} \quad (19)$$

iterations and correctly decides which of the two players is winning.

Observe that (19) is the bound (18) multiplied by 3.

VI. CONCLUDING REMARKS

We introduced a notion of condition number for stochastic mean payoff games, and bounded the complexity of value iteration in terms of this condition number. Whereas condition numbers are familiar for problems over archimedean fields, this leads to an appropriate notion of condition number for *nonarchimedean* semidefinite programming. In particular, our present results explain, at least in part, the perhaps surprising benchmarks of [4], revealing that random nonarchimedean semidefinite feasibility instances with generic valuations can be simpler to solve than their archimedean analogues. In some sense, “good conditioning” provides a quantitative version of “genericity,” and most instances in [4] are well conditioned. This raises the issue of evaluating the condition number on random instances. It is also an interesting question to investigate whether the solution of nonarchimedean SDP could be used, in general, to solve archimedean SDP, and vice versa.

REFERENCES

[1] G. Blekherman, P. A. Parrilo, and R. R. Thomas, *Semidefinite Optimization and Convex Algebraic Geometry*, ser. MOS-SIAM Ser. Optim. Philadelphia, PA: SIAM, 2013, vol. 13.
 [2] B. Gärtner and J. Matoušek, *Approximation Algorithms and Semidefinite Programming*. Heidelberg: Springer, 2012.
 [3] E. de Klerk and F. Vallentin, “On the Turing model complexity of interior point methods for semidefinite programming,” *SIAM J. Optim.*, vol. 26, no. 3, pp. 1944–1961, 2016.
 [4] X. Allamigeon, S. Gaubert, and M. Skomra, “Solving generic nonarchimedean semidefinite programs using stochastic game algorithms,” *J. Symb. Comp.*, vol. 85, pp. 25–54, 2018.
 [5] A. Condon, “The complexity of stochastic games,” *Inform. and Comput.*, vol. 96, no. 2, pp. 203–224, 1992.

[6] 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)*, ser. Lecture Notes in Comput. Sci., vol. 5878. Springer, 2009, pp. 112–121.
 [7] R. D. Nussbaum, “Hilbert’s projective metric and iterated nonlinear maps,” *Mem. Amer. Math. Soc.*, vol. 75, no. 391, 1988.
 [8] G. Cohen, S. Gaubert, and J.-P. Quadrat, “Duality and separation theorems in idempotent semimodules,” *Linear Algebra Appl.*, vol. 379, pp. 395–422, 2004.
 [9] M. Akian, S. Gaubert, and A. Guterman, “Tropical polyhedra are equivalent to mean payoff games,” *Int. J. Algebra Comput.*, vol. 22, no. 1, p. 125001 (43 pages), 2012.
 [10] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino, “A pseudo-polynomial algorithm for mean payoff stochastic games with perfect information and few random positions,” arXiv:1508.03431, 2015.
 [11] V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan, “Cyclic games and finding minimax mean cycles in digraphs,” *Zh. Vychisl. Mat. Mat. Fiz.*, vol. 28, no. 9, pp. 1406–1417, 1988.
 [12] R. Ibsen-Jensen and P. B. Miltersen, “Solving simple stochastic games with few coin toss positions,” in *Algorithms – ESA 2012*, ser. Lecture Notes in Comput. Sci., vol. 7501. Springer, 2012, pp. 636–647.
 [13] D. Grigoriev and V. V. Podolskii, “Tropical effective primary and dual Nullstellensätze,” in *Proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS)*, ser. LIPIcs. Leibniz Int. Proc. Inform., vol. 30. Wadern: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2015, pp. 379–391.
 [14] R. D. Nussbaum, “Convexity and log convexity for the spectral radius,” *Linear Algebra Appl.*, vol. 73, pp. 59–122, 1986.
 [15] M. Bodirsky and M. Mamino, “Max-closed semilinear constraint satisfaction,” in *Proceedings of the 11th International Computer Science Symposium in Russia (CSR)*, ser. Lecture Notes in Comput. Sci., vol. 9691. Springer, 2016, pp. 88–101.
 [16] L. van den Dries and P. Speissegger, “The real field with convergent generalized power series,” *Trans. Amer. Math. Soc.*, vol. 350, no. 11, pp. 4377–4421, 1998.
 [17] X. Allamigeon, S. Gaubert, and M. Skomra, “Tropical spectrahedra,” arXiv:1610.06746, 2016.
 [18] F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat, *Synchronization and Linearity: An Algebra for Discrete Event Systems*. Wiley, 1992.
 [19] P. Butkovič, *Max-linear Systems: Theory and Algorithms*, ser. Springer Monogr. Math. London: Springer, 2010.
 [20] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*, ser. Grad. Stud. Math. Providence, RI: AMS, 2015, vol. 161.
 [21] T. M. Liggett and S. A. Lippman, “Stochastic games with perfect information and time average payoff,” *SIAM Rev.*, vol. 11, no. 4, pp. 604–607, 1969.
 [22] J. Renegar, “Incorporating condition measures into the complexity theory of linear programming,” *SIAM J. Optim.*, vol. 5, no. 3, pp. 506–524, 1995.
 [23] D. Rosenberg and S. Sorin, “An operator approach to zero-sum repeated games,” *Israel J. Math.*, vol. 121, pp. 221–246, 2001. [Online]. Available: <http://dx.doi.org/10.1007/BF02802505>
 [24] A. Neyman, “Stochastic games and nonexpansive maps,” in *Stochastic Games and Applications*, ser. Nato Science Series C, A. Neyman and S. Sorin, Eds. Kluwer Academic Publishers, 2003, vol. 570, ch. 26, pp. 397–415.
 [25] E. Kohlberg, “Invariant half-lines of nonexpansive piecewise-linear transformations,” *Math. Oper. Res.*, vol. 5, no. 3, pp. 366–372, 1980.
 [26] M. Akian, S. Gaubert, and A. Hochart, “Ergodicity conditions for zero-sum games,” *Discrete Contin. Dyn. Syst.*, vol. 35, no. 9, pp. 3901–3931, 2015.
 [27] S. Sergeev, “Max-plus definite matrix closures and their eigenspaces,” *Linear Algebra Appl.*, vol. 421, no. 2, pp. 182–201, 2007.
 [28] U. Zwick and M. Paterson, “The complexity of mean payoff games on graphs,” *Theoret. Comput. Sci.*, vol. 158, no. 1–2, pp. 343–359, 1996.