

THE TROPICAL SHADOW-VERTEX ALGORITHM SOLVES MEAN PAYOFF GAMES IN POLYNOMIAL TIME ON AVERAGE

XAVIER ALLAMIGEON, PASCAL BENCHIMOL, AND STÉPHANE GAUBERT

ABSTRACT. We introduce an algorithm which solves mean payoff games in polynomial time on average, assuming the distribution of the games satisfies a flip invariance property on the set of actions associated with every state. The algorithm is a tropical analogue of the shadow-vertex simplex algorithm, which solves mean payoff games via linear feasibility problems over the tropical semiring $(\mathbb{R} \cup \{-\infty\}, \max, +)$. The key ingredient in our approach is that the shadow-vertex pivoting rule can be transferred to tropical polyhedra, and that its computation reduces to optimal assignment problems through Plücker relations.

1. INTRODUCTION

A *mean payoff game* involves two opponents, “Max” and “Min”, who alternatively move a pawn over the nodes of a weighted bipartite digraph. The latter consists of two classes of nodes, represented by squares and circles, and respectively indexed by $i \in [m]$ and $j \in [n]$ (we use the notation $[k] := \{1, \dots, k\}$). The weight of the arc (i, j) (resp. (j, i)) is a real number denoted by A_{ij} (resp. B_{ij}). We set $A_{ij} := -\infty$ (resp. $B_{ij} := -\infty$) when there is no such arc. An example of game is given in Figure 1.

When the pawn is placed over a square node i , Player Max selects an outgoing arc (i, j) , and then moves the pawn to circle node j and receives the payment A_{ij} from Player Min. Conversely, when the pawn is located on a circle node j , Player Min chooses an arc (j, i') , moves the pawn to square node i' , and Player Max pays her the amount $B_{i'j}$. We assume that A (resp. B) does not have any identically $-\infty$ row (resp. column), so that both players have at least one possible action at every node. The game starts from a circle node $j_0 = j$, and then the two players make infinitely many moves, visiting a sequence $j_0, i_1, j_1, i_2, \dots$ of nodes. The objective of Player Max is to maximize his mean payoff, defined as the liminf of the following ratio when $p \rightarrow +\infty$:

$$(1) \quad (-B_{i_1 j_0} + A_{i_1 j_1} - B_{i_2 j_1} + A_{i_2 j_2} + \dots - B_{i_p j_{p-1}} + A_{i_p j_p})/p.$$

Symmetrically, Player Min aims at minimizing her mean loss, defined as the limsup of (1) when $p \rightarrow +\infty$. Mean payoff games can be defined more generally over arbitrary (not necessarily bipartite) digraphs. This situation can be reduced to the present one [ZP96].

Mean payoff games were studied by Ehrenfeucht and Mycielski in [EM79], where they proved that these games have a value and positional optimal strategies. In more detail, for every initial state $j \in [n]$, there exists a real χ_j and positional strategies $\sigma : [m] \rightarrow [n]$ and $\tau : [n] \rightarrow [m]$, such that: (i) Player Max is certain to win a mean payoff greater than or equal to χ_j with the strategy σ (*i.e.* by choosing the arc $(i, \sigma(i))$ every time the pawn is on a square node $i \in [m]$), (ii) Player Min is sure that her mean loss is less than or equal to χ_j with the strategy τ .

The decision problem associated with mean payoff games consists in determining whether the initial state j is *winning* for Player Max, *i.e.* $\chi_j \geq 0$. The question of the existence of a polynomial time algorithm solving this problem was first raised by Gurvich, Karzanov and Khachiyan in [GKK88]. This problem was shown to be in $\text{NP} \cap \text{co-NP}$ by Zwick and Paterson in [ZP96]. While mean payoff games (and the related class of parity games) received

Date: October 15, 2018.

X. Allamigeon and S. Gaubert are partially supported by the PGM0 program of EDF and Fondation Mathématique Jacques Hadamard. P. Benchimol is supported by a PhD fellowship of DGA and École Polytechnique.




FIGURE 1. An example of mean payoff game. Circle node 1 is a winning initial state for Player Max, while circle node 2 is losing.

an important attention over the past years [GKK88, ZP96, Jur98, GG98, VJ00, BV07, JPZ08, Fri09, BCD⁺11], all the algorithms developed so far are superpolynomial, and the question raised by Gurvich *et al.* is still open.

The present work exploits the equivalence between mean payoff games and linear feasibility problems in tropical algebra. Indeed, it was shown by Akian, Gaubert and Guterman in [AGG12] that the initial state n is winning for Player Max in the game with payments matrices A, B if, and only if, there exists a solution $x \in (\mathbb{R} \cup \{-\infty\})^{n-1}$ to the following system of inequalities:

$$(2) \quad \forall i \in [m], \quad \max(A_{i1} + x_1, \dots, A_{i(n-1)} + x_{n-1}, A_{in}) \\ \geq \max(B_{i1} + x_1, \dots, B_{i(n-1)} + x_{n-1}, B_{in}).$$

The constraints of the form (2) correspond to affine inequalities over the tropical (max-plus) semiring, *i.e.* the set $\mathbb{T} := \mathbb{R} \cup \{-\infty\}$ endowed with the operations $x \oplus y := \max(x, y)$ as addition, and $x \odot y := x + y$ as multiplication. The conjunction of finitely many such inequalities defines a *tropical polyhedron*. Solving a mean payoff game consequently amounts to determine whether a tropical polyhedron is empty, which can be thought of as the tropical analogue of the feasibility problem in linear programming. This is among the motivations leading to the development of a tropical simplex method in a previous work of the authors and Joswig [ABGJ13]. Then, complexity results known for the classical simplex algorithm can be potentially transferred to the tropical setting. However, the main obstacle is to “tropicalize” the pivoting rule involved, *i.e.* to define a tropical pivoting rule which is both compatible with the classical one, and computable, if possible, with a reasonable time complexity. So far [ABGJ14], the only pivoting rules which have been tropicalized are *combinatorial*, *i.e.* they are defined in terms of the neighborhood of the current basic point in the vertex/edge graph of the polyhedron.

Contributions. We prove that the shadow-vertex simplex algorithm can be tropicalized. Following the average-case analysis of the shadow-vertex algorithm of Adler, Karp and Shamir in [AKS87], we deduce that the tropical algorithm solves mean payoff games in polynomial time on average (Section 4). The complexity bound holds when the distribution of the games satisfies a *flip invariance* property. The latter requires that the distribution of the games is left invariant by every transformation consisting, for an arbitrary node of the game, in flipping the orientation of all the arcs incident to this node. Equivalently, the probability distribution on the set of payment matrices A, B is invariant by every transformation consisting in swapping the i th row of A with the i th row of B , or the j th column of A with the j th column of B . Figure 2 provides the illustration of a discrete distribution of games satisfying the property.

The key difficulty in our approach is to show that the computation of the tropical shadow-vertex pivoting rule can be done in polynomial time (Section 3). To this end, we exploit the fact that the shadow-vertex rule is semi-algebraic, *i.e.* it is defined in terms of the signs of finitely




FIGURE 2. A distribution of game satisfying the flip invariance property (with $m = 1$ and $n = 2$), together with the payment matrices. The four configurations are supposed to be equiprobable. The nodes on which the flip operations have been performed are depicted in bold.

many polynomials. Under some genericity conditions, we deduce that the tropical rule reduces to classical linear programs over some Newton polytopes, which are actually (Minkowski sums of) bipartite perfect matching polytopes.

Related work. We are not aware of other works with such average-case complexity results on mean payoff games. In [RBKM10], Roth *et al.* made a probabilistic analysis of $n \times n$ bi-matrix games with weights chosen independently uniformly in $[0, 1]$. Under this assumption, they showed that with high probability (greater than $1 - f(n)$, with $f(n) = o(1/n^c)$ for all constant c), such games admit a pure stationary strategy equilibrium parametrized by only 4 actions. The latter can be consequently found in polynomial time. While this result indicates that complex instances of games are rare, it does not seem to us that it can be used to deduce an average-case complexity bound.

2. PRELIMINARIES

2.1. Tropical arithmetic and generalized Puiseux series. As previously discussed, the triple $(\mathbb{T}, \oplus, \odot)$ forms a semiring, and the elements $\mathbb{0} := -\infty$ and $\mathbb{1} := 0$ correspond to the zero and unit respectively. The tropical operations can be extended to matrices with entries in \mathbb{T} in a usual way, by defining $A \oplus B := (A_{ij} \oplus B_{ij})_{ij}$ and $A \odot B := (\bigoplus_k A_{ik} \odot B_{kj})_{ij}$. We also introduce the exponentiation $x^{\odot k}$ for any $x \in \mathbb{T}$ and $k \in \mathbb{N}$, which is defined as the product $x \odot x \odot \dots \odot x$ of k occurrences of x (if $k = 0$, it is set to $\mathbb{1}$).

Even if the addition \oplus does not have an inverse, it is convenient to consider tropical numbers with a negative sign. The sign is encoded in a formal way, by introducing two copies \mathbb{T}_+ and \mathbb{T}_- of $\mathbb{T} \setminus \{0\}$, respectively consisting of the positive and negative elements. The set \mathbb{T}_\pm of *tropical signed numbers* is defined as $\mathbb{T}_+ \cup \mathbb{T}_- \cup \{0\}$. The elements of \mathbb{T}_+ are simply denoted by elements $a \in \mathbb{T} \setminus \{0\}$, while the elements of \mathbb{T}_- are denoted by $\ominus a$. The *modulus* of $x \in \mathbb{T}_\pm$ is defined as $|x| := a$ if $x = a$ or $x = \ominus a$, and $|0| := 0$. Similarly, the *sign* of $x \in \mathbb{T}_\pm$ is defined by $\text{sign}(x) = +1$ if $x \in \mathbb{T}_+$, $\text{sign}(x) = -1$ if $x \in \mathbb{T}_-$, and $\text{sign}(0) = 0$.

While the tropical addition of signed elements may not be well defined, we can extend the multiplication over $x, y \in \mathbb{T}_\pm$, by defining $x \odot y$ as the unique element of \mathbb{T}_\pm with modulus $|x| \odot |y|$ and sign $(\text{sign}(x) \times \text{sign}(y))$. For instance, $(\ominus 3) \odot 4 = \ominus 7$, and $(\ominus 2) \odot (\ominus 4) = 6$. The exponentiation $x^{\odot k}$ is generalized to the case $x \in \mathbb{T}_\pm$ as well. For any $x \in \mathbb{T}_\pm$, we use the notation $\ominus x$ as a shorthand for the operation $(\ominus 1) \odot x$. The *positive* and *negative parts* x^+ and x^- of an element $x \in \mathbb{T}_\pm$ are defined by $(x^+, x^-) := (x, 0)$ if $x \in \mathbb{T}_+$, $(0, \ominus x)$ if $x \in \mathbb{T}_-$, and $(0, 0)$ if $x = 0$. We extend this notation to vectors and matrices entrywise.

A matrix $M \in \mathbb{T}_\pm^{m \times n}$ is said to be *generic* if the following maximum

$$(3) \quad \max \left\{ |M_{1\sigma(1)}| \odot \dots \odot |M_{n\sigma(n)}| : \sigma \in S_n \right\}$$

is not equal to $\mathbb{0}$, and is attained by only one permutation σ in the symmetric group S_n . A matrix $A \in \mathbb{T}_{\pm}^{m \times n}$ is *strongly non-degenerate* if all of its square submatrices are generic. In particular, the coefficients of A are not null (tropically).

Generalized Puiseux series. Tropical arithmetic can be intuitively illustrated by the arithmetic over asymptotic orders of magnitude. For instance, if we denote by $\Theta(t^a)$ the equivalence class of real functions of t which belong to some interval $[Kt^a, K't^a]$ when $t \rightarrow +\infty$ ($0 < K \leq K'$), we have $\Theta(t^a) + \Theta(t^b) = \Theta(t^{\max(a,b)})$, and $\Theta(t^a) \times \Theta(t^b) = \Theta(t^{a+b})$. We use generalized Puiseux series as a way to manipulate such orders of magnitude in a formal way.

A (real) *generalized Puiseux series* (or *Puiseux series* for short) is a formal series \mathbf{x} in the indeterminate t of the form $x_{\alpha_1}t^{\alpha_1} + x_{\alpha_2}t^{\alpha_2} + \dots$, where $\alpha_i \in \mathbb{R}$, $x_{\alpha_i} \in \mathbb{R} \setminus \{0\}$, and the sequence of the α_i is decreasing, and either finite or unbounded. The set of generalized Puiseux series forms a field that we denote \mathbb{K} . Given a Puiseux series \mathbf{x} as above, the largest exponent α_1 is called the *valuation* of \mathbf{x} , and is denoted by $\text{val}(\mathbf{x})$. By convention, the valuation of the null series $\mathbf{x} = 0$ is defined as $\mathbb{0} = -\infty$. Thus the valuation $\text{val}(\cdot)$ maps \mathbb{K} to \mathbb{T} .

A Puiseux series \mathbf{x} is *positive*, which is denoted by $\mathbf{x} > 0$, if the coefficient $x_{\text{val}(\mathbf{x})}$ of the term with largest exponent in \mathbf{x} is positive. We denote by \leq the total order over \mathbb{K} defined by $\mathbf{x} \leq \mathbf{y}$ if $\mathbf{x} = \mathbf{y}$ or $\mathbf{y} - \mathbf{x} > 0$. Then, we define the *signed valuation* $\text{sval}(\mathbf{x})$ of \mathbf{x} as the element of \mathbb{T}_{\pm} given by $\text{val}(\mathbf{x})$ if $\mathbf{x} > 0$, $\ominus(\text{val}(\mathbf{x}))$ if $\mathbf{x} < 0$, and $\mathbb{0}$ if $\mathbf{x} = 0$. Given $x \in \mathbb{T}_{\pm}$, we also denote by $\text{sval}^{-1}(x)$ the set of Puiseux series \mathbf{x} such that $\text{sval}(\mathbf{x}) = x$. Valuation, signed valuation and its inverse are extended to vectors and matrices coordinate-wise.

As discussed above, the arithmetic operations over \mathbb{K} and \mathbb{T} are related. For instance, for all $\mathbf{x}, \mathbf{y} \geq 0$, we have $\text{val}(\mathbf{x} + \mathbf{y}) = \max(\text{val}(\mathbf{x}), \text{val}(\mathbf{y}))$. Similarly, if $\mathbf{x}, \mathbf{y} \in \mathbb{K}$, then $\text{sval}(\mathbf{x}\mathbf{y}) = \text{sval}(\mathbf{x}) \odot \text{sval}(\mathbf{y})$. More generally, the valuation will be used to transfer “classical” results to the tropical setting. In particular, convex polyhedra and linear programs over generalized Puiseux series essentially behave as over \mathbb{R} (\mathbb{K} is a real-closed field [Mar07], so Tarski’s transfer principle applies). The simplex algorithm can be defined over \mathbb{K} as usual, and the valuation map will allow us to relate it with the tropical simplex algorithm.

General notations. Given a matrix A of dimension $m \times n$, and two subsets $I \subset [m]$ and $J \subset [n]$, we denote by $A_{I \times J}$ the submatrix formed by the rows and the columns of A respectively indexed by $i \in I$ and $j \in J$. If $J = [n]$, we simply use the notation A_I . Similarly, if $i \in [n]$, A_i represents the i -th line of A . The transpose matrix of A is denoted by A^{\top} , and the cardinality of a set I by $|I|$. Given s_1, \dots, s_n , we denote by $\text{diag}(s_1, \dots, s_n)$ (resp. $\text{tdiag}(s_1, \dots, s_n)$) the matrix of dimension $n \times n$, with coefficients s_i on the diagonal, and $\mathbb{0}$ (resp. $\mathbb{0}$) elsewhere.

2.2. The tropical simplex method. The tropical simplex method, introduced in [ABGJ13], allows to solve tropical analogues of linear programming problems:

$$\begin{aligned} \text{LP}(A, b, c) \quad & \text{minimize} \quad c^{\top} \odot x \quad (x \in \mathbb{T}^n) \\ & \text{subject to} \quad x \geq \mathbb{0}, \quad A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^- \end{aligned}$$

where $A \in \mathbb{T}_{\pm}^{m \times n}$, $b \in \mathbb{T}_{\pm}^m$, and $c \in \mathbb{T}_{\pm}^n$. We denote by \mathcal{P} the tropical polyhedron defined by the constraints of $\text{LP}(A, b, c)$. Note that the inequalities $x \geq \mathbb{0}$ are trivially satisfied by any $x \in \mathbb{T}^n$, hence they are superfluous. However, as we shall see, they are involved in the definition of tropical basic points, which is why we need to keep them.

Similarly to the classical simplex method, the principle of the tropical simplex method is to pivot over the (feasible) tropical basic points, while decreasing the objective function $x \mapsto c^{\top} \odot x$. It handles tropical linear programs which satisfy a certain non-degeneracy assumption. Here, we will make the following sufficient assumption:

Assumption A. The matrices $(A \ b)$ and $\begin{pmatrix} A \\ c \end{pmatrix}$ are strongly non-degenerate.

In this setting, a *basis* is a couple (I, J) where $I \subset [m]$, $J \subset [n]$ satisfy $|I| + |J| = n$. Note that we will often manipulate (I, J) through the disjoint union $I \uplus J$ of I and J . By a tropical




FIGURE 3. Intersection of two tropical hyperplanes in general position. The hyperplanes are given by the equalities $x_1 = \max(x_2, 4)$ (blue) and $\max(-5 + x_1, x_2) = 2$ (orange).

analogue of Cramer theorem [Plu90], it can be shown that, under Assumption A, the system

$$(4) \quad \begin{aligned} A_I^+ \odot x \oplus b_I^+ &= A_I^- \odot x \oplus b_I^- \\ x_J &= \mathbb{0} \end{aligned}$$

admits at most one solution in \mathbb{T}^n . If this system admits a solution \bar{x} , the latter is referred to as the *basic point* associated with the basis (I, J) . When \bar{x} belongs to \mathcal{P} , it is called a *feasible basic point*, and (I, J) is a *feasible basis*.

Remark 1. Every equality in the system described in (4) defines a *tropical hyperplane*. Assumption A ensures that these hyperplanes are in general position, so that the intersection of n such hyperplanes is either empty, or reduced to a singleton. We refer to Figure 3 for an illustration. This provides a geometric interpretation of basic points in terms of the arrangement of the tropical hyperplanes associated with the system $x \geq \mathbb{0}$, $A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-$.

The execution of the tropical simplex method on $\text{LP}(A, b, c)$ is related with the execution of the classical simplex method on a lifting linear program over Puiseux series. More precisely, a *lift* of $\text{LP}(A, b, c)$ is a linear program over Puiseux series of the form:

$$\begin{aligned} \text{LP}(A, b, c) \quad & \text{minimize} \quad c^\top x \quad (x \in \mathbb{K}^n) \\ & \text{subject to} \quad x \geq 0, \quad Ax + b \geq 0 \end{aligned}$$

where $A \in \text{sval}^{-1}(A)$, $b \in \text{sval}^{-1}(b)$, and $c \in \text{sval}^{-1}(c)$. Let us denote by \mathcal{P} the convex polyhedron defined by the inequalities of $\text{LP}(A, b, c)$. Then, \mathcal{P} and \mathcal{P} have precisely the same (feasible) bases, and the map $x \mapsto \text{val}(x)$ induces a one-to-one correspondence between their basic points [ABGJ13, Proposition 17]. Besides, if $x^* \in \mathcal{P}$ minimizes the function $x \mapsto c^\top x$, then $\text{val}(x^*) \in \mathcal{P}$ minimizes $x \mapsto c^\top \odot x$. Note that under Assumption A, the linear program $\text{LP}(A, b, c)$ is non-degenerate, in the sense that no minor of $(A \ b)$ and $\begin{pmatrix} A \\ c \end{pmatrix}$ is null.

Moreover, both simplex methods iterate over basic points in the same way. Starting from a basic point of basis (I, J) , they pivot to an adjacent basic point associated with a basis (I', J') such that $I' \uplus J' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k_{\text{in}}\}$, for some $k_{\text{out}} \in I \uplus J$, $k_{\text{in}} \notin I \uplus J$. The element k_{out} is called the *leaving variable*, and is provided by the pivoting rule. The integer k_{in} is uniquely determined when the problem is non-degenerate.¹

As a consequence, under Assumption A, if the classical and tropical simplex methods both start from the basis (I, J) and select the same leaving variable k_{out} , they pivot to the same basis

¹Our terminology follows the one of the dual simplex method. This inversion comes from the fact that our notion of basis actually corresponds to the set of “non-basic” variables in a linear program written with slack variables.

```

1: procedure PIVOT( $(I, J), k_{\text{out}}$ )
2:   for all  $k_{\text{in}} \in ([m] \setminus I) \uplus ([n] \setminus J)$  do
3:     Let  $(I', J')$  be defined by  $I' \uplus J' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k_{\text{in}}\}$ 
4:     if the system  $A_{I'}^+ \odot x \oplus b_{I'}^+ = A_{I'}^- \odot x \oplus b_{I'}^-$ ,  $x_{J'} = \mathbb{0}$  admits a solution  $\bar{x} \in \mathbb{T}^n$ , and  $\bar{x} \in \mathcal{P}$  then
5:       return  $(I', J')$ 
6:     end
7:   done
8: end

```

FIGURE 4. A naive implementation of the tropical pivoting operation

(I', J') . In other words, the tropical simplex method traces the image under the map $\text{val}(\cdot)$ of the path followed by the classical simplex method, provided that they use compatible pivoting rules, *i.e.* at any feasible basis, both rules select the same leaving variable.

We recall that, given a tropical basic point with basis (I, J) and a leaving variable k_{out} , the operation of pivoting to the next tropical basic point can be done in time $O(n(m+n))$ [ABGJ13, Theorem 33]. However, this pivoting operation is limited to basic points with no $-\infty$ coefficients. We present in Figure 4 a simpler, but more expensive, pivoting operation, which handles the general case. The algorithm $\text{PIVOT}((I, J), k_{\text{out}})$ consists in testing all the m possibly entering variables k_{in} in the complement of $I \uplus J$. By the non-degeneracy assumptions, only one variable can lead to a feasible basis. Each candidate basic point can be computed in time $O(n^3)$ (see [Plu90]), and the feasibility can be checked in $O(mn)$ by testing the m inequalities in $A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-$. The total complexity of our pivoting operation is therefore in $O(mn(m+n^2))$.

Remark 2. The geometric interpretation of the pivoting operation in the tropical setting is analogous to the classical setting. Let x and y be two adjacent tropical basic points, respectively associated with the bases (I, J) and (I', J') , so that $I' \uplus J' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k_{\text{in}}\}$ for some $k_{\text{out}} \in I \uplus J$, $k_{\text{in}} \notin I \uplus J$. Let $K \subset [m]$ and $L \subset [n]$ such that $K \uplus L = (I \uplus J) \setminus \{k_{\text{out}}\}$. Then, the set \mathcal{E} of points $z \in \mathcal{P}$ which satisfy the equalities

$$\begin{aligned} A_K^+ \odot z \oplus b_K^+ &= A_K^- \odot z \oplus b_K^- \\ z_L &= \mathbb{0} \end{aligned}$$

is called a *tropical edge* of \mathcal{P} . Geometrically, it coincides with the *tropical segment* between the two points x, y , which consists of the set $\{\lambda \odot x \oplus \mu \odot y : \lambda, \mu \in \mathbb{T}, \lambda \oplus \mu = \mathbb{1}\}$. As shown in [ABGJ13, Proposition 18], this tropical edge is equal to the image under the map $\text{val}(\cdot)$ of the edge \mathcal{E} of the polyhedron \mathcal{P} which connects the two basic points associated with the bases (I, J) and (I', J') .

Example 3. We provide in Figure 5 an example of tropical polyhedron in dimension 2. It is defined by the following five inequalities:

$$\begin{aligned} \max(x, 1+y) &\geq 3 \\ y &\geq \max(-10+x, 1) \\ \max(y, 4) &\geq -3+x \\ 8 &\geq \max(x, 2+y) \\ 4+x &\geq \max(y, 5) \end{aligned}$$

or, equivalently, by the system $x \geq \mathbb{0}$, $A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-$ where:

$$A = \begin{pmatrix} 0 & 1 \\ \ominus(-10) & 0 \\ \ominus(-3) & 0 \\ \ominus 0 & \ominus 2 \\ 4 & \ominus 0 \end{pmatrix} \quad b = \begin{pmatrix} \ominus 3 \\ \ominus 1 \\ 4 \\ 8 \\ \ominus 5 \end{pmatrix}.$$




FIGURE 5. A tropical polyhedron (in gray), two basic points (in blue and orange) and a tropical edge (in green).

Note that the inequalities $x \geq 0$ are never active (our polyhedron is bounded in \mathbb{R}^2), so that all feasible bases are of the form (I, \emptyset) , where $I \subset [5]$ has cardinality 2. The basic points associated with the bases $(\{2, 3\}, \emptyset)$ and $(\{3, 4\}, \emptyset)$ are depicted in blue and orange respectively. The tropical edge between them is represented in green.

3. TROPICALIZING THE SHADOW-VERTEX SIMPLEX ALGORITHM

3.1. The classical shadow-vertex pivoting rule. Given $\mathbf{u}, \mathbf{v} \in \mathbb{K}^n$, the shadow-vertex rule aims at solving the following parametric family of linear programs for increasing values of $\lambda \geq 0$:

$$(5) \quad \begin{aligned} & \text{minimize} && (\mathbf{u}^\top - \lambda \mathbf{v}^\top) \mathbf{x} && (\mathbf{x} \in \mathbb{K}^n) \\ & \text{subject to} && \mathbf{x} \geq 0, \mathbf{A}\mathbf{x} + \mathbf{b} \geq 0 \end{aligned}$$

The vectors \mathbf{u} and \mathbf{v} are respectively called *objective* and *co-objective* vectors. The input of (5) is supposed to satisfy a genericity property. Here, we will assume that no minor of $(\mathbf{A} \ \mathbf{b})$ and $(\mathbf{A}^\top \ \mathbf{u} \ \mathbf{v})$ is null.

When λ is continuously increased from 0, the basic points of \mathcal{P} minimizing the function $\mathbf{x} \mapsto (\mathbf{u}^\top - \lambda \mathbf{v}^\top) \mathbf{x}$ form a sequence $\bar{\mathbf{x}}^0, \dots, \bar{\mathbf{x}}^p$. The shadow-vertex rule amounts to iterate over this sequence. It relies on the reduced cost vectors w.r.t. the objective and co-objective vectors. Given a basis (I, J) , the *reduced cost vector* $\mathbf{y}^{(I, J)} \in \mathbb{K}^{I \uplus J}$ w.r.t. the objective vector \mathbf{u} is defined as the unique solution \mathbf{y} of the system $(\mathbf{A}'_{I \uplus J})^\top \mathbf{y} = \mathbf{u}$, where $\mathbf{A}' = \begin{pmatrix} \mathbf{Id} \\ \mathbf{A} \end{pmatrix}$, and \mathbf{Id} is the identity matrix. The reduced cost vector $\mathbf{z}^{(I, J)}$ w.r.t. the co-objective vector \mathbf{v} can be defined similarly, by replacing \mathbf{u} by \mathbf{v} . Then, at basis (I, J) , the shadow-vertex rule selects the leaving variable $k_{\text{out}} \in I \uplus J$ such that $\mathbf{y}_{k_{\text{out}}}^{(I, J)} > 0$, $\mathbf{z}_{k_{\text{out}}}^{(I, J)} > 0$, and:

$$(6) \quad \mathbf{y}_{k_{\text{out}}}^{(I, J)} / \mathbf{z}_{k_{\text{out}}}^{(I, J)} = \min \{ \mathbf{y}_l^{(I, J)} / \mathbf{z}_l^{(I, J)} : l \in I \uplus J, \mathbf{y}_l^{(I, J)} > 0 \text{ and } \mathbf{z}_l^{(I, J)} > 0 \}.$$

Note that k_{out} is unique under the non-degeneracy assumptions. We refer to [Bor87, Chapter 1.3] for a proof of (6). In the following, we will denote by $\rho_{\text{sv}}(\mathbf{A}, \mathbf{u}, \mathbf{v})$ the function which, given a basis (I, J) , returns the leaving variable provided by the shadow-vertex rule with objective and co-objective vectors \mathbf{u} and \mathbf{v} . If there is no $l \in I \uplus J$ such that $\mathbf{y}_l^{(I, J)} > 0$ and $\mathbf{z}_l^{(I, J)} > 0$, the function $\rho_{\text{sv}}(\mathbf{A}, \mathbf{u}, \mathbf{v})(I, J)$ will be supposed to return the special value **None**. It can be shown that this happens if, and only if, the basic point associated with the basis (I, J) maximizes the co-objective function $\mathbf{x} \mapsto \mathbf{v}^\top \mathbf{x}$ (see [Bor87, Chapter 1.3]).

3.2. Semi-algebraic pivoting rules and their tropical counterpart. In this section, we focus on the problem of finding a tropical pivoting rule $\rho_{\text{sv}}^{\text{trop}}$ compatible with the shadow-vertex

rule ρ_{sv} . More formally, we aim at defining a function $\rho_{\text{sv}}^{\text{trop}}(A, u, v)$ parametrized by a tropical matrix $A \in \mathbb{T}_{\pm}^{m \times n}$, and objective and co-objective vectors $u, v \in \mathbb{T}_{\pm}^n$, such that:

$$\rho_{\text{sv}}^{\text{trop}}(A, u, v)(I, J) = \rho_{\text{sv}}(\mathbf{A}, \mathbf{u}, \mathbf{v})(I, J) \quad \text{for any basis } (I, J),$$

for all $\mathbf{A} \in \text{sval}^{-1}(A)$, $\mathbf{u} \in \text{sval}^{-1}(u)$, and $\mathbf{v} \in \text{sval}^{-1}(v)$. In this case, $\rho_{\text{sv}}^{\text{trop}}$ will be said to be *compatible with ρ_{sv} on the instance (A, u, v)* .

Tropical polynomials. The connection we use between the classical and tropical worlds relies on polynomials over generalized Puiseux series.

Let $P \in \mathbb{K}[X_1, \dots, X_p]$ be a multivariate polynomial, and suppose that it is of the form $\sum_{\alpha \in S} c_{\alpha} X_1^{\alpha_1} \dots X_p^{\alpha_p}$, where $S \subset \mathbb{N}^p$, and $c_{\alpha} \in \mathbb{K} \setminus \{0\}$ for all $\alpha \in S$. The set S is called the *support* of P . We associate a tropical polynomial $\text{trop}(P) \in \mathbb{T}_{\pm}[X_1, \dots, X_p]$ defined as the following formal expression:

$$\text{trop}(P) := \bigoplus_{\alpha \in S} c_{\alpha} \odot X_1^{\odot \alpha_1} \odot \dots \odot X_p^{\odot \alpha_p},$$

with $c_{\alpha} := \text{sval}(c_{\alpha})$ for all $\alpha \in S$. Given $x \in \mathbb{T}_{\pm}^p$, we say that the polynomial $\text{trop}(P)$ *vanishes* on x if the following maximum

$$(7) \quad \max\{|c_{\alpha}| \odot |x_1|^{\odot \alpha_1} \odot \dots \odot |x_p|^{\odot \alpha_p} : \alpha \in S\}$$

is reached at least twice, or is equal to \emptyset . If $\text{trop}(P)$ does not vanish on x , we define $\text{trop}(P)(x) \in \mathbb{T}_{\pm}$ as follows:

$$\text{trop}(P)(x) := c_{\alpha^*} \odot x_1^{\odot \alpha_1^*} \odot \dots \odot x_p^{\odot \alpha_p^*},$$

where α^* is the unique element of S reaching the maximum in (7). The following lemma relates the values of P and $\text{trop}(P)$:

Lemma 4. *Let $x \in \mathbb{T}_{\pm}^p$, and suppose that $\text{trop}(P)$ does not vanish on x . Then, for any $\mathbf{x} \in \text{sval}^{-1}(x)$, we have $\text{sval}(P(\mathbf{x})) = \text{trop}(P)(x)$. In particular, the sign of $P(\mathbf{x})$ is equal to the sign of $\text{trop}(P)(x)$.*

Proof. Let $\mathbf{x} \in \text{sval}^{-1}(x)$. Given $\alpha \in S$, the valuation of the term $c_{\alpha} x_1^{\alpha_1} \dots x_p^{\alpha_p}$ is equal to $|c_{\alpha}| \odot |x_1|^{\odot \alpha_1} \odot \dots \odot |x_p|^{\odot \alpha_p}$. Similarly, its signed valuation is given by $c_{\alpha} \odot x_1^{\odot \alpha_1} \odot \dots \odot x_p^{\odot \alpha_p}$. As the maximum in (7) is reached by only one element $\alpha^* \in S$, we deduce that $c_{\alpha^*} x_1^{\alpha_1^*} \dots x_p^{\alpha_p^*}$ is the unique monomial with largest valuation in $P(\mathbf{x})$. Thus, the signed valuation of $P(\mathbf{x})$ coincides with the signed valuation of $c_{\alpha^*} x_1^{\alpha_1^*} \dots x_p^{\alpha_p^*}$, and is equal to $\text{trop}(P)(x) = c_{\alpha^*} \odot x_1^{\odot \alpha_1^*} \odot \dots \odot x_p^{\odot \alpha_p^*}$. \square

Following this, we can introduce determinants of tropical matrices. Let us define

$$\text{tdet}_n(X) := \bigoplus_{\sigma \in S_n} \text{tsign}(\sigma) \odot X_{1\sigma(1)} \odot \dots \odot X_{n\sigma(n)},$$

where $\text{tsign}(\sigma) := 1$ if the permutation σ is even, $\ominus 1$ otherwise. The polynomial tdet_n is simply denoted tdet when there is no ambiguity. If $M \in \mathbb{T}_{\pm}^{n \times n}$, the *tropical determinant* of M is defined as $\text{tdet}(M)$ when the polynomial tdet does not vanish on M . Note that the latter condition is equivalent to the fact that M is generic. In this case, $\text{tdet}(M)$ can be computed in time complexity $O(n^3)$, by solving an assignment problem over the bipartite graph with node set $[n] \uplus [n]$, in which every arc (i, j) is equipped with the weight $|M_{ij}|$. Indeed, the maximum weight matching provides the unique permutation $\sigma^* \in S_n$ reaching the maximum in (3), and by definition, $\text{tdet}(M)$ is given by $\text{tsign}(\sigma^*) \odot M_{1\sigma^*(1)} \odot \dots \odot M_{n\sigma^*(n)}$.


Example 5. The determinant of (2×2) -matrices is given by the polynomial $\det_2 = X_{1,1}X_{2,2} - X_{1,2}X_{2,1}$, and the corresponding tropical polynomial is $\text{tdet}_2 = (X_{1,1} \odot X_{2,2}) \oplus (\ominus(X_{1,2} \odot X_{2,1}))$. Let us consider the tropical matrix $M := \begin{pmatrix} 3 & \ominus 2 \\ 1 & \ominus 1 \end{pmatrix}$, and let $\mathbf{M} := \begin{pmatrix} 2t^3 + \dots & -t^2 + \dots \\ 4t + \dots & -9t + \dots \end{pmatrix}$ be an arbitrary

lift of M (the dots represent terms of the series which have a smaller exponent, which we left unspecified). The tropical polynomial tdet_2 does not vanish on M , since we have:

$$(8) \quad \max(|M_{1,1}| \odot |M_{2,2}|, |M_{1,2}| \odot |M_{2,1}|) = \max(3 \odot 1, 1 \odot 2) = \max(4, 3) .$$

Hence, the term reaching the maximum in (8) is associated with the monomial $X_{1,1} \odot X_{2,2}$ of tdet_2 , so that $\text{tdet}_2(M) = M_{1,1} \odot M_{2,2} = \ominus 4$. On the other hand, the determinant of \mathbf{M} is of the form $-18t^4 + \dots + 4t^3 + \dots$. Consequently, we indeed have $\text{sval}(\det_2(\mathbf{M})) = \text{tdet}_2(M)$, as expected.

Moreover, the term $|M_{1,1}| \odot |M_{2,2}|$ attaining the maximum in (8) is given by the maximum weight assignment $(1, 1), (2, 2)$ in the following bipartite graph with weights $|M_{ij}|$:



The shadow-vertex rule as a semi-algebraic rule. We claim that the shadow-vertex rule is a *semi-algebraic rule*, in the sense that the leaving variable returned by $\rho_{\text{sv}}(\mathbf{A}, \mathbf{u}, \mathbf{v})(I, J)$ only depends on the current basis (I, J) and on the signs of finitely many polynomials taken on the matrix $\mathbf{M} := \begin{pmatrix} \mathbf{A} \\ \mathbf{u}^\top \\ \mathbf{v}^\top \end{pmatrix}$. To make the notations simpler, we fix a basis (I, J) , and we respectively denote by \mathbf{y} and \mathbf{z} the reduced cost vectors $\mathbf{y}^{(I,J)}$ and $\mathbf{z}^{(I,J)}$. We also define $\bar{J} := [n] \setminus J$.

Let us denote by $P_{K \times L}$ the polynomial given by the $(K \times L)$ -minor of the matrix $X = (X_{ij})$ of formal variables, for any $K \subset [m+2]$ and $L \subset [n]$ such that $|K| = |L|$. For instance, if $K = \{1, 2\}$ and $L = \{3, 4\}$, $P_{K \times L}$ is given by the determinant of the submatrix $\begin{pmatrix} X_{1,3} & X_{1,4} \\ X_{2,3} & X_{2,4} \end{pmatrix}$, i.e. $P_{K \times L} = X_{1,3}X_{2,4} - X_{2,3}X_{1,4}$. For all $l \in I \uplus J$, we define two polynomials Q_l and R_l as follows:

$$\begin{aligned} Q_i &:= P_{(I \setminus \{i\} \cup \{m+1\}) \times \bar{J}} & R_i &:= P_{(I \setminus \{i\} \cup \{m+2\}) \times \bar{J}} & \text{when } i \in I, \\ Q_j &:= P_{(I \cup \{m+1\}) \times (\bar{J} \cup \{j\})} & R_j &:= P_{(I \cup \{m+2\}) \times (\bar{J} \cup \{j\})} & \text{when } j \in J. \end{aligned}$$

Lemma 6. For all $l \in I \uplus J$,

$$\mathbf{y}_l = s_l Q_l(\mathbf{M}) / P_{I \times \bar{J}}(\mathbf{M}), \quad \mathbf{z}_l = s_l R_l(\mathbf{M}) / P_{I \times \bar{J}}(\mathbf{M}),$$

where s_l is a constant in $\{\pm 1\}$ which only depends on the integer l and the sets I and J .

Proof. We restrict our attention to the vector \mathbf{y} (the proof is similar for the vector \mathbf{z}). Recall that \mathbf{y} is given by the following system:

$$\begin{aligned} \mathbf{y}_J + (\mathbf{A}_{I \times J})^\top \mathbf{y}_I &= \mathbf{u}_J \\ (\mathbf{A}_{I \times \bar{J}})^\top \mathbf{y}_I &= \mathbf{u}_{\bar{J}} \end{aligned}$$

By the latter part, for all $i \in I$,

$$\mathbf{y}_i = (-1)^{n - \text{idx}(i)} \det \begin{pmatrix} \mathbf{A}_{(I \setminus \{i\}) \times \bar{J}} \\ (\mathbf{u}_{\bar{J}})^\top \end{pmatrix} / \det(\mathbf{A}_{I \times \bar{J}}),$$

where $\text{idx}(i)$ represents the index of i in the ordered set I . It follows that for all $j \in J$, we have:

$$\mathbf{y}_j \det(\mathbf{A}_{I \times \bar{J}}) = \mathbf{u}_j \det(\mathbf{A}_{I \times \bar{J}}) - \sum_{i \in I} (-1)^{n - \text{idx}(i)} \mathbf{A}_{ij} \det \begin{pmatrix} \mathbf{A}_{(I \setminus \{i\}) \times \bar{J}} \\ (\mathbf{u}_{\bar{J}})^\top \end{pmatrix}.$$

By developing the determinant of $\begin{pmatrix} \mathbf{A}_{I \times (\bar{J} \cup \{j\})} \\ (\mathbf{u}_{\bar{J} \cup \{j\}})^\top \end{pmatrix}$ w.r.t. the column $\begin{pmatrix} \mathbf{A}_{I \times \{j\}} \\ \mathbf{u}_j \end{pmatrix}$, we obtain that:

$$\mathbf{y}_j \det(\mathbf{A}_{I \times \bar{J}}) = (-1)^{n+1 - \text{idx}'(j)} \det \begin{pmatrix} \mathbf{A}_{I \times (\bar{J} \cup \{j\})} \\ (\mathbf{u}_{\bar{J} \cup \{j\}})^\top \end{pmatrix},$$

where $\text{idx}'(j)$ stands for the index of j in the ordered set $\bar{J} \cup \{j\}$. \square

As a consequence of Lemma 6, the properties $\mathbf{y}_l > 0$, $\mathbf{z}_l > 0$ can be tested by determining the signs of $Q_l(\mathbf{M})$, $R_l(\mathbf{M})$ and $P_{I \times \bar{J}}(\mathbf{M})$. Moreover, we have:

$$\mathbf{y}_l / \mathbf{z}_l = Q_l(\mathbf{M}) / R_l(\mathbf{M}) .$$

Hence, the comparison of two ratios $\mathbf{y}_k / \mathbf{z}_k$ and $\mathbf{y}_l / \mathbf{z}_l$ involved in the shadow-vertex rule can be made by evaluating the sign of a polynomial of the form $T_{kl} := Q_k R_l - Q_l R_k$ on the matrix \mathbf{M} . This shows that the shadow-vertex rule is semi-algebraic.

Tropical shadow-vertex rule. Following the previous discussion, we can express $\rho_{\text{sv}}(\mathbf{A}, \mathbf{u}, \mathbf{v})$ as a function defined in terms of the signs of some minors $\det(\mathbf{M}_{K \times L})$, and of the signs of the $T_{kl}(\mathbf{M})$ ($k \neq l$). Given tropical entries (A, u, v) , we simply define $\rho_{\text{sv}}^{\text{trop}}(A, u, v)$ as the same function, in which the minors of \mathbf{M} have been substituted by the corresponding tropical minors of the matrix $M := \begin{pmatrix} A \\ u^\top \\ v^\top \end{pmatrix}$, and the $T_{kl}(\mathbf{M})$ have been replaced by $\text{trop}(T_{kl})(M)$.

In more detail, the function $\rho_{\text{sv}}^{\text{trop}}(A, u, v)(I, J)$ returns the unique element $k_{\text{out}} \in \Lambda$ such that

$$(9) \quad \text{sign}(\text{trop}(T_{k_{\text{out}} l}(M))) = -s_{k_{\text{out}}} s_l \quad \text{for all } l \in \Lambda \setminus \{k_{\text{out}}\} ,$$

where Λ is the set of the elements $l \in I \uplus J$ such that $\text{sign}(\text{trop}(Q_l)(M)) = \text{sign}(\text{trop}(R_l)(M)) = s_l \text{sign}(\text{tdet}(M_{I \times \bar{J}}))$. The latter condition is the tropical counterpart of the conditions $\mathbf{y}_l, \mathbf{z}_l > 0$ in the definition of ρ_{sv} . Equation (9) is the analogue of $\mathbf{y}_{k_{\text{out}}} / \mathbf{z}_{k_{\text{out}}} < \mathbf{y}_l / \mathbf{z}_l$ for all $l \in \Lambda$, $l \neq k_{\text{out}}$. If the set Λ is empty, we set $\rho_{\text{sv}}^{\text{trop}}(A, u, v)(I, J)$ to the special value **None**.

The main result of this section is the following:

Theorem 7. *If the matrix $\begin{pmatrix} A \\ u^\top \\ v^\top \end{pmatrix}$ is strongly non-degenerate, then $\rho_{\text{sv}}^{\text{trop}}$ is compatible with ρ_{sv} on the instance (A, u, v) .*

Moreover, for all bases (I, J) , the leaving variable returned by $\rho_{\text{sv}}^{\text{trop}}(A, u, v)(I, J)$ can be computed in time $O(n^4)$.

Proof. Let $\mathbf{A} \in \text{sval}^{-1}(A)$, $\mathbf{u} \in \text{sval}^{-1}(u)$, and $\mathbf{v} \in \text{sval}^{-1}(v)$. By assumption, the matrix $M = \begin{pmatrix} A \\ u^\top \\ v^\top \end{pmatrix}$ is strongly non-degenerate, so that the sign of every tropical minor $\text{tdet}(M_{K \times L})$ coincides with the sign of the corresponding minor of $\mathbf{M} := \begin{pmatrix} \mathbf{A} \\ \mathbf{u}^\top \\ \mathbf{v}^\top \end{pmatrix}$ by Lemma 4. Consequently, by Lemma 6, the set Λ precisely consists of the elements $l \in I \uplus J$ such that $\mathbf{y}_l > 0$ and $\mathbf{z}_l > 0$. As discussed earlier, each tropical minor can be computed in time $O(n^3)$. Hence, the set Λ can be determined in time $O(n^4)$. It now remains to examine the case of the polynomials $\text{trop}(T_{kl})$, and to show in particular that they do not vanish on M . For the sake of brevity, we restrict to the case $k, l \in I$. The general case can be handled in a similar way.

Let us write the polynomial $\text{trop}(T_{kl})$ under the form $\bigoplus_{\alpha \in S} c_\alpha \odot X^{\odot \alpha}$, where S is the support of T_{kl} (we use the notation $X^{\odot \alpha}$ as a shorthand of $\bigodot_{ij} X_{ij}^{\odot \alpha_{ij}}$). By definition, $\text{trop}(T_{kl})$ does not vanish on M if, and only if, there exists a unique solution to the following maximization problem:

$$(10) \quad \begin{array}{ll} \text{maximize} & |c_\alpha| \odot |M|^{\odot \alpha} \\ \text{subject to} & \alpha \in S \end{array}$$

Observe that the coefficients in T_{kl} are integers. Hence, as elements of the field \mathbb{K} , they are constant Puiseux series, with valuation $0 = \mathbb{1}$. Besides, $M_{ij} \neq 0$ for all (i, j) , thanks to the strong non-degeneracy of M . Then, we can simply rewrite $|c_\alpha| \odot |M|^{\odot \alpha} = \sum_{ij} |M_{ij}| \alpha_{ij}$ for all $\alpha \in S$. As a consequence, Problem (10) can be solved by considering the following classical linear program:

$$(11) \quad \begin{array}{ll} \text{maximize} & \sum_{ij} |M_{ij}| \alpha_{ij} \\ \text{subject to} & \alpha \in \text{New}(T_{kl}) \end{array}$$

```

1: procedure PCBC( $\mathbf{A}, \mathbf{b}$ )
2:    $\mathbf{u} := (\epsilon, \epsilon^2, \dots, \epsilon^n)^\top$   $\triangleright 0 < \epsilon \ll 1$ 
3:    $\bar{\mathbf{x}} := (0, \dots, 0)^\top$ 
4:   for  $k = 1$  to  $m$  do
5:     if  $\mathbf{A}_k \bar{\mathbf{x}} + \mathbf{b}_k < 0$  then
6:       Starting from  $\bar{\mathbf{x}}$ , iterate over the basic points and edges of  $\mathcal{P}^{(k-1)}$  using the rule
        $\rho_{\text{sv}}(\mathbf{A}_{[k-1]}, \mathbf{u}, (\mathbf{A}_k)^\top)$ , until finding a point  $\bar{\mathbf{x}}'$  such that  $\mathbf{A}_k \bar{\mathbf{x}}' + \mathbf{b}_k = 0$ .
7:       if there is no such point  $\bar{\mathbf{x}}'$  then return “Empty”
8:       else  $\bar{\mathbf{x}} := \bar{\mathbf{x}}'$ 
9:     end
10:  done
11:  return “Non-empty”
12: end

```

FIGURE 6. Parametric Constraint-by-Constraint algorithm

where $\text{New}(T_{kl}) \subset \mathbb{R}^{(m+2) \times n}$ is the *Newton polytope* of the polynomial T_{kl} , defined as the convex hull of its support S . Since the set of vertices of the polytope $\text{New}(T_{kl})$ is a subset of S , it is easy to show that $\text{trop}(T_{kl})$ does not vanish on M if, and only if, Problem (11) admits a unique solution α^* . In this case, we have $\alpha^* \in S$, and the sign of $\text{trop}(T_{kl})(M)$ is immediately given by the sign of the term $c_{\alpha^*} \odot M^{\odot \alpha^*}$.

It now remains to check that Problem (11) indeed has a unique solution, and that the latter can be found efficiently. To this aim, we use Plücker quadratic relations (see for instance [GKZ94, Chapter 3]), which provide the identity $T_{kl} = P_{I \times \bar{J}} P_{(I \setminus \{k, l\} \cup \{m+1, m+2\}) \times \bar{J}}$. This implies that the Newton polytope of T_{kl} consists in the Minkowski sum of the two polytopes $\Delta_1 := \text{New}(P_{I \times \bar{J}})$ and $\Delta_2 := \text{New}(P_{(I \setminus \{k, l\} \cup \{m+1, m+2\}) \times \bar{J}})$. As a result, Problem (11) can be decomposed into the following two linear programs:

$$(12) \quad \begin{array}{ll} \text{maximize} & \sum_{ij} |M_{ij}| \alpha_{ij} \\ \text{subject to} & \alpha \in \text{New}(\Delta_i) \end{array} \quad \text{for } i \in \{1, 2\} .$$

More precisely, the set of optimal solutions of Problem (11) is precisely the Minkowski sum of the set of optimal solutions of the two problems given in (12). The polytopes $\text{New}(\Delta_i)$ are bipartite perfect matching polytopes. Consequently, the two problems in (12) correspond to optimal assignment problems, and can be solved in $O(n^3)$. Besides, they both admit a unique solution thanks to the genericity condition on M .

To summarize, $\text{trop}(T_{kl})$ does not vanish on M , and the sign of $\text{trop}(T_{kl})(M)$ can be computed in time $O(n^3)$. By Lemmas 4 and 6, $\text{sign}(\text{trop}(T_{kl})(M)) = -s_k s_l$ if, and only if, $\mathbf{y}_k / \mathbf{z}_k < \mathbf{y}_l / \mathbf{z}_l$. This completes the analysis of the polynomial $\text{trop}(T_{kl})$.

We deduce that $\rho_{\text{sv}}^{\text{trop}}$ and ρ_{sv} are compatible. The output k_{out} of $\rho_{\text{sv}}^{\text{trop}}(A, u, v)(I, J)$ can be computed by determining the smallest element of the set Λ according to the abstract ordering relation \prec defined by $k \prec l \iff \text{sign}(\text{trop}(T_{kl})(M)) = -s_k s_l$. Every comparison has time complexity $O(n^3)$, and so the result can be obtained in time $O(n^4)$. \square

4. AVERAGE-CASE COMPLEXITY OF MEAN PAYOFF GAMES

4.1. Tropicalization of the Parametric Constraint-by-Constraint algorithm. The average-case analysis of [AKS87] applies to the so-called *Parametric Constraint-by-Constraint* (PCBC) algorithm. We first recall the principle of this algorithm. We restrict the presentation to polyhedral feasibility problems, following our motivation to their tropical counterparts and mean payoff games.

The PCBC algorithm is given in Figure 6. Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, and $k \in \{0, \dots, m\}$, we denote by $\mathcal{P}^{(k)}$ the polyhedron defined by the first $n + k$ inequalities of the system $\mathbf{x} \geq 0$, $\mathbf{A}\mathbf{x} + \mathbf{b} \geq 0$. The PCBC algorithm consists in determining by induction on $k = 1, \dots, m$ whether the polyhedron $\mathcal{P}^{(k)}$ is empty. The invariant of the loop from Lines 4 to 10 is that $\bar{\mathbf{x}}$ is

```

1: procedure TROPPCBC( $A, b$ )
2:    $u := (\epsilon, \epsilon^{\odot 2}, \dots, \epsilon^{\odot n})^\top$   $\triangleright -\infty < \epsilon \ll 0$ 
3:    $\bar{x} := (0, \dots, 0)^\top$ 
4:   for  $k = 1$  to  $m$  do
5:     if  $A_k^+ \odot \bar{x} \oplus b_k^+ < A_k^- \odot \bar{x} \oplus b_k^-$  then
6:       Starting from  $\bar{x}$ , iterate over the tropical basic points and edges of  $\mathcal{P}^{(k-1)}$  using the tropical
       rule  $\rho_{sv}^{\text{trop}}(A_{[k-1]}, u, (A_k)^\top)$  until finding a point  $\bar{x}' \in \mathcal{P}^{(k-1)}$  such that  $A_k^+ \odot \bar{x}' \oplus b_k^+ = A_k^- \odot \bar{x}' \oplus b_k^-$ .
7:       if there is no such point  $\bar{x}'$  then return “Empty”
8:       else  $\bar{x} := \bar{x}'$ 
9:     end
10:  done
11:  return “Non-empty”
12: end

```

FIGURE 7. Tropicalization of the PCBC algorithm

the (unique) basic point of $\mathcal{P}^{(k-1)}$ minimizing the function $x \mapsto u^\top x$, where u is an objective vector fixed throughout the whole execution of PCBC. At the k -th iteration of the loop, when \bar{x} does not satisfy the constraint $A_k x + b_k \geq 0$, the simplex algorithm equipped with the shadow-vertex pivoting rule is used. The co-objective vector is set to $(A_k)^\top$. The simplex algorithm thus follows a path in $\mathcal{P}^{(k-1)}$ consisting of basic points and the edges between them. We stop it as soon as it discovers a point $\bar{x}' \in \mathcal{P}^{(k-1)}$ such that $A_k \bar{x}' + b_k = 0$ on the path. This point is obviously a basic point of $\mathcal{P}^{(k)}$. It follows from the definition of the shadow-vertex rule that \bar{x}' minimizes the objective function $x \mapsto u^\top x$ over $\mathcal{P}^{(k)}$. Then, \bar{x}' can be used as a starting point for the execution of the simplex algorithm during the $(k+1)$ -th iteration. If no such point \bar{x}' is discovered, the simplex algorithm stops at a basic point x'' associated with a basis (I'', J'') such that $\rho_{sv}(A_{[k-1]}, u, (A_k)^\top)(I'', J'') = \text{None}$.² In this case, x'' maximizes the function $x \mapsto A_k x$ over $\mathcal{P}^{(k-1)}$, which shows that $A_k x + b_k < 0$ for all $x \in \mathcal{P}^{(k-1)}$. Then, the algorithm returns “Empty”.

The objective vector u is set to $(\epsilon, \epsilon^2, \dots, \epsilon^n)^\top$, where $\epsilon > 0$ is a sufficiently small scalar. Since $u_j > 0$ for all $j \in [n]$, the vector $(0, \dots, 0)^\top$ is a basic point of $\mathcal{P}^{(0)} = (\mathbb{R}_+)^n$ minimizing $x \mapsto u^\top x$.

The PCBC algorithm is still correct when applied on inputs A, b with entries in \mathbb{K} . This suggests to tropicalize it by using the tropical simplex algorithm equipped with the pivoting rule developed in Section 3. This is the purpose of the algorithm TROPPCBC given in Figure 7. Its principle is analogous to PCBC. It manipulates the sequence of tropical polyhedra $\mathcal{P}^{(k)}$ ($0 \leq k \leq m$), which are respectively defined by the first $n+k$ inequalities of the system $x \geq 0$, $A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-$. It also involves an objective vector of the form $u := (\epsilon, \epsilon^{\odot 2}, \dots, \epsilon^{\odot n})^\top$, with $\epsilon < 0$.

Let us describe in more detail the operations performed at Line 6. For each visited basic point x of $\mathcal{P}^{(k-1)}$, the tropical rule $\rho_{sv}^{\text{trop}}(A_{[k-1]}, u, (A_k)^\top)$ determines a variable k_{out} leaving the basis (I, J) associated with x . The tropical simplex algorithm then pivots along the edge \mathcal{E} formed by the points $z \in \mathcal{P}^{(k-1)}$ which activate all the inequalities indexed by $l \in (I \uplus J) \setminus \{k_{\text{out}}\}$ in the system defining $\mathcal{P}^{(k-1)}$ (see Remark 2). There exists a point $\bar{x}' \in \mathcal{E}$ such that $A_k^+ \odot \bar{x}' \oplus b_k^+ = A_k^- \odot \bar{x}' \oplus b_k^-$ if, and only if, the pair (I', J') given by $I' \uplus J' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k\}$ is a feasible basis of $\mathcal{P}^{(k)}$. Indeed, such an \bar{x}' is precisely characterized as the basic point of $\mathcal{P}^{(k)}$ of basis (I', J') . Thus, its existence can be checked in time $O(n(m+n^2))$, as explained in Section 2.2. If there is no such \bar{x}' in \mathcal{E} , we use the algorithm $\text{PIVOT}((I, J), k_{\text{out}})$ to compute the next basis of $\mathcal{P}^{(k-1)}$.

Note that the condition at Line 7 is satisfied when there is no basic point of $\mathcal{P}^{(k-1)}$ to be visited anymore, *i.e.* when the tropical pivoting rule $\rho_{sv}^{\text{trop}}(A_{[k-1]}, u, (A_k)^\top)$ returns **None**.

²As noted in [AKS86, Section 4], if the simplex algorithm encounters a ray of the polyhedron of $\mathcal{P}^{(k-1)}$ on the path, it necessarily finds a point $\bar{x}' \in \mathcal{P}^{(k-1)}$ such that $A_k \bar{x}' + b_k = 0$.

In order to use the tropical shadow-vertex rule on the instances $(A_{[k-1]}, u, (A_k)^\top)$, we verify that the matrix $\begin{pmatrix} A \\ u^\top \end{pmatrix}$ is strongly non-degenerate. The following lemma shows that this property holds assuming that ϵ is small enough:

Lemma 8. *Suppose that the matrix A is strongly non-degenerate, and $\epsilon < n(\min_{ij} |A_{ij}| - \max_{ij} |A_{ij}|)$. Then, the matrix $\begin{pmatrix} A \\ u^\top \end{pmatrix}$ is strongly non-degenerate.*

Proof. Let $M \in \mathbb{T}_\pm^{K \times L}$ be a square submatrix of $\begin{pmatrix} A \\ u^\top \end{pmatrix}$. If M is a submatrix of A , then it is clear that M is generic. Now, suppose that M involves the last line u^\top (i.e. $m+1 \in K$), and that the maximum

$$\max_{k \in K} \left\{ \bigodot_{k \in K} |M_{k\sigma(k)}| : \sigma \text{ is a bijection from } K \text{ to } L \right\}$$

is reached at least by two distinct bijections σ^* and τ^* . If $\sigma^*(m+1) = \tau^*(m+1)$, this immediately shows that the $(K \setminus \{m+1\}) \times (L \setminus \{\sigma^*(m+1)\})$ -submatrix of A is degenerate. Thus, we can suppose that $\sigma^*(m+1)$ and $\tau^*(m+1)$ are distinct, for instance $\sigma^*(m+1) > \tau^*(m+1)$. However, as $\epsilon < 0$, we have:

$$\bigodot_{k \in K} |M_{k\sigma^*(k)}| \leq \sigma^*(m+1)\epsilon + n \max_{ij} |A_{ij}| < \tau^*(m+1)\epsilon + n \min_{ij} |A_{ij}| \leq \bigodot_{k \in K} |M_{k\tau^*(k)}|.$$

In any case, we get a contradiction. We deduce that M is generic. \square

Note that, under the assumptions of Lemma 8, if we choose $\mathbf{A} \in \text{sval}^{-1}(A)$ and $\mathbf{u} \in \text{sval}^{-1}(u)$, no minor of the matrix $\begin{pmatrix} \mathbf{A} \\ \mathbf{u}^\top \end{pmatrix}$ is null. This ensures that the application of the classical shadow-vertex pivoting rule also makes sense in the PCBC algorithm.

Thanks to the compatibility of the tropical shadow-vertex rule with its classical counterpart (Theorem 7), we immediately obtain the following result:

Proposition 9. *Let $A \in \mathbb{T}_\pm^{m \times n}$, $b \in \mathbb{T}_\pm^m$ such that $(A \ b)$ is strongly non-degenerate, and let $\epsilon < n(\min_{ij} |A_{ij}| - \max_{ij} |A_{ij}|)$.*

Then, the algorithm TROP-PCBC correctly determines whether the tropical polyhedron $\{x \in \mathbb{T}^n : A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-\}$ is empty.

Moreover, for all $\mathbf{A} \in \text{sval}^{-1}(A)$, $\mathbf{b} \in \text{sval}^{-1}(b)$ and $\epsilon \in \text{sval}^{-1}(\epsilon)$, the total number of basic points visited by TROP-PCBC(A, b) and by PCBC(\mathbf{A}, \mathbf{b}) are equal.

Proof. First note that the conditions of Theorem 7 are satisfied, thanks to Lemma 8. We are going to show by induction that the algorithms PCBC and TROP-PCBC iterate over the same sequence of bases. Initially, at Line 3, both algorithms start from the basis $(\emptyset, [n])$.

Now, consider the k -th iteration of the loop between Lines 4 and 10. By induction hypothesis, the points \bar{x} and \bar{x} are basic points of $\mathcal{P}^{(k-1)}$ and $\mathcal{P}^{(k-1)}$ respectively, associated with the same basis (I_0, J_0) . The point \bar{x} (resp. \bar{x}) satisfies the condition $\mathbf{A}_k \bar{x} + \mathbf{b}_k \geq 0$ (resp. $A_k^+ \odot \bar{x} \oplus b_k^+ \geq A_k^- \odot \bar{x} \oplus b_k^-$) if, and only if, (I_0, J_0) is a feasible basis of $\mathcal{P}^{(k)}$ (resp. $\mathcal{P}^{(k)}$). As the polyhedra $\mathcal{P}^{(k)}$ and $\mathcal{P}^{(k)}$ have the same feasible bases (see [ABGJ13, Proposition 17]), we deduce that the two conditions $\mathbf{A}_k \bar{x} + \mathbf{b}_k \geq 0$ and $A_k^+ \odot \bar{x} \oplus b_k^+ \geq A_k^- \odot \bar{x} \oplus b_k^-$ are equivalent.

If none of these conditions is satisfied, the two algorithms PCBC and TROP-PCBC execute Line 6 and run the classical and tropical shadow-vertex algorithms. Assume that the two latter algorithms are located at basic points \mathbf{x} and x of $\mathcal{P}^{(k-1)}$ and $\mathcal{P}^{(k-1)}$ respectively, associated with the same basis (I, J) . Note that \mathbf{x} is the final point of the path followed by the classical simplex algorithm if, and only if, x is the final point of the path followed by the tropical simplex algorithm. Indeed, these two statements amount to $\rho_{\text{sv}}(\mathbf{A}_{[k-1]}, \mathbf{u}, \mathbf{A}_k^\top)(I, J) = \text{None}$ and $\rho_{\text{sv}}^{\text{trop}}(A_{[k-1]}, u, A_k^\top)(I, J) = \text{None}$ respectively. The equivalence then follows from Theorem 7.

If \mathbf{x} , or equivalently x , is not the final point of the path, the pivoting rules ρ_{sv} and $\rho_{\text{sv}}^{\text{trop}}$ returns the same leaving variable $k_{\text{out}} \in I \uplus J$, still by Theorem 7. In this case, the classical (resp. tropical) simplex algorithm pivots along the edge \mathcal{E} (resp. \mathcal{E}) formed by the points of $\mathcal{P}^{(k-1)}$ (resp. $\mathcal{P}^{(k-1)}$) which activate the inequalities indexed by $l \in (I \uplus J) \setminus \{k_{\text{out}}\}$. Let $I' \subset [m]$ and $J' \subset [n]$ such that $I' \uplus J' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k\}$.

The existence of a point $\bar{x}' \in \mathcal{E}$ (resp. $\bar{x}' \in \mathcal{E}$) such that $\mathbf{A}_k \bar{x}' + \mathbf{b}_k = 0$ (resp. $A_k^+ \odot \bar{x}' \oplus b_k^+ = A_k^- \odot \bar{x}' \oplus b_k^-$) is equivalent to the fact that the basis (I', J') is a feasible basis of $\mathcal{P}^{(k)}$ (resp. of $\mathcal{P}^{(k)}$). Using again the fact that the polyhedra $\mathcal{P}^{(k)}$ and $\mathcal{P}^{(k)}$ have the same feasible bases, we deduce that the classical simplex algorithm finds a point \bar{x}' such that $\mathbf{A}_k \bar{x}' + \mathbf{b}_k = 0$ when pivoting along the edge \mathcal{E} if, and only if, the tropical simplex algorithm discovers a point $\bar{x}' \in \mathcal{E}$ which satisfies $A_k^+ \odot \bar{x}' \oplus b_k^+ = A_k^- \odot \bar{x}' \oplus b_k^-$. In this case, \bar{x}' and \bar{x}' are basic points (of $\mathcal{P}^{(k)}$ and $\mathcal{P}^{(k)}$ respectively) associated with the same basis (I', J') . If no such points \bar{x}' and \bar{x}' exist, the edge \mathcal{E} is necessarily bounded (see Footnote 2). Thus, the tropical edge $\mathcal{E} = \text{val}(\mathcal{E})$ is also bounded. As a result, the two simplex algorithms both reach new basic points of $\mathcal{P}^{(k-1)}$ and $\mathcal{P}^{(k-1)}$ respectively. These points are associated with the same basis (I'', J'') , given by $I'' \uplus J'' = (I \uplus J) \setminus \{k_{\text{out}}\} \cup \{k_{\text{in}}\}$ for some $k_{\text{in}} \notin I \uplus J$ (the entering variables in the classical and tropical cases are necessarily identical, by unicity).

This completes the proof by induction, and shows the second part of the proposition. The correctness of the algorithm TROPPCBC immediately follows from the correctness of PCBC, and the fact that $\mathcal{P} = \emptyset$ if, and only if, $\mathcal{P} = \emptyset$. \square

Remark 10. As stated in Proposition 9, the scalar ϵ is supposed to be small enough. We point out that there is no need to choose or manipulate ϵ explicitly in the algorithm TROPPCBC. Indeed, as shown in the proof of Theorem 7, ϵ is only involved in optimal assignment problems which arise in the computation of the leaving variable returned by the tropical shadow-vertex rule $\rho_{\text{sv}}^{\text{trop}}(A_{[k-1]}, u, (A_k)^\top)$ at Line 6 ($k \in [m]$). Let us fix k , and let $M := \begin{pmatrix} A_{[k-1]} \\ u^\top \\ A_k \end{pmatrix} \in \mathbb{T}_{\pm}^{(k+1) \times n}$.

The optimal assignment problems to be solved are associated to weighted bipartite graphs, with node sets $K \subset [k+1]$ and $L \subset [n]$ and weights $|M_{hl}|$ for $h \in K$ and $l \in L$. Let G be such a graph. It involves weights with a dependency on ϵ if, and only if, the set K contains the node k , which corresponds to the index of the row vector u^\top in the matrix M . More precisely, the arcs whose weight depends on ϵ are precisely of the form (k, l) , for all $l \in L$. Their respective weights are $\epsilon^{\odot l} = l \times \epsilon$. It is clear that for any sufficiently small $\epsilon < 0$, the (unique) optimal assignment σ^* in G is obtained by mapping the node k to the smallest element l^* of L , and then by solving the optimal assignment problem in the graph G' obtained from G by removing the nodes k and l^* and their incident arcs. Since the weights of G' do not depend on ϵ , the optimal assignment problem in G' can be solved numerically, in a standard way.

In other words, the dependency on ϵ in the optimal assignment problems can be handled in a symbolic way. This is comparable to the “lexicographic” treatment of the scalar ϵ in the PCBC algorithm described in [AKS87, Section 6.1].

4.2. Average-case analysis. Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$ such that no minor of the matrix $(\mathbf{A} \ \mathbf{b})$ is null, the probabilistic analysis of [AKS87] applies to polyhedra of the form

$$\mathcal{P}_{S,S'}(\mathbf{A}, \mathbf{b}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq 0, (S\mathbf{A}S')\mathbf{x} + S\mathbf{b} \geq 0\},$$

where $S = \text{diag}(s_1, \dots, s_m)$, $S' = \text{diag}(s'_1, \dots, s'_n)$, and the s_i and s'_j are i.i.d. entries with values in $\{\pm 1\}$ such that each of them is equal to $+1$ (resp. -1) with probability $1/2$. Equivalently, the 2^{m+n} polyhedra of the form $\mathcal{P}_{S,S'}(\mathbf{A}, \mathbf{b})$ are equiprobable.

Theorem 11 ([AKS87]). *For any fixed choice of \mathbf{A} and \mathbf{b} such that no minor of $(\mathbf{A} \ \mathbf{b})$ is null, provided that ϵ is sufficiently small, the total number of basic points visited by PCBC($S\mathbf{A}S'$, $S\mathbf{b}$) is bounded by $O(\min(m^2, n^2))$ on average.*

It can be verified that the proof of Theorem 11 is still valid when we replace \mathbb{R} by any real-closed field K . Alternatively, it can be shown that Theorem 11 can be expressed as a first-order sentence, so that Tarski’s principle can be used to transpose it to any real-closed field. As a consequence of Proposition 9, the algorithm TROPPCBC should visit a quadratic number of tropical basic points on average. This is the way we translate the result of Adler *et al.* to mean payoff games. The probability distribution of games is expressed over their payments matrices A, B , and must satisfy the following requirements:

- Assumption B.** (i) for all $i \in [m]$ (resp. $j \in [n]$), the distribution of the matrices A, B is invariant by the exchange of the i -th row (resp. j -th column) of A and B .
(ii) almost surely, A_{ij} and B_{ij} are distinct and not equal to \emptyset for all $i \in [m]$, $j \in [n]$. In this case, we introduce the signed matrix $W = (W_{ij}) \in \mathbb{T}_{\pm}^{m \times n}$, defined by $W_{ij} := A_{ij}$ if $A_{ij} > B_{ij}$, and $\ominus B_{ij}$ if $A_{ij} < B_{ij}$.
(iii) almost surely, the matrix W is strongly non-degenerate.

Let us briefly discuss the requirements of Assumption B.

Condition (i) corresponds to the flip invariance property. It handles discrete distributions (see Figure 2) as well as continuous ones. In particular, if the distribution of the payment matrices admits a density function f , Condition (i) can be expressed as the invariance of f by exchange operations on its arguments. For instance, if $m = 1$ and $n = 2$, the flip invariance holds if, and only if, for almost all a_{ij}, b_{ij} , $f(a_{1,1}, a_{1,2}, b_{1,1}, b_{1,2}) = f(b_{1,1}, b_{1,2}, a_{1,1}, a_{1,2}) = f(b_{1,1}, a_{1,2}, a_{1,1}, b_{1,2}) = f(a_{1,1}, b_{1,2}, b_{1,1}, a_{1,2})$.

The requirements $A_{ij}, B_{ij} \neq \emptyset$ for all i, j in Condition (ii) ensure that the flip operations always provide games in which the two players have at least one action to play from every position. The matrix W can be intuitively thought of as a tropical subtraction “ $A \ominus B$ ”, and the conditions $A_{ij} \neq B_{ij}$ ensure that W is well defined. Then, the following result holds:

Lemma 12. *If $A_{ij} \neq B_{ij}$ for all i, j , and W is defined as in Condition (ii) of Assumption B, then the initial state n is winning in the game with matrices A, B if, and only if, $\text{TROPCCBC}(W_{[m] \times [n-1]}, W_{[m] \times \{n\}})$ returns “Non-empty”.*

Proof. Given $a, b, c, d \in \mathbb{T}$ such that $a \neq c$, it can be easily proved that the inequality $\max(a + x_1, b) \geq \max(c + x_1, d)$ over x_1 is equivalent to $b \geq \max(c + x_1, d)$ if $a < c$, and $\max(a + x_1, b) \geq d$ if $a > c$. Using this principle, we deduce that the two systems $A_{[m] \times [n-1]} \odot x \oplus A_{[m] \times \{n\}} \geq B_{[m] \times [n-1]} \odot x \oplus B_{[m] \times \{n\}}$ and $W_{[m] \times [n-1]}^+ \odot x \oplus W_{[m] \times \{n\}}^- \geq W_{[m] \times [n-1]}^- \odot x \oplus W_{[m] \times \{n\}}^+$ are equivalent. As a result, by [AGG12, Theorem 3.2], the initial state n is winning if, and only if, the tropical polyhedron defined by the latter system is non-empty. This provides the expected result, thanks to the first part of Proposition 9. \square

Finally, Condition (iii) is the tropical counterpart of the non-degeneracy assumption used in [AKS87] to establish the average-case complexity bound.

We point out that the set of matrices A, B which do not satisfy the requirements stated in Conditions (ii) and (iii) has measure zero. As a consequence, these two conditions do not impose important restrictions on the distribution of A, B , and they can rather be understood as genericity conditions.

We are now ready to establish our polynomial bound on the average-case complexity of mean payoff games.

Theorem 13. *Under a distribution satisfying Assumption B, TROPCCBC determines in polynomial time on average whether an initial state is winning for Player Max in the mean payoff game with payment matrices A, B .*

Proof. Without loss of generality, we assume that the initial state is the circle node n .

Let us fix two payment matrices A, B satisfying Conditions (ii) and (iii) of Assumption B, and let W as defined in Condition (ii). Starting from the pair (A, B) of matrices, the successive applications of row/column exchange operations precisely yield 2^{m+n-1} different pairs of matrices. In particular, without loss of generality, we can assume that the n -th columns of A and B have not been switched. Then, the pair of matrices that we obtained are of the form $(A_{s,s'}, B_{s,s'})$, where $s \in \{\mathbb{1}, \ominus \mathbb{1}\}^m$, $s' \in \{\mathbb{1}, \ominus \mathbb{1}\}^{n-1}$, and $A_{s,s'}$ and $B_{s,s'}$ are the matrices obtained from A and B respectively, by exchanging the rows i and the columns j such that $s_i = \ominus \mathbb{1}$ and $s'_j = \ominus \mathbb{1}$. The (i, j) -entries of $A_{s,s'}$ and $B_{s,s'}$ are distinct, and so we can define a matrix $W_{s,s'}$ in the same way we have built W from A and B . Observe that $(W_{s,s'})_{[m] \times [n-1]} = S \odot W_{[m] \times [n-1]} \odot S'$ and $(W_{s,s'})_{[m] \times \{n\}} = S \odot W_{[m] \times \{n\}}$, where $S := \text{tdiag}(s_1, \dots, s_m)$ and $S' := \text{tdiag}(s'_1, \dots, s'_{n-1})$.

Thus, by Lemma 12, the 2^{m+n-1} games obtained by the successive flipping operations can be solved by calling the algorithm $\text{TROP}(\text{CBC})(S \odot W_{[m] \times [n-1]} \odot S', S \odot W_{[m] \times \{n\}})$.

Let $\mathbf{W} \in \text{sval}^{-1}(W)$ be a fixed lift of W . Thanks to Condition (iii), no minor of \mathbf{W} is null. Besides, as explained in Remark 10, we do not explicitly manipulate the scalar ϵ in the algorithm $\text{TROP}(\text{CBC})$. Instead, we use a symbolic technique which simulates the behavior of the tropical shadow-vertex rule for any choice of ϵ small enough. This ensures that for all $\epsilon \in \text{sval}^{-1}(\epsilon)$, the conditions of Theorem 11 are satisfied on the instance $(\mathbf{W}_{[m] \times [n-1]}, \mathbf{W}_{[m] \times \{n\}})$. Note that the 2^{m+n-1} instances $(\mathbf{S}\mathbf{W}_{[m] \times [n-1]}\mathbf{S}', \mathbf{S}\mathbf{W}_{[m] \times \{n\}})$, where \mathbf{S}, \mathbf{S}' are diagonal matrices with diagonal coefficients in $\{\pm 1\}$, are respectively lifts of the instances $(S \odot W_{[m] \times [n-1]} \odot S', S \odot W_{[m] \times \{n\}})$. Thanks to the second part of Proposition 9 and Theorem 11, it follows that the total number of visited basic points when calling the algorithm $\text{TROP}(\text{CBC})(S \odot W_{[m] \times [n-1]} \odot S', S \odot W_{[m] \times \{n\}})$ for all $s \in \{\mathbb{1}, \ominus \mathbb{1}\}^m$, $s' \in \{\mathbb{1}, \ominus \mathbb{1}\}^{n-1}$ is bounded by $O(2^{m+n-1} \min(m^2, n^2))$. Moreover, every iteration of the tropical simplex algorithm at Line 6 of $\text{TROP}(\text{CBC})$ consists in determining the leaving variable returned by the tropical shadow-vertex rule, and pivoting to the next basis or computing the point \bar{x}' . The complexity of the former step is $O(n^3)$ by the second part of Theorem 7, and the complexity of the latter step is $O(mn(m+n^2))$. Hence, every iteration is polynomial time. In total, solving the 2^{m+n-1} games associated with the matrices $(A_{s,s'}, B_{s,s'})$ can be done in time $O(2^{m+n-1} mn(m+n^2) \min(m^2, n^2))$.

Let T be the random variable corresponding to the time complexity of our method to solve the game with payment matrices A, B drawn from a distribution satisfying Assumption B. Similarly, given $s \in \{\mathbb{1}, \ominus \mathbb{1}\}^m$, $s' \in \{\mathbb{1}, \ominus \mathbb{1}\}^{n-1}$, let $T_{s,s'}$ be the random variable representing the time complexity to solve the game with matrices $A_{s,s'}, B_{s,s'}$, where A, B are drawn from the latter distribution. Thanks to Condition (i), $\mathbb{E}[T] = \mathbb{E}[T_{s,s'}]$ for all s, s' , and so:

$$\mathbb{E}[T] = \frac{1}{2^{m+n-1}} \times \mathbb{E}\left[\sum_{s,s'} T_{s,s'}\right] \leq \frac{1}{2^{m+n-1}} \times (K 2^{m+n-1} mn(m+n^2) \min(m^2, n^2))$$

for a certain constant $K > 0$. This concludes the proof. \square

5. CONCLUSION

We have defined a tropical analogue of the shadow-vertex simplex algorithm, and shown that every iteration has polynomial time complexity. As a corollary, we have established a polynomial-time average-case result on mean payoff games, based on the analysis of Adler, Karp and Shamir of the classical shadow-vertex algorithm. The main restriction of the model is the flip invariance property. It is an open question whether the tropical approach can be applied with other probabilistic models. In particular, it would be interesting to transfer smoothed complexity results, *e.g.* [ST04], to the tropical setting. The results of Section 3 also suggest that there is a general method to tropicalize any semi-algebraic pivoting rule, based on the characterization of the Newton polytopes involved. This will be addressed in a future work.

Acknowledgments. The authors thank the anonymous reviewers for their helpful comments which contributed to improve the presentation of the paper. We also thank Michael Joswig for many insightful discussions.

REFERENCES

- [ABGJ13] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Tropicalizing the simplex algorithm. Submitted. E-print arXiv:1308.0454, 2013.
- [ABGJ14] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Combinatorial simplex algorithms can solve mean payoff games. Submitted. E-print arXiv:1309.5925, 2014.
- [AGG12] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *Int. J. Algebr. Comput.*, 22(1):125001, 2012.
- [AKS86] I. Adler, R. M. Karp, and R. Shamir. A family of simplex variants solving an $m \times d$ linear program in expected number of pivot steps depending on d only. *Mathematics of Operations Research*, 11(4):570–590, 1986.

- [AKS87] I. Adler, R. M. Karp, and R. Shamir. A simplex variant solving an $m \times d$ linear program in $O(\min(m^2, d^2))$ expected number of pivot steps. *Journal of Complexity*, 3(4):372 – 387, 1987.
- [BCD⁺11] L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J.F. Raskin. Faster algorithms for mean-payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.
- [Bor87] K. H. Borgwardt. *The simplex method: a probabilistic analysis*, volume 1 of *Algorithms and Combinatorics*. Springer Verlag, 1987.
- [BV07] H. Björklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Appl. Math.*, 155:210–229, 2007.
- [EM79] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
- [Fri09] O. Friedmann. An exponential lower bound for the parity game strategy improvement algorithm as we know it. In *LICS*, pages 145–156. IEEE, August 2009.
- [GG98] S. Gaubert and J. Gunawardena. The duality theorem for min-max functions. *C. R. Acad. Sci. Paris.*, 326, Série I:43–48, 1998.
- [GKK88] V. A. Gurvich, A. V. Karzanov, and L. G. Khachivan. Cyclic games and an algorithm to find mini-max cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28(5):85–91, 1988.
- [GKZ94] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Mathematics: Theory & Applications. Birkhäuser Boston, 1994.
- [JPZ08] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM Journal on Computing*, 38(4):1519–1532, 2008.
- [Jur98] M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
- [Mar07] T. Markwig. A field of generalised puiseux series for tropical geometry. *Preprint arXiv:0709.3784*, 2007.
- [Plu90] M. Plus. Linear systems in $(\max, +)$ -algebra. In *Proceedings of the 29th Conference on Decision and Control*, Honolulu, Dec. 1990.
- [RBKM10] A. Roth, M. F. Balcan, A. Kalai, and Y. Mansour. On the equilibria of alternating move games. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 805–816, Philadelphia, PA, USA, 2010. SIAM.
- [ST04] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, May 2004.
- [VJ00] J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In E. A. Emerson and A. P. Sistla, editors, *Computer Aided Verification*, volume 1855 of *LNCS*, pages 202–215. Springer Berlin Heidelberg, 2000.
- [ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.

INRIA AND CMAP, ÉCOLE POLYTECHNIQUE, CNRS, 91128 PALAISEAU CEDEX FRANCE
 E-mail address: `firstname.lastname@inria.fr`