SIAM J. DISCRETE MATH.
Vol. 29, No. 2, pp. 751–795

TROPICALIZING THE SIMPLEX ALGORITHM*

XAVIER ALLAMIGEON†, PASCAL BENCHIMOL†, STÉPHANE GAUBERT†, AND MICHAEL JOSWIG‡

Abstract. We develop a tropical analogue of the simplex algorithm for linear programming. In particular, we obtain a combinatorial algorithm to perform one tropical pivoting step, including the computation of reduced costs, in $O(n(m+n))$ time, where m is the number of constraints and n is the dimension.

Key words. tropical geometry, linear programming, simplex method

AMS subject classifications. 14T05, 90C05

DOI. 10.1137/130936464

1. Introduction.

The tropical semiring $(\mathbb{T}, \oplus, \odot)$ is the set $\mathbb{T} = \mathbb{R} \cup \{-\infty\}$ endowed with the two operations $a \oplus b = \max(a, b)$ and $a \odot b = a + b$. We are interested in the tropical equivalent of linear programming. In other words, our goal is to give an algorithm for minimizing a tropical linear form $\max(c_1 + x_1, \ldots, c_n + x_n)$ over a tropical polyhedron. The latter is the set of solutions $x \in \mathbb{T}^n$ of finitely many inequalities of the form

$$\max(a_1 + x_1, \ldots, a_n + x_n, a_{n+1}) \geqslant \max(b_1 + x_1, \ldots, b_n + x_n, b_{n+1}).$$

All the coefficients a_j, b_j, c_j are elements of $\mathbb{T}$. An example is depicted in Figure 1.

Several avenues lead to this research. First, the classical simplex method belongs to the most relevant algorithms, for both its applicability and its theoretical implications. So it is natural to explore variants and derivations, including tropical ones. In the form that we are studying this leads to a class of minmax problems which are also interesting from a purely complexity-theoretic point of view. In [AGG12] it is shown that a tropical analogue of the feasibility problem in linear optimization is polynomial-time equivalent to deciding which player has a winning strategy in a mean-payoff game. The latter decision problem is among the few problems in NP as well as co-NP (see Zwick and Paterson [ZP96]) for which no polynomial-time algorithm is known. This game-theoretic perspective leads to a second approach to tropical linear programming. A third train of thought is more geometrical. Viro suggested investigating the tropical aspects of real algebraic geometry [Vir01]. Nonetheless, the main focus of tropical geometry so far concerns the tropicalization of algebraic varieties which are defined over the complex numbers (or Puiseux series with complex coefficients). More recently, however, the tropicalization of real semi-algebraic sets has been

*Received by the editors September 10, 2013; accepted for publication (in revised form) November 3, 2014; published electronically April 16, 2015.
http://www.siam.org/journals/sidma/29-2/93646.html

†INRIA and CMAP, École Polytechnique, 91128 Palaiseau Cedex, France (xavier.allamigeon@inria.fr, pascal.benchimol@polytechnique.edu, stephane.gaubert@inria.fr). The second author's research was supported by a Ph.D. fellowship of DGA and École Polytechnique. The first and third authors are partially supported by the PGMO program of EDF and Fondation Mathématique Jacques Hadamard.

‡Institut für Mathematik, TU Berlin, Str. des 17. Juni 136, 10623 Berlin, Germany (joswig@math.tu-berlin.de). This author's research was supported by Einstein Foundation Berlin and the German Research Foundation (DFG).

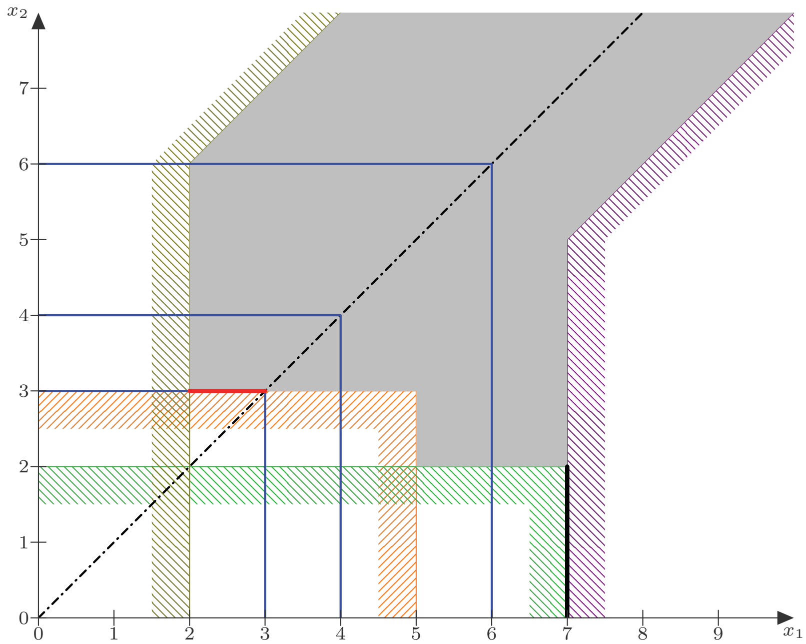

Fig. 1. *A tropical linear program. The feasible set is the union of the gray shaded area with the thick black halfline. Three level sets for the objective function* $\max(x_1, x_2)$ *are depicted in blue. The thick red segment is the set of optima.*

studied by Alessandrini [Ale13]. In this vein our work seeks to contribute to understanding the tropicalizations of the most simple semialgebraic sets: convex polyhedra. A related motivation arises from linear programming over ordered fields, the complexity of which is a well-known open question [Meg89, section 2]. Ordered fields arise naturally when dealing with perturbations of classical linear programs [Jer73, FAA02].

Tropical polyhedra or tropically convex sets have appeared in different guises in the works of several authors, including [Zim77, CG79, LMS01, CGQ04, BH04]; the present work is specially motivated by the approach of Develin and Sturmfels [DS04], in which tropical polyhedra are studied by combinatorial means, and by the further work of Develin and Yu [DY07], who showed that tropical polyhedra are precisely the images by the valuation of (convex) polyhedra over the field of Puiseux series.

With this in mind, the most natural approach for tropical linear programming probably is to do linear programming over real Puiseux series and to tropicalize, i.e., to devise a method which traces the valuation of the path followed by the simplex algorithm over real Puiseux series. This is exactly what we do here. What makes our algorithm interesting is that the method itself does not manipulate Puiseux series (explicit lifts to real Puiseux series are not needed). Instead it directly processes the tropical input and “stays tropical” throughout the computations. In this way, the arithmetical operations remain elementary.

In order to make our ideas more apparent, and to avoid technical details which are too cumbersome to attack in a direct fashion, in the present paper we assume that our tropical linear program is primally and dually nondegenerate. Further, we assume that each point in the feasible region has finite coordinates only. Any tropical linear program satisfying these properties will be called *standard*; see Assumptions

Algorithm 1. Phase II tropical simplex algorithm.

Input: A matrix $A \in \mathbb{T}_{\pm}^{m \times n}$, a column vector $b \in \mathbb{T}_{\pm}^{m}$, an unsigned row vector $c \in \mathbb{T}^n$. A tropical basic point x^I of $\mathcal{P}(A, b)$, and the corresponding set $I \subset [m]$.
Output: A tropical basic point of $\mathcal{P}(A, b)$ that is minimal with respect to c.

compute the tropical reduced costs y associated with I
while y has a tropically negative entry **do**
    choose $i_{\mathsf{out}} \in I$ such that $y_{i_{\mathsf{out}}}$ is tropically negative
    $K \leftarrow I \setminus \{i_{\mathsf{out}}\}$
    pivot along the tropical edge $\mathcal{E}_K$ to the tropical basic point $x^{I'}$ for a set of the form $I' = K \cup \{i_{\mathsf{ent}}\}$
    $I \leftarrow I'$
    compute the tropical reduced costs y associated with I
return x^I

4, 5, and 6 below. We defer all ramifications which come from looking at degenerate input or infinite coefficients to a subsequent second paper. Our main result is the following theorem.

Theorem 1.1. *Consider a standard tropical linear program with n variables and m inequalities. Then, the tropical simplex algorithm (Algorithm 1) terminates and returns an optimal solution for any tropical pivoting rule. Every iteration (pivoting and computing reduced costs) can be done in time $O(n(m + n))$. Moreover, the algorithm traces the image by the valuation map of the path followed by the classical simplex algorithm applied to any lift of this program to the field of real Puiseux series, with a compatible pivoting rule.*

In particular, under the assumptions of Theorem 1.1, linear programs over Puiseux series are implicitly solved by the tropical simplex algorithm. By definition, a tropical pivoting rule selects a variable of tropically negative reduced cost. A classical pivoting rule is said to be compatible with the former tropical pivoting rule if they select the same variables. Tropical pivot rules are the topic of section 4.

Our tropical simplex algorithm relies on several tools of independent interest. For instance, Corollary 3.6 shows that, again under the general position assumption, the cells of an arrangement of hyperplanes over the field of real Puiseux series are in one-to-one correspondence with the cells of the arrangement of the associated tropical hyperplanes. This leads to the notion of *tropical basic points* and *tropical edges* of a system of tropical affine inequalities. Unlike the classical case, a tropical basic point may not be tropically extreme; see Proposition 3.11 and Remark 3.12 below. This stems from the lack of a good notion for a general "face" of a tropical polyhedron; see the discussions in [Jos05, DY07]. This is related to competing notions of rank [DSS05, AGG09].

A fundamental discrepancy to the classical simplex algorithm is that a tropical edge in $\mathbb{T}^n$ may have a more complex geometrical structure as, indeed, it consists of up to n ordinary segments. These segments can be determined from *tangent digraphs*, which encode a local description of a tropical polyhedron; tangent digraphs were initially introduced in the form of directed hypergraphs in [AGG13]. The cornerstone of the tropical simplex algorithm is a new combinatorial characterization of the tangent digraph (Proposition 4.3) at a point inside a tropical edge. In particular, this entails an incremental computation of tangent digraphs from one ordinary segment to another (Proposition 4.9), leading to an $O(n(m+n))$ time method for one full pivoting step; see Theorem 4.13. Finally, we define the tropical reduced cost vector, which allows one to certify the optimality of a given basic point. We show that the vector of reduced costs can be computed by solving a system of signed tropical linear equations and that the running time of this step is also bounded by $O(n(m+n))$; see Theorem 5.7.

Let us finally point out some related work. The study of the analogues of linear programs over ordered semirings was undertaken in the book [Zim81]; in particular, a duality theorem for a special class of linear programs can be found there. The idea of looking for analogues of convex programming results over "extremal" (a variant of tropical) structures is also apparent in [Zim76]. Several recent works have proposed algorithms to solve various tropical programming problems. In [BA08], a dichotomy algorithm is developed, allowing one to solve tropical linear programming problems by a reduction to linear feasibility problems. In [GKS12], more general linear-fractional programming problems are studied, in which one maximizes the difference of two tropical linear forms. A policy iteration algorithm based on a parametric mean payoff game is given there. These policies seem to have interesting connections with basic points. However, our present approach leads to a fundamentally different method: we move along edges in the graph of the tropical polytope, whereas policy iteration type algorithms often take "great leaps" in the same graph; also, one iteration of the present algorithm takes only $O(n(m+n))$ time, whereas every iteration in [GKS12] requires us to solve a mean payoff game. Yet another different class of algorithms for solving tropical linear feasibility problems relies on cyclic projection [CGB03, GS07, AGNS11]. Recall also that the tropical linear feasibility problem is equivalent to mean payoff games, for which a number of algorithms are available, like pumping [GKK88], value iteration [ZP96, AGG12], or policy iteration [GG98, DG06, BV07, Cha09]. A reduction of mean payoff games to classical linear programs with exponentially large coefficients is established in [Sch09]. The asymptotic simplex method developed in [FAA02] solves arbitrary linear programs on Laurent series (which is sufficient for tropical linear programs with rational coefficients). Each iteration of their method requires $O(s(m+n)^2)$ operations, where $s \leqslant (m+n)$ is the maximum taken over the valuations of all Puiseux series arising during the computation. Our tropical simplex algorithm shows a better complexity per iteration since the factor of s is dispensed with, but the approach of [FAA02] does not require any genericity assumptions.

This paper is organized as follows. Section 2 describes our notation and collects the relevant known facts about convex polyhedra over real Puiseux series. For the reader's convenience we introduce a running example which we refer to throughout this paper. In section 3 we characterize the key players in our algorithms: tropical basic points and tropical edges. The core of our paper is section 4, where we describe the tropical pivot. Section 5 discusses tropical reduced costs. Finally, Theorem 1.1 is proved in section 6.

Our algorithm for solving tropical linear programs is outlined in Algorithm 1. It directly corresponds to Phase II of the classical simplex method over real Puiseux series. Phase II starts from a given tropical basic point and proceeds along improving edges toward an optimal tropical basic point, The Phase I problem, to find a first tropical basic point, will be addressed in a sequel to this work. While classically Phase I can be reduced to Phase II, in general, this requires solving a degenerate linear program. As explained above, this is out of the scope of the present paper.

2. Preliminaries.

2.1. Tropical arithmetic. The domain for our computations is the set $\mathbb{T} = \mathbb{R} \cup \{-\infty\}$. The neutral elements for the tropical "addition" and "multiplication" are $\mathbb{0} := -\infty$ and $\mathbb{1} := 0$, respectively. The usual definition of matrix operations carries over to tropical matrices. Given two matrices $A = (a_{ij})$ and $B = (b_{ij})$, we denote by $A \oplus B$ and $A \odot B$ the matrices with entries $a_{ij} \oplus b_{ij}$ and $\bigoplus_k a_{ik} \odot b_{kj}$, respectively. We also denote by $A^\top$ the transpose of the matrix A, by A_i the ith row of A, and by A_I the submatrix of A formed from the rows $i \in I$. For the sake of simplicity, we identify vectors of size n with $n\times 1$-matrices. Given $a = (a_{1j}) \in \mathbb{T}^{1\times n}$ and $x \in \mathbb{T}^n$, we denote by $\arg\max(a \odot x)$ the set of indices $i \in [n] = \{1, \dots, n\}$ attaining the maximum in

$$a \odot x = \max_{j\in[n]}(a_{1j} + x_j).$$

The usual total order $\leqslant$ on $\mathbb{R}$ extends to $\mathbb{T}$. This induces a partial ordering of tropical vectors by entrywise comparisons. The topology induced by the order makes $(\mathbb{T}, \oplus, \odot)$ a topological semiring.

In the following, we will think of the n-fold product space $\mathbb{T}^n$ as a semimodule over $\mathbb{T}$, where scalars act tropically on vectors by $(\lambda, x) \mapsto \lambda \odot x := (\lambda + x_1, \dots, \lambda + x_n)$ and the tropical vector addition is $(x, y) \mapsto x \oplus y := (\max(x_1, y_1), \dots, \max(x_n, y_n))$.

2.1.1. Signed tropical numbers. It will be convenient to use the set of *signed tropical numbers* [Plu90], denoted here by $\mathbb{T}_\pm$. The latter set consists of two copies of $\mathbb{T}$, called the set of *positive tropical numbers* and the set of *negative tropical numbers*, respectively. These two copies are glued by identifying the element $\mathbb{0}$. Positive and negative tropical numbers are written as a and $\ominus a$, respectively, for some $a \in \mathbb{T}$. By definition, the numbers a and $\ominus a$ are different, unless $a = \mathbb{0}$. Their *sign* is $\operatorname{sign}(a) = 1$ and $\operatorname{sign}(\ominus a) = -1$ when a is not $\mathbb{0}$ and $\operatorname{sign}(\mathbb{0}) = 0$. The *modulus* of $x \in \{a, \ominus a\}$ is defined as $|x| := a$. The multiplication $x \odot y$ of two elements $x, y \in \mathbb{T}_\pm$ yields the element whose modulus is $|x| + |y|$ and whose sign is the product $\operatorname{sign}(x)\operatorname{sign}(y)$. The *positive part* and the *negative part* of an element $x \in \mathbb{T}_\pm$ are the tropical numbers x^+ and x^- defined by

$$x^+ = \begin{cases} |x| & \text{if } x \text{ is positive,} \\ \mathbb{0} & \text{otherwise,} \end{cases} \qquad x^- = \begin{cases} \mathbb{0} & \text{if } x \text{ is positive,} \\ |x| & \text{otherwise.} \end{cases}$$

Modulus, positive part, and negative part extend to matrices entrywise. It was shown in [Plu90] that signed tropical numbers can be embedded in a semiring, called the *symmetrized tropical semiring*. Indeed, the sum of two signed tropical numbers with opposite signs but identical modulus cannot be defined as a signed tropical number; one needs to enlarge $\mathbb{T}_\pm$ with a third type of elements, called *balanced elements*, to represent such sums. We will defer the discussion of the symmetrized tropical semiring until section 5.1, since the additional technicalities can be spared in the first three quarters of this paper. In particular, the addition of signed tropical numbers will not be used before section 5.1.

2.1.2. General position. The *permanent* of the square matrix $M = (m_{ij}) \in \mathbb{T}^{n\times n}$ is given by

$$\operatorname{tper}(M) := \bigoplus_{\sigma\in\operatorname{Sym}(n)} m_{1\sigma(1)} \odot \cdots \odot m_{n\sigma(n)} = \max_{\sigma\in\operatorname{Sym}(n)} m_{1\sigma(1)} + \cdots + m_{n\sigma(n)}, \tag{2.1}$$

where $\mathrm{Sym}(n)$ is the set of all permutations of $[n]$. Computing the tropical permanent amounts to finding a permutation which attains the maximum in (2.1). Such a permutation is a solution of the assignment problem with costs (m_{ij}). It can found in time $O(n^3)$ using the Hungarian method; see [Sch03, section 17.3]. A square matrix is said to be *tropically singular* if $\mathrm{tper}(M) = \mathbb{0}$ or if the maximum is attained at least twice in (2.1).

A slightly more restrictive notion of singularity arises when signs are taken into account. A signed matrix $M \in \mathbb{T}_\pm^{n\times n}$ is *tropically sign singular* if $\mathrm{tper}(|M|) = \mathbb{0}$ or if the maximum in $\mathrm{tper}(|M|)$ is attained on two distinct permutations σ and π such that the terms $\mathrm{tsign}(\sigma) \odot m_{1\sigma(1)} \odot \cdots \odot m_{n\sigma(n)}$ and $\mathrm{tsign}(\pi) \odot m_{1\pi(1)} \odot \cdots \odot m_{n\pi(n)}$ have opposite tropical signs, where $\mathrm{tsign}(\sigma) = \mathbb{1}$ if σ is an even permutation and $\mathrm{tsign}(\sigma) = \ominus\mathbb{1}$ otherwise. The notion of tropical sign singularity of a matrix appeared in different forms in [GM84], [Plu90], and [Jos05, section 4].

We call a rectangular matrix $W \in \mathbb{T}_\pm^{m\times n}$ *tropically generic* if for every square submatrix U of W either $\mathrm{tper}(|U|) = \mathbb{0}$ or $|U|$ is not tropically singular. Similarly, the matrix W is *tropically sign generic* if $\mathrm{tper}(|U|) = \mathbb{0}$ or U is not tropically sign singular, again for all square submatrices U.

Example 2.1. Consider the following matrix with signed tropical entries:

$$W = \begin{pmatrix} -5 & -3 & \ominus 0 \\ \ominus(-7) & -5 & 0 \\ -7 & -2 & \ominus 0 \\ -2 & \ominus(-6) & \ominus 0 \end{pmatrix}.$$

The matrix W is not tropically generic. Indeed, consider its submatrix W' formed from the first two rows and the first two columns. We have $\mathrm{tper}(|W'|) = ((-5) \odot (-5)) \oplus (|\ominus(-7)| \odot (-3)) = (-10) \oplus (-10)$, thus $|W'|$ is tropically singular. However, W' is not tropically sign singular, as the terms $\mathbb{1} \odot (-5) \odot (-5) = -10$ and $\ominus\mathbb{1} \odot \ominus(-7) \odot (-3) = -10$ associated with the maximizing permutations in $\mathrm{tper}(|W'|)$ have the same tropical sign.

Now consider the submatrix $W'' = \left(\begin{smallmatrix} \ominus(-7) & 0 \\ -7 & \ominus 0 \end{smallmatrix}\right)$ formed from the second and third rows and the first and last columns of W. We have $\mathrm{tper}(|W''|) = \big(|\ominus(-7)| \odot |\ominus 0|\big) \oplus \big((-7) \odot 0\big) = (-7) \oplus (-7)$ and the two terms $\mathbb{1} \odot \ominus(-7) \odot \ominus 0 = -7$ and $\ominus\mathbb{1} \odot (-7) \odot 0 = \ominus(-7)$ have opposite tropical signs. Thus W'' is not tropically sign singular, and therefore W is not tropically sign generic.

2.2. Tropically convex sets and tropical polyhedra. A set $S \subset \mathbb{T}^n$ is said to be *tropically convex* if $\lambda \odot x \oplus \mu \odot y \in S$ for all $x, y \in S$ and $\lambda, \mu \in \mathbb{T}$ such that $\lambda \oplus \mu = \mathbb{1}$. The set S is said to be *tropical cone* when the same conclusion holds even if the requirement that $\lambda \oplus \mu = \mathbb{1}$ is omitted. A tropical cone is *polyhedral* if it is finitely generated. These notions are analogous to the classical ones, since the condition $\lambda, \mu \geqslant \mathbb{0}$ is trivially satisfied. Given $V \subset \mathbb{T}^n$, we denote by $\mathrm{tconv}(V)$ the smallest (inclusionwise) tropically convex subset of $\mathbb{T}^n$ containing V. Similarly, $\mathrm{tpos}(V)$ denote the smallest tropical cone of $\mathbb{T}^n$ containing V.

2.2.1. Tropical half-spaces and s-hyperplanes. An *(affine) tropical half-space* is a subset of $\mathbb{T}^n$ of the form

$$\max(\alpha_1 + x_1, \ldots, \alpha_n + x_n, \alpha_{n+1}) \geqslant \max(\beta_1 + x_1 \ldots, \beta_n + x_n, \beta_{n+1}), \tag{2.2}$$

where $\alpha, \beta \in \mathbb{T}^{n+1}$. When $\alpha_{n+1} = \beta_{n+1} = \mathbb{0}$, it is said to be a *linear tropical half-space*. Throughout this paper, we assume that half-spaces are defined by nontrivial inequalities.

Assumption 1. There is at least one nonnull coefficient in the inequality (2.2), i.e.,

$$\max\left(\max_{j\in[n+1]}\alpha_j,\ \max_{j\in[n+1]}\beta_j\right) > \mathbb{0}\,.$$

Without loss of generality (see [GK11, Lemma 1]), we also always assume that half-spaces are induced by an inequality satisfying the following condition.

Assumption 2. Each variable appears on at most one side of the inequality (2.2), i.e.,

$$\min(\alpha_j,\beta_j) = \mathbb{0} \text{ for all } j\in[n+1]\,.$$

Then, we can concisely describe a tropical half-space with a signed row vector $a = (a_{1j}) \in \mathbb{T}_\pm^{1\times n}$ and a signed scalar $b \in \mathbb{T}_\pm$ as

$$\begin{aligned}\mathcal{H}^{\geqslant}(a,b) :&= \bigl\{x\in\mathbb{T}^n \mid a_{11}^+\odot x_1 \oplus\cdots\oplus a_{1n}^+\odot x_n\oplus b^+ \\ &\qquad\geqslant a_{11}^-\odot x_1\oplus\cdots\oplus a_{1n}^-\odot x_n\oplus b^-\bigr\}\\ &= \{x\in\mathbb{T}^n \mid a^+\odot x\oplus b^+ \geqslant a^-\odot x\oplus b^-\}.\end{aligned}$$

A *signed tropical hyperplane*, or *s-hyperplane*, is defined as the set of the solutions $x\in\mathbb{T}^n$ of an equality of the form

$$\mathcal{H}(a,b) = \{x\in\mathbb{T}^n \mid a^+\odot x\oplus b^+ = a^-\odot x\oplus b^-\}, \tag{2.3}$$

where $a\in\mathbb{T}_\pm^{1\times n}$ and $b\in\mathbb{T}_\pm$. When $\mathcal{H}^{\geqslant}(a,b)$ is a nonempty proper subset of $\mathbb{T}^n$, its boundary is $\mathcal{H}(a,b)$.

Remark 2.2. The set $\mathcal{H}(a,b)$ is said to be *signed* because it corresponds to the tropicalization of the intersection of a usual hyperplane with the nonnegative orthant over Puiseux series; see section 2.3. A tropical (unsigned) hyperplane is defined by an unsigned row vector $a = (a_{1j}) \in \mathbb{T}^{1\times n}$ and an unsigned scalar $b\in\mathbb{T}$ as the set of all points $x\in\mathbb{T}^n$ such that the maximum is attained at least twice in $a\odot x\oplus b = \max(a_{11}+x_1,\dots,a_{1n}+x_n,b)$; see [RGST05]. This corresponds to the tropicalization of an entire ordinary hyperplane.

2.2.2. Tropical polyhedra. A *tropical polyhedron* is the intersection of finitely many tropical affine half-spaces. It will be denoted by a signed matrix $A\in\mathbb{T}_\pm^{m\times n}$ and a signed vector $b\in\mathbb{T}_\pm^m$ as

$$\mathcal{P}(A,b) := \{x\in\mathbb{T}^n \mid A^+\odot x\oplus b^+\geqslant A^-\odot x\oplus b^-\} = \bigcap_{i\in[m]}\mathcal{H}^{\geqslant}(A_i,b_i).$$

If all those tropical half-spaces are linear, i.e., if b is identically $\mathbb{0}$, that intersection is a tropical polyhedral cone.

Example 2.3. The tropical polyhedron depicted in Figure 1 is defined by the following matrix and vector:

$$A = \begin{pmatrix} -5 & -3\\ \ominus(-7) & -5\\ -7 & -2\\ -2 & \ominus(-6)\end{pmatrix} \text{ and } b = \begin{pmatrix}\ominus 0\\ 0\\ \ominus 0\\ \ominus 0\end{pmatrix}$$

The half-space depicted in orange in Figure 1 is $\mathcal{H}^{\geqslant}(A_1,b_1) = \{x\in\mathbb{T}^2 \mid \max(x_1-5, x_2-3)\geqslant 0\}$. Its boundary is the signed hyperplane $\mathcal{H}(A_1,b_1) = \{x\in\mathbb{T}^2 \mid \max(x_1-5,x_2-3) = 0\}$. The last three rows yield the inequalities

$$\max(x_2, 0) \geqslant x_1 - 7\,,$$
$$\max(x_1 - 7, x_2 - 2) \geqslant 0\,,$$
$$x_1 \geqslant \max(x_2 - 6, 0)\,,$$

which define the half-spaces respectively depicted in purple, green, and khaki in Figure 1.

A point x in a tropical polyhedron $\mathcal{P}(A, b)$ clearly satisfies the inequalities $x_j \geqslant \mathbb{0}$ for all $j \in [n]$. Although redundant, including these inequalities in the representation of a tropical polyhedron is occasionally useful.

Assumption 3. For all $j \in [n]$, all points $x \in \mathcal{P}(A, b)$ satisfy $x_j > \mathbb{0}$ or the nonnegativity constraint $x_j \geqslant \mathbb{0}$ appears in the external representation of $\mathcal{P}(A, b)$, i.e., there exists a row index $i \in [m]$ such that $(A_i\ b_i)$ is the row vector whose jth entry is $\mathbb{1}$ while all other entries are $\mathbb{0}$.

The Minkowski–Weyl theorem holds in the tropical case: a tropical polyhedron can be defined either externally (i.e., by means of half-spaces), or internally as the convex hull of finitely many points and rays.

THEOREM 2.4 (see [GK11, Theorem 2]). *A subset $\mathcal{P} \subset \mathbb{T}^n$ is a tropical polyhedron if and only if there exist two finite sets $V, R \subset \mathbb{T}^n$ such that*

$$\mathcal{P} = \{x \oplus y \mid x \in \operatorname{tconv}(V) \textit{ and } y \in \operatorname{tpos}(R)\}.$$

It will be convenient to homogenize a tropical polyhedron $\mathcal{P}(A, b)$ into the tropical polyhedral cone $\mathcal{C} := \{x \in \mathbb{T}^{n+1} \mid W^+ \odot x \geqslant W^- \odot x\}$, where $W := (A\ b)$. As a tropical cone, $\mathcal{C}$ is closed under tropical scalar multiplication. For this reason, we identify $\mathcal{C}$ with its image in the *tropical projective space*

$$\mathbb{TP}^n \;:=\; \left\{\mathbb{R} \odot x \mid x \in \mathbb{T}^{n+1} \setminus \{(\mathbb{0}, \ldots, \mathbb{0})\}\right\}\,.$$

The points of the tropical polyhedron $\mathcal{P}(A, b)$ are associated with elements of the tropical polyhedral cone $\mathcal{C}$ by the following bijection:

$$\begin{aligned} \mathcal{P}(A, b) &\longrightarrow \{y \in \mathcal{C} \mid y_{n+1} = \mathbb{1}\}, \\ x &\longmapsto (x, \mathbb{1}). \end{aligned} \tag{2.4}$$

The points of the form $(x, \mathbb{0})$ in $\mathcal{C}$ correspond to the rays in the recession cone of $\mathcal{P}(A, b)$; see [GK11].

Remark 2.5. Let $R \in \mathbb{T}^{m \times n}$ be a matrix with finite coefficients only. Then $\mathcal{P} = \operatorname{tpos}(R)$ is a tropical polyhedral cone in $\mathbb{T}^n$ such that the image of $\mathcal{P} \cap \mathbb{R}^n$ under the canonical projection from $\mathbb{R}^n$ to the *tropical torus* $\{\mathbb{R} \odot x \mid x \in \mathbb{R}^n\}$ is a "tropical polytope" in the sense of Develin and Sturmfels [DS04]. Via this identification, the tropical linear half-spaces which are nonempty proper subsets of $\mathbb{T}^n$ correspond to the "tropical half-spaces" studied in [Jos05]. The tropical projective space defined above compactifies the tropical torus (with boundary).

2.3. Puiseux series. The set $\mathbb{R}\{\{t\}\}$ of *(generalized) Puiseux series* with real coefficients is the set of formal power series

$$\boldsymbol{x} = \sum_{\alpha \in \mathbb{R}} x_\alpha t^\alpha$$

with $x_\alpha \in \mathbb{R}$ such that the support $\{x_\alpha \mid x_\alpha \neq 0\}$ is either a finite set or the set of valuations of a increasing unbounded sequence. By definition of its support, every nonnull Puiseux series $\boldsymbol{x}$ admits a smallest exponent $\alpha_{\min} \in \mathbb{R}$. The real number $-\alpha_{\min}$ is called the *valuation* of $\boldsymbol{x}$ and is denoted $\operatorname{val}(\boldsymbol{x})$. By convention, we set $\operatorname{val}(0) = -\infty$. The *leading coefficient*, denoted $\operatorname{lc}(\boldsymbol{x})$, is the coefficient $x_{\alpha_{\min}}$ of the smallest exponent $\alpha_{\min} = -\operatorname{val}(\boldsymbol{x})$ when $\boldsymbol{x} \neq 0$ and is 0 otherwise. Throughout the paper, we write $\mathbb{K}$ instead of $\mathbb{R}\{\{t\}\}$.

The set of generalized Puiseux series, equipped with the sum and product of formal power series, constitutes a field. It can be identified with a subfield of the field of Hahn series, i.e., formal power series with arbitrary well-ordered support. A variant of this field was also considered by Hardy under the name of "generalized Dirichlet series." Our approach follows Markwig [Mar10].

The n-fold Cartesian product $\mathbb{K}^n$ is a $\mathbb{K}$-vector space when equipped with the scalar multiplication $(\boldsymbol{\lambda}, \boldsymbol{x}) \mapsto \boldsymbol{\lambda}\boldsymbol{x} := (\boldsymbol{\lambda}\boldsymbol{x}_1, \dots, \boldsymbol{\lambda}\boldsymbol{x}_n)$ and the vector addition $(\boldsymbol{x}, \boldsymbol{y}) \mapsto \boldsymbol{x} + \boldsymbol{y} := (\boldsymbol{x}_1 + \boldsymbol{y}_1, \dots, \boldsymbol{x}_n + \boldsymbol{y}_n)$.

A Puiseux series $\boldsymbol{x}$ is said to be *positive* if $\operatorname{lc}(\boldsymbol{x}) > 0$, and we write $\boldsymbol{x} > 0$ in this case. Similarly, we write $\boldsymbol{x} > \boldsymbol{y}$ if $\boldsymbol{x} - \boldsymbol{y} > 0$. This definition turns $\mathbb{K}$ into an ordered field. The topology induced by this order makes $\mathbb{K}$ a topological field.

The valuation is a map from $\mathbb{K}$ to $\mathbb{T}$ which satisfies

$$\begin{aligned} \operatorname{val}(\boldsymbol{x}\boldsymbol{y}) &= \operatorname{val}(\boldsymbol{x}) \odot \operatorname{val}(\boldsymbol{y}), \\ \operatorname{val}(\boldsymbol{x} + \boldsymbol{y}) &\leqslant \operatorname{val}(\boldsymbol{x}) \oplus \operatorname{val}(\boldsymbol{y}). \end{aligned}$$

Equality occurs in the last inequality if and only if the leading terms of $\boldsymbol{x}$ and $\boldsymbol{y}$ do not cancel. In particular, cancellation never occurs whenever $\boldsymbol{x}$ and $\boldsymbol{y}$ share the same sign. This property is the main reason for using Puiseux series to study the tropical semiring. Indeed, the map val defines a homomorphism from the semiring $\mathbb{K}_+$ of nonnegative Puiseux series to the tropical semiring. This homomorphism is order preserving, that is,

$$\text{if } \boldsymbol{x} \geqslant \boldsymbol{y} \geqslant 0, \text{ then } \operatorname{val}(\boldsymbol{x}) \geqslant \operatorname{val}(\boldsymbol{y}) \,.$$

It is convenient to equip the valuation with a sign information. We define the *signed valuation* map by

$$\begin{aligned} \operatorname{sval} : \mathbb{K} &\longrightarrow \mathbb{T}_{\pm}, \\ \boldsymbol{x} &\longmapsto \begin{cases} \operatorname{val}(\boldsymbol{x}) & \text{if } \boldsymbol{x} \geqslant 0, \\ \ominus \operatorname{val}(\boldsymbol{x}) & \text{otherwise.} \end{cases} \end{aligned}$$

A *lift* of a signed tropical number $x \in \mathbb{T}_{\pm}$ is a Puiseux series $\boldsymbol{x}$ such that $\operatorname{sval}(\boldsymbol{x}) = x$. Clearly, such a lift is by no means unique. The set of all lifts will be denoted $\operatorname{sval}^{-1}(x)$. The signed valuation map is extended to vectors and matrices by componentwise application. In the following, any Puiseux series will be written in bold and its signed valuation with a standard font, e.g., $x = \operatorname{sval}(\boldsymbol{x})$.

2.4. Puiseux linear programming solves tropical linear programming. Hyperplanes, half-spaces, and convex polyhedra can be defined over an arbitrary ordered field. The most basic results used in linear programming (Farkas' lemma, Minkowski–Weyl, strong duality, etc.) are of algebraic nature. Their proofs only rely on the axioms of ordered fields and consequently are also valid in this setting; see, for instance, [Jer73, Meg87, FAA02]. Actually, in the present paper, we deal with the field $\mathbb{K}$ of Puiseux series with real coefficients, which is known to be real closed [Mar10], i.e., each nonnegative element is a square, and every polynomial with odd degree has at least one root. For such a field, stronger results follow from Tarski's principle: any first-order sentence that is valid over the reals is also valid over an arbitrary real closed field and thus valid over $\mathbb{K}$. We refer to [Tar51, Sei54] for further details; see also [BPCR06] for a recent overview. In order to have a concise name we call ordinary polyhedra defined over $\mathbb{K}$ *Puiseux polyhedra*.

In this section, we examine how tropical polyhedra are related with Puiseux polyhedra in $\mathbb{K}^n_+$ via the valuation map. In [DY07, Proposition 2.1], Develin and Yu prove that a tropical polyhedral cone $\operatorname{tpos}(R)$ can be lifted to a Puiseux polyhedral cone in $\mathbb{K}^n_+$ by lifting the set R of generators. This result can be trivially extended to arbitrary tropical polyhedra, thanks to the tropical Minkowski–Weyl theorem (Theorem 2.4), by lifting the whole internal representation. Alternatively, we shall see that a tropical polyhedron can also be lifted to a Puiseux polyhedron in $\mathbb{K}^n_+$ by lifting its external representation by half-spaces. As a consequence, an optimal solution to a tropical linear program can be found by solving a linear program over Puiseux series.

We denote by $\boldsymbol{H}(\boldsymbol{a}, \boldsymbol{b})$ the hyperplane over $\mathbb{K}^n$ defined by the equality $\boldsymbol{a}\boldsymbol{x}+\boldsymbol{b} = 0$, where $\boldsymbol{a} \in \mathbb{K}^{1\times n}$ and $\boldsymbol{b} \in \mathbb{K}$. The hyperplane $\boldsymbol{H}(\boldsymbol{a}, \boldsymbol{b})$ induces the half-space $\boldsymbol{H}^{\geqslant}(\boldsymbol{a}, \boldsymbol{b})$ by replacing the equality constraint by the inequality $\geqslant$. We will denote Puiseux polyhedra as follows:

$$\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) := \{\boldsymbol{x} \in \mathbb{K}^n \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geqslant 0\},$$

where $\boldsymbol{A} \in \mathbb{K}^{m\times n}$ and $\boldsymbol{b} \in \mathbb{K}^m$.

We now consider a tropical linear program:

$$\begin{array}{ll} \text{minimize} & c \odot x \\ \text{subject to} & x \in \mathcal{P}(A, b), \end{array} \tag{2.5}$$

where $A \in \mathbb{T}^{m\times n}_\pm$, $b \in \mathbb{T}^m_\pm$ are signed matrices and $c \in \mathbb{T}^{1\times n}$ is an unsigned row vector.

Proposition 2.6. *There is a way to associate to every tropical linear program of the form* (2.5) *satisfying Assumption* 3 *a Puiseux linear program*

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \end{array} \tag{2.6}$$

satisfying $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$, $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ *and* $\boldsymbol{c} \in \operatorname{sval}^{-1}(c)$, *so that*

(i) *the image by the valuation of the feasible set of the linear program* (2.6) *is precisely the feasible set of the tropical linear program* (2.5)*; in particular,* (2.6) *is feasible if and only if* (2.5) *is feasible;*

(ii) *the valuation of any optimal solution of* (2.6) *(if any) is an optimal solution of* (2.5)*.*

Notice that the converse of (ii) does not necessarily hold. That is, there are tropical linear programs with optimal solutions which do not arise as projections from any lift; see Example 2.8 below.

Proof. To begin with, we will exhibit lifts $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ of the external representation such that $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})) = \mathcal{P}(A, b)$. The inclusion $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})) \subset \mathcal{P}(A, b)$ is satisfied for any lifts $\boldsymbol{A}$, $\boldsymbol{b}$. Indeed, consider a point $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Then, by Assumption 3, the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is included in the nonnegative orthant $\mathbb{K}^n_+$. Let $(\boldsymbol{A}\ \boldsymbol{b}) = (\boldsymbol{A}^+\ \boldsymbol{b}^+) - (\boldsymbol{A}^-\ \boldsymbol{b}^-)$, where the entries of $(\boldsymbol{A}^+\ \boldsymbol{b}^+)$ and $(\boldsymbol{A}^-\ \boldsymbol{b}^-)$ are nonnegative. Every point $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ satisfies $\boldsymbol{A}^+\boldsymbol{x}+\boldsymbol{b}^+ \geqslant \boldsymbol{A}^-\boldsymbol{x}+\boldsymbol{b}^-$. Since the Puiseux series on each side of these inequalities are nonnegative, the valuation preserves their ordering and $A^+ \odot x \oplus b^+ \geqslant A^- \odot x \oplus b^-$.

We claim that the reverse inclusion holds for any lift of the form $(\boldsymbol{A}\ \boldsymbol{b}) = (\boldsymbol{A}^+\ \boldsymbol{b}^+) - (\boldsymbol{A}^-\ \boldsymbol{b}^-)$ defined, for $i \in [m]$ and $j \in [n]$, by

$$\begin{aligned} \boldsymbol{A}^+ &= (\alpha t^{-a_{ij}^+}) \text{ and } \boldsymbol{A}^- = (t^{-a_{ij}^-}), \\ \boldsymbol{b}^+ &= (\alpha t^{-b_i^+}) \text{ and } \boldsymbol{b}^- = (t^{-b_i^-}), \end{aligned} \tag{2.7}$$

where α is a real number strictly greater than $n+1$, and $A = (a_{ij})$. To see this, observe that for any $x \in \mathcal{P}(A, b)$ the lift $\boldsymbol{x} = (t^{-x_1}, \dots, t^{-x_n})$ belongs to the Puiseux polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Indeed, for any $i \in [m]$ we have

$$\boldsymbol{A}_i^- \boldsymbol{x} + \boldsymbol{b}_i^- = \sum_{j=1}^{n} t^{-a_{ij}^- - x_j} + t^{-b_i^-} \leqslant (n+1) t^{-(a_i^- \odot x \oplus b_i^-)} < \alpha t^{-(a_i^- \odot x \oplus b_i^-)}$$

and

$$\boldsymbol{A}_i^+ \boldsymbol{x} + \boldsymbol{b}_i^+ = \alpha \left(\sum_{j=1}^{n} t^{-a_{ij}^+ - x_j} + t^{-b_i^+} \right) \geqslant \alpha t^{-(a_i^+ \odot x \oplus b_i^+)} \geqslant \alpha t^{-(a_i^- \odot x \oplus b_i^-)},$$

thus $\boldsymbol{A}_i \boldsymbol{x} + \boldsymbol{b}_i > 0$. This shows that a lift to real Puiseux series does exist.

We need to prove the claimed properties of such a lift. Let $\boldsymbol{A}$ and $\boldsymbol{b}$ as above. Since $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})) = \mathcal{P}(A, b)$, the Puiseux linear program (2.6) is feasible if and only if the tropical one (2.5) is feasible. Now take any $\boldsymbol{c} \in \operatorname{sval}^{-1}(c)$, e.g., $\boldsymbol{c}_j = t^{-c_j}$ for $j \in [n]$. If (2.6) admits an optimal solution $\boldsymbol{x}^*$, then $\boldsymbol{cx} \geqslant \boldsymbol{cx}^* \geqslant 0$ for all $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Since c is nonnegative, $c \odot x \geqslant c \odot \operatorname{val}(\boldsymbol{x}^*)$ for all $x \in \operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})) = \mathcal{P}(A, b)$. This concludes the proof. $\square$

Remark 2.7. Observe that in Proposition 2.6, the Puiseux linear program (2.6) cannot be unbounded. Indeed, for all lifts $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$, we have $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \subset \mathbb{K}_+^n$ thanks to Assumption 3. Since c has tropically positive entries, its lift $\boldsymbol{c}$ also has positive entries. Then the inequality $\boldsymbol{cx} \geqslant 0$ holds for all $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ and thus provides a lower bound for the minimization problem (2.6).

We offer a visualization for Proposition 2.6 in Figure 2. As model theory explains, polyhedra over real Puiseux series, for most purposes, pretty much are the same as classical polyhedra over the reals. Therefore, we can also visualize the lift in Proposition 2.6 as a classical polyhedron. More precisely, the constraints of the lift are provided by (2.6), and we can choose the lifted objective vector to be $\boldsymbol{c}_j = t^{-c_j}$ for $j \in [n]$. Now, replacing the parameter t by a real number provides a real linear program. If that real number is sufficiently small, the classical linear program over the reals is combinatorially equivalent to the Puiseux one, that is, they share the same vertex-facet incidences and hence they share the same vertex-edge graph; the optimal vertices are in bijection.

Example 2.8. Throughout the rest of this paper, we will illustrate our results on the following tropical linear program:

$$\begin{array}{lll} & \text{minimize} & \max(x_1 - 2, x_2, x_3 - 1) \\ \mathcal{H}_1 & \text{subject to} & \max(0, x_2 - 1) \geqslant \max(x_1 - 1, x_3 - 1) \\ \mathcal{H}_2 & & x_3 \geqslant \max(0, x_2 - 2) \\ \mathcal{H}_3 & & x_2 \geqslant 0 \\ \mathcal{H}_4 & & x_1 \geqslant \max(0, x_2 - 3) \\ \mathcal{H}_5 & & 0 \geqslant x_2 - 4\ . \end{array}$$

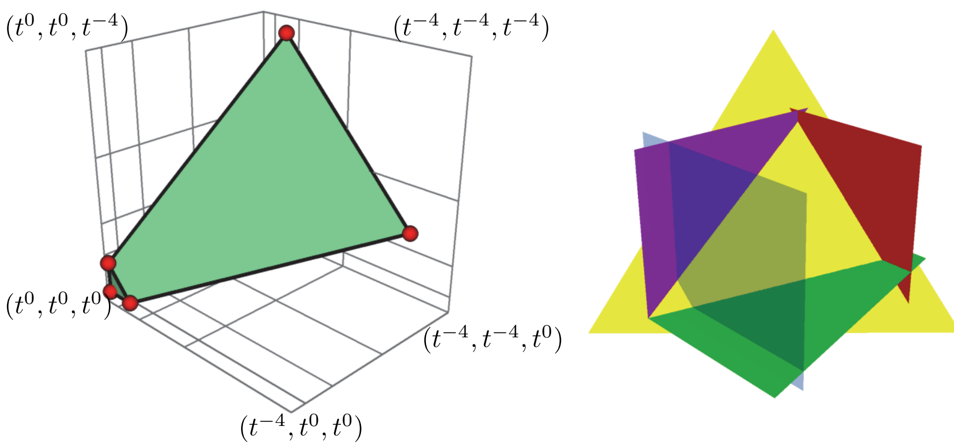

Fig. 2. *A lift of the tropical polyhedron defined by the inequalities from Example* 2.8 *and its external representation.*

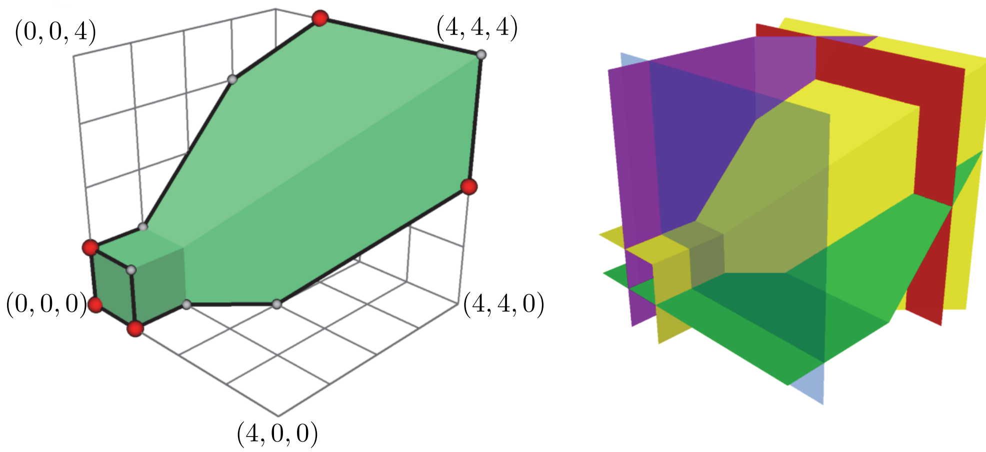

Fig. 3. *The tropical polyhedron defined by the inequalities from Example* 2.8 *and its external representation.*

These constraints define the tropical polyhedron represented in Figure 3. A lift of this tropical polyhedron is depicted in Figure 2. The optimal valuation of this tropical linear program is 0 and the set of optimal solutions is the ordinary square:

$$\{(x_1, x_2, x_3) \in \mathbb{T}^3 \mid 0 \leqslant x_1 \leqslant 1 \text{ and } x_2 = 0 \text{ and } 0 \leqslant x_3 \leqslant 1\}.$$

However, over Puiseux series, there is a unique optimum. It is the point located in the intersection of three hyperplanes obtained by lifting the inequalities $(\mathcal{H}_2)$, $(\mathcal{H}_3)$, and $(\mathcal{H}_4)$. This point has valuation $(0, 0, 0)$, which is an optimum for the tropical linear program. Corollary 3.6 and Proposition 5.5 below assert this does not depend on the choice of the lift. The reason is that our example satisfies the standard conditions mentioned in Theorem 1.1.

We will also present several results in the homogeneous setting. They will be illustrated on the tropical cone defined by the following inequalities:

$$
\begin{aligned}
\max(x_4, x_2 - 1) &\geqslant \max(x_1 - 1, x_3 - 1), \\
x_3 &\geqslant \max(x_4, x_2 - 2), \\
x_2 &\geqslant x_4, \\
x_1 &\geqslant \max(x_4, x_2 - 3), \\
x_4 &\geqslant x_2 - 4.
\end{aligned}
\tag{2.9}
$$

This cone corresponds to the previous polyhedron by the correspondence given in (2.4), i.e., the coordinate x_4 plays the role of the affine component. For the sake of simplicity, the linear half-spaces in (2.9) are still referred to as $(\mathcal{H}_1)$–$(\mathcal{H}_5)$.

2.5. The simplex method. A Puiseux linear program can be solved using the classical simplex method. We briefly recall the basic facts. Let $I \subset [m]$ be a subset of cardinality n such that the submatrix $\boldsymbol{A}_I$, formed from the rows with indices in I, is nonsingular. The intersection $\bigcap_{i\in I} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$ contains a unique point, which we denote as $\boldsymbol{x}^I$. When $\boldsymbol{x}^I$ belongs to the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, it is called a *(feasible) basic point*.

Remark 2.9. A basis is usually defined by a partition of the (explicitly bounded) variables $(\boldsymbol{s}_1, \dots, \boldsymbol{s}_m)$ in "basic" and "nonbasic" variables, where $\boldsymbol{s} = \boldsymbol{A}\boldsymbol{x}+\boldsymbol{b}$. Observe that I corresponds to the "nonbasic" variables as it indexes the zero coordinates of $\boldsymbol{s}$. The set I can also be interpreted as the set of "basic" variables in the dual program. For any set $I \subset [m]$ we let

$$
\boldsymbol{\mathcal{P}}_I(\boldsymbol{A}, \boldsymbol{b}) := \bigcap_{i\in I} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i) \cap \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})\,.
$$

A subset $K \subset [m]$ of cardinality $n-1$ defines the *(feasible) edge* $\boldsymbol{\mathcal{E}}_K := \boldsymbol{\mathcal{P}}_K(\boldsymbol{A}, \boldsymbol{b})$ when $\bigcap_{i\in K} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$ is an affine line that intersects $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Notice that an edge defined in this way may have "length zero," i.e., as a set it consists of only a single point. A basic point $\boldsymbol{x}^I$ is contained in the n edges defined by the sets $I \setminus \{k\}$ for $k \in I$. The edge $I \setminus \{k\}$ belongs to the line directed by the vector $\boldsymbol{d}^k$ with coordinates

$$
\boldsymbol{d}^k_j = (-1)^{k+j} \frac{\det \boldsymbol{M}_{kj}}{\det \boldsymbol{A}_I}\ , \tag{2.10}
$$

where $\boldsymbol{M}_{kj}$ is the matrix obtained from $\boldsymbol{A}_I$ by deleting its kth row and jth column.

As we are minimizing, moving along the edge $I \setminus \{k\}$ from the basic point $\boldsymbol{x}^I$ improves the objective function if the *reduced cost* $\boldsymbol{y}_k = \boldsymbol{c}\boldsymbol{d}^k$ is negative. The vector of reduced costs $\boldsymbol{y} = (\boldsymbol{y}_k)_{k\in I}$ forms a solution of the following linear system of equations:

$$
-\boldsymbol{A}_I^\top \boldsymbol{y} + \boldsymbol{c} = 0\ . \tag{2.11}
$$

Each iteration of the simplex method starts on some basic point $\boldsymbol{x}^I$. An edge $I \setminus \{k\}$ with a negative reduced cost is selected. If no such edge exists, then the basic point is optimal by the strong duality theorem of linear programming [Sch03, section 5.5] (which holds in any ordered field such as real Puiseux series). Otherwise, the algorithm *pivots*, i.e., moves to the other end of the selected edge. Pivoting amounts to finding the *length* $\boldsymbol{\mu} \in \mathbb{K}$ of the edge, which is given by

$$\boldsymbol{\mu} = \inf \left\{ \frac{\boldsymbol{A}_i \boldsymbol{x}^I + \boldsymbol{b}_i}{-\boldsymbol{A}_i \boldsymbol{d}^k} \mid i \in [m] \setminus I \text{ and } \boldsymbol{A}_i \boldsymbol{d}^k < 0 \right\} . \tag{2.12}$$

If the edge is bounded, i.e., if there exists an $i \in [m] \setminus I$ such that $\boldsymbol{A}_i \boldsymbol{d}^k < 0$, then the algorithm reaches a new basic point. Otherwise, the linear program is unbounded, and the valuation $\boldsymbol{\mu}$ is ∞.

Remark 2.10. Basic points and directions are provided by determinants. If $(A\ b) = \operatorname{val}(\boldsymbol{A}\ \boldsymbol{b})$ is tropically generic, this amounts to computing tropical permanents. However, the length $\boldsymbol{\mu}$ of the edge cannot be computed only with valuations. This difficulty can be observed already in dimension one. Consider the Puiseux polyhedron defined by the inequalities

$$\boldsymbol{x} \leqslant 1 \text{ and } \boldsymbol{x} \geqslant t^2 \text{ and } \boldsymbol{x} \geqslant t^3 \ .$$

Minimizing $\boldsymbol{c} = 1$ and starting from the basic point $\boldsymbol{x} = 1$, the direction of the single pivot is $\boldsymbol{d} = -1$. The pivoting step must decide where the edge ends, and in this case the edge length is given by $\boldsymbol{\mu} = \min(1 - t^2, 1 - t^3) = 1 - t^2$. Yet, the valuation of $1 - t^2$ and $1 - t^3$ yields zero in both cases. This shows that in order to find the correct minimum t^2, it does not suffice to look at the valuations of the optimal solutions of the Puiseux lift.

3. Tropical basic points and tropical edges. Geometrically speaking, the classical simplex method traces the vertex-edge graph of an ordinary polyhedron from one basic point along a directed path to an optimal solution, which again is basic. Basic points and edges over Puiseux series are cells of the arrangement of hyperplanes $\{\boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)\}_{i \in [m]}$. It turns out that, under some genericity assumptions, the valuation of these cells can be described by intersecting tropical half-spaces in $\{\mathcal{H}^{\geqslant}(A_i, b_i)\}_{i \in [m]}$ and s-hyperplanes in $\{\mathcal{H}(A_i, b_i)\}_{i \in [m]}$. This result will be proved in Corollary 3.6 below.

3.1. The tangent digraph. Consider a matrix $W = (w_{ij}) \in \mathbb{T}_{\pm}^{m \times (n+1)}$. For every point $x \in \mathbb{T}^{n+1}$ with no $\mathbb{0}$ entries, we define the *tangent graph* $\mathcal{G}_x(W)$ at the point x with respect to W as a bipartite graph over the following two disjoint sets of nodes: the "coordinate nodes" $[n+1]$ and the "hyperplane nodes" $\{i \in [m] \mid W_i^+ \odot x = W_i^- \odot x > \mathbb{0}\}$. There is an edge between the hyperplane node i and the coordinate node j when $j \in \arg\max(|W_i| \odot x)$.

The *tangent digraph* $\vec{\mathcal{G}}_x(W)$ is an oriented version of $\mathcal{G}_x(W)$, where the edge between the hyperplane node i and the coordinate node j is oriented from j to i when w_{ij} is tropically positive, and from i to j when w_{ij} is tropically negative (if a tangent digraph contains an edge between i and j then $w_{ij} \neq \mathbb{0}$).

Examples of tangent digraphs are given in Figure 4 (where hyperplane nodes are denoted $\mathcal{H}_i$). The term "tangent" comes from the fact that $\vec{\mathcal{G}}_x(W)$ is a combinatorial encoding of the tangent cone at x in the tropical cone $\mathcal{C} = \mathcal{P}(W, \mathbb{0})$; see [AGG13]. The tangent digraph is the same for any two points in the same cell of the arrangement of tropical hyperplanes given by the inequalities. The tangent graph $\mathcal{G}_x(W)$ corresponds to the "types" introduced in [DS04] but relative only to the hyperplanes given by the tight inequalities at x.

When there is no risk of confusion, we will denote by $\mathcal{G}_x$ and $\vec{\mathcal{G}}_x$ the tangent graph and digraph, respectively.

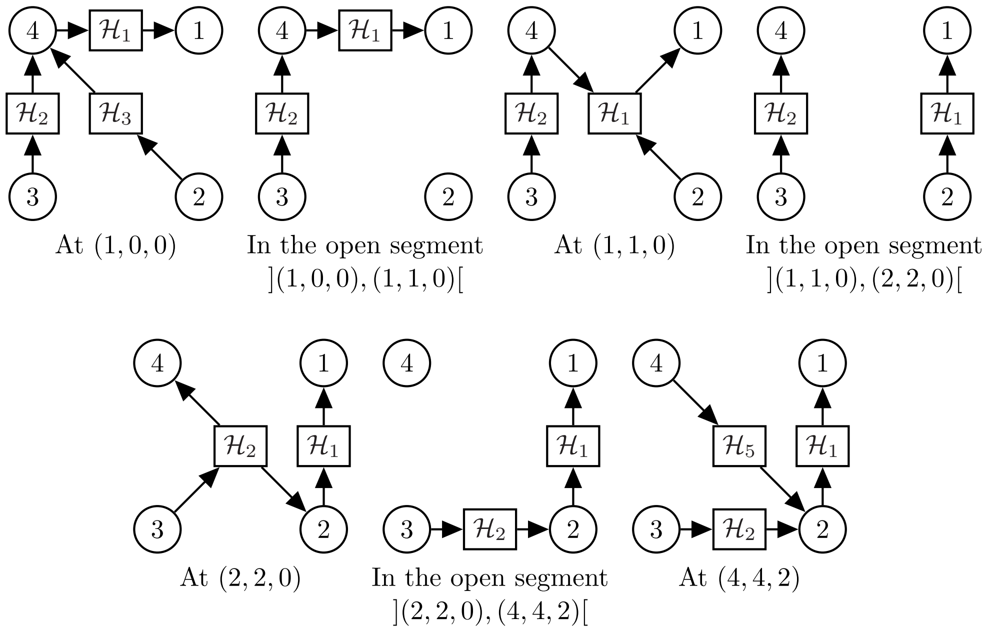

FIG. 4. *Tangent digraphs at various points of the tropical cone obtained by homogenization of the tropical polyhedron defined by the inequalities from Example* 2.8. *Hyperplane nodes are rectangles and coordinate nodes are circles.*

Example 3.1. Let W be the matrix formed by the coefficients of the system (2.9), and consider the point $x = (1, 0, 0, 0)$ (corresponding to $(1, 0, 0)$ via the bijection (2.4)). The inequalities $(\mathcal{H}_1)$, $(\mathcal{H}_2)$, and $(\mathcal{H}_3)$ are tight at x. They read

$$\begin{aligned} \max(\underline{x_4}, x_2 - 1) &\geqslant \max(\underline{x_1 - 1}, x_3 - 1), \\ \underline{x_3} &\geqslant \max(\underline{x_4}, x_2 - 2), \\ \underline{x_2} &\geqslant \underline{x_4}, \end{aligned}$$

where we marked the positions where the maxima are attained. The tangent digraph $\vec{\mathcal{G}}_x(W)$ is depicted in the top left of Figure 4. For instance, the first inequality provides the arcs from coordinate node 4 to hyperplane node $\mathcal{H}_1$ and from $\mathcal{H}_1$ to coordinate node 1.

If I and J are respectively subsets of the hyperplane and coordinate nodes of $\mathcal{G}_x$, a *matching* between I and J is a subgraph of $\mathcal{G}_x$ with node set $I \cup J$ in which every node is incident to exactly one edge.

LEMMA 3.2. *Let $W \in \mathbb{T}_{\pm}^{m\times(n+1)}$ and $x \in \mathbb{T}^{n+1}$ be a point with no $\mathbb{0}$ entries. Suppose the tangent graph $\mathcal{G}_x$ contains a matching between the hyperplane nodes I and the coordinate nodes J. Then the submatrix W' of W formed from rows I and columns J is such that $|W'|$ has a finite tropical permanent. Moreover, the matching yields a maximizing permutation in the latter tropical permanent.*

Proof. Let $\{(i_1, j_1), \ldots, (i_q, j_q)\}$ be a matching between the hyperplane nodes $I = \{i_1, \ldots, i_q\}$ and the coordinate nodes $J = \{j_1, \ldots, j_q\}$. By definition of the tangent graph, for all $p \in [q]$, we have

$$|w_{i_p j_p}| + x_{j_p} \geqslant |w_{i_p l}| + x_l \quad \text{for all } l \in [n+1] \ .$$

Since x has no $\mathbb{0}$ entries, this implies $\sum_{p=1}^{q} |w_{i_p j_p}| \geqslant \sum_{p=1}^{q} |w_{i_p \sigma(i_p)}|$ for any bijection $\sigma : I \to J$. Thus the tropical permanent of $|W'|$ is $\sum_{p=1}^{q} |w_{i_p j_p}|$, which is obtained with the bijection $i_p \mapsto j_p$.

Consider a $p \in [q]$. By definition of the tangent graph $|W_{i_p}| \odot x > \mathbb{0}$. Moreover, we suppose that x has finite entries. Thus $|w_{i_p j_p}| > \mathbb{0}$. As a consequence, $\sum_{p=1}^{q} |w_{i_p j_p}|$ is finite. ∎

LEMMA 3.3. *Let $W \in \mathbb{T}_{\pm}^{m\times(n+1)}$ and $x \in \mathbb{T}^{n+1}$ be a point with no $\mathbb{0}$ entries. If the tangent graph $\mathcal{G}_x$ contains an undirected cycle, then the matrix W contains a square submatrix W' such that $|W'|$ is tropically singular and* $\operatorname{tper}(|W'|) > \mathbb{0}$. *Moreover, if the cycle is directed in the tangent digraph $\vec{\mathcal{G}}_x$, then W' is tropically sign singular.*

Proof. To prove the first statement, let $j_1, i_1, j_2, \dots, i_q, j_{q+1} = j_1$ be an undirected cycle in $\mathcal{G}_x$. By Assumption 2 we have $q \geqslant 2$. Up to restricting to a subcycle, we may assume that the cycle is simple, i.e., the indices $i_1, \dots, i_q$ and $j_1, \dots, j_q$ are pairwise distinct. As a consequence, the maps $\sigma : i_p \mapsto j_p$ and $\tau : i_p \mapsto j_{p+1}$ for $p \in [q]$ are bijections. The sets of edges $\{(i_p, j_p) \mid p \in [q]\}$ and $\{(i_p, j_{p+1}) \mid p \in [q]\}$ are two distinct matchings between the hyperplane nodes $i_1, \dots, i_p$ and the coordinate nodes $j_1, \dots, j_p$. Let W' be the submatrix of W formed from rows $i_1, \dots, i_q$ and columns $j_1, \dots, j_q$. By Lemma 3.2, the bijections σ and τ are both maximizing in $\operatorname{tper}(|W'|)$. Lemma 3.2 also shows that $\operatorname{tper}(|W'|) > \mathbb{0}$.

Now suppose that the cycle is directed. Then, $w_{i_p j_p}$ is tropically positive and $w_{i_p j_{p+1}}$ is tropically negative for all $p \in [q]$. Consequently, the tropical signs of $w_{i_1 j_1} \odot \cdots \odot w_{i_q j_q}$ and $w_{i_1 j_2} \odot \cdots \odot w_{i_q j_{q+1}}$ differ by $(-1)^q$. Moreover, τ is obtained from σ by a cyclic permutation of order q, so their signs differ by $(-1)^{q+1}$. As a result, the terms $\operatorname{tsign}(\sigma) \odot w_{i_1 j_1} \odot \cdots \odot w_{i_q j_q}$ and $\operatorname{tsign}(\tau) \odot w_{i_1 j_2} \odot \cdots \odot w_{i_q j_{q+1}}$ have opposite tropical signs. This completes the proof. ∎

3.2. Cells of an arrangement of signed tropical hyperplanes. Our next result shows how the tangent digraph can be used to get sufficient control on the lift of the points in a tropical polyhedron to points in some Puiseux polyhedron, while dealing mostly with inequality descriptions. Throughout this section we assume that the extended matrix $(A\ b)$ is tropically sign generic.

A tropical polyhedron $\mathcal{P}(A, b)$ is always the image under the valuation map of a polyhedron over Puiseux series. Indeed, consider an internal representation $\operatorname{tconv}(V) \oplus \operatorname{tpos}(R)$ of $\mathcal{P}(A, b)$ (which exists by Theorem 2.4). The result of Develin and Yu, [DY07, Proposition 2.1] implies that any lift $\boldsymbol{V}, \boldsymbol{R}$ of the sets V, R provides a Puiseux polyhedron $\boldsymbol{\mathcal{P}} = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R})$ such that $\operatorname{val}(\boldsymbol{\mathcal{P}}) = \mathcal{P}(A, b)$.

One can also lift a tropical polyhedron through its inequality representation. For example, for any $A \in \mathbb{T}_{\pm}^{m\times n}$ and $b \in \mathbb{T}_{\pm}^{m}$ satisfying Assumption 3, Proposition 2.6 provides Puiseux matrices $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ such that $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})) = \mathcal{P}(A, b)$. However, the latter equality may fail for arbitrary lifts $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$.

Example 3.4. Consider the tropical polyhedron

$$\mathcal{P} = \{x \in \mathbb{T}^2 \mid \max(0, x_2) \geqslant x_1,\ \max(0, x_1) \geqslant x_2,\ x_1 \geqslant \mathbb{0},\ x_2 \geqslant \mathbb{0}\}\ . \tag{3.1}$$

One possible lift of this representation in terms of inequalities is the Puiseux polyhedron

$$\boldsymbol{\mathcal{P}} = \{\boldsymbol{x} \in \mathbb{K}^2 \mid 2 + \boldsymbol{x}_2 \geqslant 2\boldsymbol{x}_1,\ 2 + \boldsymbol{x}_1 \geqslant 2\boldsymbol{x}_2,\ \boldsymbol{x}_1 \geqslant 0,\ \boldsymbol{x}_2 \geqslant 0\}\ . \tag{3.2}$$

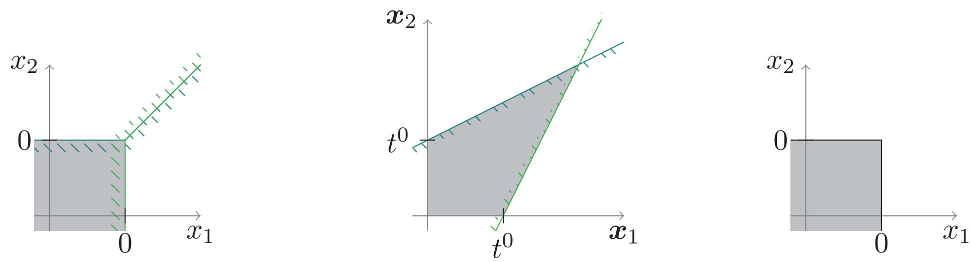

FIG. 5. *Left: the tropical polyhedron $\mathcal{P}$ described in* (3.1)*; middle: the Puiseux polyhedron $\boldsymbol{\mathcal{P}}$ obtained by lifting the inequality representation of $\mathcal{P}$ as in* (3.2)*; right: the set* $\operatorname{val}(\boldsymbol{\mathcal{P}})$*, which is stricly contained in $\mathcal{P}$.*

Since $\boldsymbol{\mathcal{P}}$ is contained in the nonnegative orthant, we have $\operatorname{val}(\boldsymbol{\mathcal{P}}) \subseteq \mathcal{P}$. Here this inclusion is strict. The set $\operatorname{val}(\boldsymbol{\mathcal{P}})$ consists of all the nonpositive points in $x \in \mathbb{T}^2$ with $x_1 \leqslant 0$ and $x_2 \leqslant 0$, while $\mathcal{P}$ additionally contains the half-line $\{(\lambda, \lambda) \mid \lambda > 0\}$. To show this, suppose that there exists $(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \boldsymbol{\mathcal{P}}$ such that $\operatorname{val}(\boldsymbol{x}_1) = \operatorname{val}(\boldsymbol{x}_2) = \lambda > 0$. Let $u_1 t^\lambda$ and $u_2 t^\lambda$ be the leading terms of $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, respectively. Then the inequality $2 + \boldsymbol{x}_1 \geqslant 2\boldsymbol{x}_2$ implies that $u_1 \geqslant 2u_2$, while $2 + \boldsymbol{x}_2 \geqslant 2\boldsymbol{x}_1$ imposes that $u_2 \geqslant 2u_1$, and we obtain a contradiction as $u_1, u_2 > 0$. See Figure 5.

THEOREM 3.5. *Suppose that $(A\ b)$ is tropically sign generic. Then the identity*

$$\operatorname{val}\Big(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n\Big) \; = \; \mathcal{P}(A, b)$$

holds for any $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$.

Proof. Let $W = (A\ b)$. For any $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$, let $\boldsymbol{W} = (\boldsymbol{A}\ \boldsymbol{b})$. We first prove the result for the cones $\mathcal{C} = \mathcal{P}(W, \mathbb{0})$ and $\boldsymbol{\mathcal{C}} = \boldsymbol{\mathcal{P}}(\boldsymbol{W}, 0)$. The inclusion $\operatorname{val}(\boldsymbol{\mathcal{C}} \cap \mathbb{K}_+^{n+1}) \subset \mathcal{C}$ is trivial. Conversely, let $x \in \mathcal{C}$. Up to removing the columns j of W with $x_j = \mathbb{0}$, we can assume that x has no $\mathbb{0}$ entries. We construct a lift $\boldsymbol{x}$ of x in the Puiseux cone $\boldsymbol{\mathcal{C}} \cap \mathbb{K}_+^{n+1}$ using the tangent digraph $\vec{\mathcal{G}}_x$ with hyperplane node set I. We claim that it is sufficient to find a vector $v \in \mathbb{R}^{n+1}$ satisfying the following conditions:

$$\sum_{j \in \arg\max(|W_i| \odot x)} \operatorname{lc}(\boldsymbol{w}_{ij}) v_j > 0 \quad \text{for all } i \in I\ , \tag{3.3}$$

$$v_j > 0 \quad \text{for all } j \in [n+1]\ , \tag{3.4}$$

where $\boldsymbol{W} = (\boldsymbol{w}_{ij})$.

Indeed, given such a vector v, consider the lift $\boldsymbol{x} = (v_j t^{-x_j})_j$ of x. Clearly $\boldsymbol{x} \in \mathbb{K}_+^{n+1}$. If $i \in I$, then (3.3) ensures that the leading coefficient of $\boldsymbol{W}_i \boldsymbol{x}$ is positive. If $i \notin I$, two cases can occur. Either $W_i^+ \odot x = W_i^- \odot x = \mathbb{0}$ and thus $\boldsymbol{W}_i \boldsymbol{x} = 0$. Otherwise, $W_i^+ \odot x > W_i^- \odot x$, so the leading term of $\boldsymbol{W}_i \boldsymbol{x}$ is positive. We conclude that $\boldsymbol{W}_i \boldsymbol{x} \geqslant 0$ for all $i \in [m]$. This proves the claim.

Let $F = (f_{ij}) \in \mathbb{R}^{I \times (n+1)}$ be the real matrix defined by $f_{ij} = \operatorname{lc}(\boldsymbol{w}_{ij})$ when $j \in \arg\max(|W_i| \odot x)$ and $f_{ij} = 0$ otherwise. We claim that there exists $v \in \mathbb{R}^{n+1}$ such that $Fv > 0$ and $v > 0$ or, equivalently, that the following polyhedron is not empty:

$$\{v \in \mathbb{R}^{n+1} \mid Fv \geqslant 1,\ v \geqslant 1\}.$$

By contradiction, suppose that the latter polyhedron is empty. Then, by Farkas' lemma [Sch03, section 5.4], there exists $\alpha \in \mathbb{R}_+^I$ and $\lambda \in \mathbb{R}_+^{n+1}$ such that

$$F^\top \alpha + \lambda \leqslant 0, \tag{3.5}$$

$$\sum_{i \in I} \alpha_i + \sum_{j \in [n+1]} \lambda_j > 0. \tag{3.6}$$

Note that if α is the 0 vector, then by (3.6) there exists a $\lambda_j > 0$ for some $j \in [n+1]$, which contradicts (3.5). Thus, the set $K = \{i \in I \mid \alpha_i > 0\}$ is not empty. Let $J \subset [n+1]$ be defined by

$$J := \bigcup_{i \in K} \arg\max(W_i^+ \odot x) = \bigcup_{i \in K} \{j \mid f_{ij} > 0\}.$$

By definition of the tangent digraph, every hyperplane node in K has an incoming arc from a coordinate node in J. Moreover, for every $j \in J$, the inequality (3.5) yields

$$\sum_{i \in I} f_{ij} \alpha_i \leqslant 0 \ .$$

This sum contains a positive term $f_{ij}\alpha_i$ (by definition of J). Consequently, it must also contain a negative term $f_{kj}\alpha_k$. Equivalently, $k \in K$ and $f_{kj} < 0$, which means that the coordinate node j has an incoming arc from the hyperplane node k. It follows that the tangent digraph $\vec{\mathcal{G}}_x$ contains a directed cycle (through nodes $K \cup J$). Then, by Lemma 3.3, the matrix W contains a tropically sign singular submatrix with $\operatorname{tper}(|W|) > \mathbb{0}$. This proves the claim.

Now we consider the polyhedron $\mathcal{P}(A, b)$. The inclusion $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n) \subset \mathcal{P}(A, b)$ is still valid. Conversely, given $x \in \mathcal{P}(A, b)$, the point $x' = (x, \mathbb{1}) \in \mathbb{T}^{n+1}$ belongs to the cone $\mathcal{C}$. By the previous proof, there exists a lift $\boldsymbol{x}'$ of x' in $\boldsymbol{\mathcal{C}} \cap \mathbb{K}_+^{n+1}$. Since $\operatorname{val}(\boldsymbol{x}'_{n+1}) = \mathbb{1}$, the point $\boldsymbol{x} = (\boldsymbol{x}'_1/\boldsymbol{x}'_{n+1}, \dots, \boldsymbol{x}'_n/\boldsymbol{x}'_{n+1})$ is well-defined. Furthermore, $\boldsymbol{x}$ clearly satisfies $\operatorname{val}(\boldsymbol{x}) = x$ and it belongs to $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n$. ∎

Theorem 3.5 shows that valuation commutes with intersection for half-spaces in general position. This is still true for mixed intersections of half-spaces and (s-)hyperplanes. Similar to our notation for Puiseux polyhedra, we let

$$\mathcal{P}_I(A, b) := \bigcap_{i \in I} \mathcal{H}(A_i, b_i) \cap \mathcal{P}(A, b).$$

Corollary 3.6. *Suppose that $(A\ b)$ is tropically sign generic. Then, for all $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$, $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$, and $I \subset [m]$,*

$$\operatorname{val}\Big(\boldsymbol{\mathcal{P}}_I(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n\Big) \ = \ \mathcal{P}_I(A, b)\,. \tag{3.7}$$

Proof. We first prove the result when $I = [m]$. In this case, the claim is about the intersection of all (Puiseux or signed tropical) hyperplanes in the arrangement. The first inclusion $\operatorname{val}\big(\bigcap_{i=1}^m \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i) \cap \mathbb{K}_+^n\big) \subset \bigcap_{i=1}^m \mathcal{H}(A_i, b_i)$ is trivial. Conversely, let $x \in \bigcap_{i=1}^m \mathcal{H}(A_i, b_i)$. Note that there is nothing to prove if that intersection is empty. The point x belongs to the tropical polyhedron $\mathcal{P}(A, b)$. By Theorem 3.5, x admits a lift in the Puiseux polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n$. But observe that the choice of tropical signs for the rows of $(A\ b)$ is arbitrary. Indeed, if $(A'\ b')$ is obtained

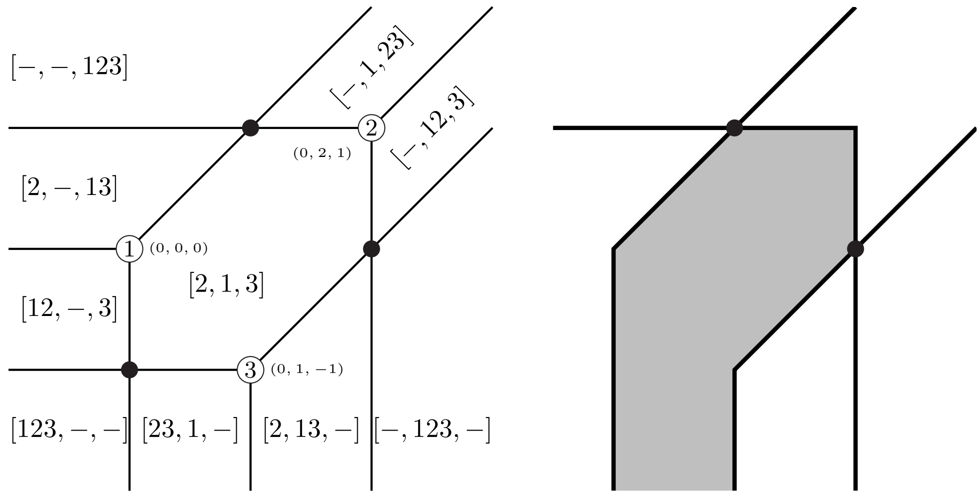

FIG. 6. *Unsigned (left) and signed (right) cell decompositions induced by the three tropical s-hyperplanes in Example* 3.7.

by multiplying some rows of $(A\ b)$ by $\ominus\mathbb{1}$, then $(A'\ b')$ satisfies the conditions of Theorem 3.5 and x belongs to $\mathcal{P}(A', b')$. Thus, for any sign pattern $s \in \{-1, +1\}^m$, there exists a lift $\boldsymbol{x}^s$ of x which belongs to the Puiseux polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}^s, \boldsymbol{b}^s) \cap \mathbb{K}_+^n$, where $(\boldsymbol{A}^s\ \boldsymbol{b}^s) = \left(\begin{smallmatrix} s_1 & & \\ & \ddots & \\ & & s_m \end{smallmatrix}\right)(\boldsymbol{A}\ \boldsymbol{b})$.

Since the Puiseux points $\boldsymbol{x}^s$ are nonnegative with valuation x, any point in their convex hull is also nonnegative with valuation x. We claim that the convex hull $\operatorname{conv}\{\boldsymbol{x}^s \mid s \in \{-1, +1\}^m\}$ contains a point in the intersection $\bigcap_{i=1}^m \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$. We prove the claim by induction on the number m of hyperplanes.

If $m = 1$, we obtain two points $\boldsymbol{x}^+$ and $\boldsymbol{x}^-$ on each side of the hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_1, \boldsymbol{b}_1)$, and it is easy to see that their convex hull intersects the hyperplane. Now, suppose we have $m \geqslant 2$ hyperplanes. Let S^+ (resp., S^-) be the set of all sign patterns $s \in \{-1, +1\}^m$ with $s_m = +1$ (resp., $s_m = -1$). By induction, the convex hull $\operatorname{conv}\{\boldsymbol{x}^s \mid s \in S^+\}$ contains a point $\boldsymbol{x}^+$ in the intersection of the first $m-1$ hyperplanes $\bigcap_{i=1}^{m-1} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$. Similarly, $\operatorname{conv}\{\boldsymbol{x}^s \mid s \in S^-\}$ contains a point $\boldsymbol{x}^-$ in $\bigcap_{i=1}^{m-1} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$. The points $\boldsymbol{x}^+$ and $\boldsymbol{x}^-$ are on opposite sides of the last hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_m, \boldsymbol{b}_m)$, and thus their convex hull intersects $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_m, \boldsymbol{b}_m)$.

When $I \subsetneq [m]$, the previous proof can be generalized by considering only the sign patterns $s \in \{-1, +1\}^m$ such that $s_i = +1$ for all $i \notin I$. $\square$

By Corollary 3.6, the intersection of the nonnegative orthant $\mathbb{K}_+^n$ with the cells of the arrangement of Puiseux hyperplanes $\{\boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)\}_{i \in [m]}$ induces a cellular decomposition of $\mathbb{T}^n$ into tropical polyhedra. We call this collection of tropical polyhedra the *signed cells* of the arrangement of tropical s-hyperplanes $\{\mathcal{H}(A_i, b_i)\}_{i \in [m]}$. Notice that the signed cells form an intersection poset thanks to Corollary 3.6.

The signed cell decomposition coarsens the cell decomposition introduced in [DS04], which partitions $\mathbb{T}^n$ into ordinary polyhedra. Here we call the latter cells *unsigned*. In particular, the one-dimensional signed cells are unions of (closed) one-dimensional unsigned cells. However, some one-dimensional unsigned cells may not belong to any one-dimensional signed cell. In the example depicted in Figure 3, this is the case for the ordinary line segment $[(1, 0, 1), (1, 1, 1)]$.

Example 3.7. Consider the tropical polyhedral cone $\mathcal{C}$ in $\mathbb{T}^3$ given by the three homogenous constraints

$$x_2 \geqslant \max(x_1, x_3), \tag{3.8}$$

$$x_1 \geqslant \max(x_2 - 2, x_3 - 1), \tag{3.9}$$

$$\max(x_1, x_3 + 1) \geqslant x_2 - 1 \ . \tag{3.10}$$

This gives rise to an arrangement of three tropical s-hyperplanes in which $\mathcal{C}$ forms one signed cell; see Figure 6 (right) for a visualization in the $x_1 = 0$ plane. Each tropical s-hyperplane yields a unique unsigned tropical hyperplane. An *open sector* is one connected component of the complement of an unsigned tropical hyperplane. The ordinary polyhedral complex arising from intersecting the open sectors of an arrangement of unsigned tropical hyperplanes is the *type decomposition* of Develin and Sturmfels [DS04]. In our example the type decomposition has 10 unsigned maximal cells; in Figure 6 (left), we marked them with labels as in [DS04].

The *apices* of the unsigned tropical hyperplanes arising from the three constraints above are $p_1 = (0, 0, 0)$, $p_2 = (0, 2, 1)$, and $p_3 = (0, 1, -1)$. The tropical convex hull of p_1, p_2, and p_3, with respect to min as the tropical addition, is the topological closure of the unsigend bounded cell $[2, 1, 3]$.

The signed cell C is precisely the union of the two maximal unsigned cells $[2, 1, 3]$ and $[23, 1, -]$ together with the (relatively open) bounded edge of type $[23, 1, 3]$ sitting in between. The other signed cells come about by replacing "$\geqslant$" by "$\leqslant$" in some subset of the constraints above. For instance, exchanging "$\geqslant$" by "$\leqslant$" in (3.8) and keeping the other two yields the signed cell which is the union of the three unsigned cells $[2, -, 13]$, $[12, -, 3]$, $[123, -, -]$ and two (relatively open) edges in between. Altogether there are six maximal signed cells in this case.

The proper notion of a "face" of a tropical polyhedron is a subject of active research; see [Jos05] and [DY07]. Notice that, like the tangent digraphs, the signed and unsigned cells depend on the arrangement of s-hyperplanes, while several different arrangements may describe the same tropical polyhedron. For example,

$$\{x \in \mathbb{T}^2 \mid x_1 \oplus x_2 \leqslant 1\} = \{x \in \mathbb{T}^2 \mid x_1 \leqslant 1 \text{ and } x_2 \leqslant 1\} \, . \tag{3.11}$$

Even if a canonical external representation exists (see [AK13]), it may not satisfy the genericity conditions of Corollary 3.6. Thus this approach does not easily lead to a meaningful notion of faces for tropical polyhedra.

Lacking a good notion of a "face," the following two results introduce suitable concepts which are good enough for our algorithms.

PROPOSITION-DEFINITION 3.8 (tropical basic points). *Suppose that* $(A\ b)$ *is tropically sign generic and that Assumption* 3 *holds. Let* I *be a subset of* $[m]$ *of cardinality* n *such that* $\operatorname{tper}(|A_I|) > \mathbb{0}$. *If the set* $\mathcal{P}_I(A, b)$ *is not empty, it contains a unique point, called a* (feasible) tropical basic point *of* $\mathcal{P}(A, b)$.

The tropical basic points *of* $\mathcal{P}(A, b)$ *are exactly the valuations of the basic points of* $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ *for any lift* $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$.

Proof. This is a straightforward consequence of Corollary 3.6 and the definition of basic points. Assumption 3 ensures that the sets of the form $\boldsymbol{\mathcal{P}}_I(\boldsymbol{A}, \boldsymbol{b})$, for any $I \subset [m]$, are contained in $\mathbb{K}_+^n$. ∎

Remark 3.9. Alternatively, the fact that $\mathcal{P}_I(A, b)$, if it is not empty, contains a unique element, follows from the Cramer theorem in the symmetrized tropical semiring [Plu90], see Theorem 5.1 below. This also implies that the technical condition that $x_j > \mathbb{0}$ in Assumption 3 is not needed to derive the uniqueness of this element.

PROPOSITION-DEFINITION 3.10 (tropical edges). *Suppose that* $(A\ b)$ *is tropically sign generic and that Assumption* 3 *holds. Let* K *be a subset of* $[m]$ *of cardinality* $n-1$ *such that* A_K *has a tropically sign nonsingular maximal minor. If the set* $\mathcal{P}_K(A,b)$ *is not empty, then it is called a* (feasible) tropical edge.

The tropical edges *of* $\mathcal{P}(A,b)$ *are exactly the valuation of the edges of* $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ *for any lift* $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$.

Proof. The arguments are the same as in the proof of Proposition-Definition 3.8. ◻

The correspondence between basic points and edges with their tropical counterparts is illustrated in Figures 2 and 3, where basic points are depicted by red dots and edges by black lines. These definitions are meaningful only if $(A\ b)$ is tropically sign generic. Otherwise, $\mathcal{P}(A,b)$ may have no basic points in the sense of Proposition-Definition 3.8. For instance, the set $\{x \in \mathbb{T}^2 \mid x_1 \geqslant x_2 \oplus 0 \text{ and } x_2 \geqslant x_1 \oplus 0\}$ does not contain such a point. Notice that our genericity assumptions imply that the tropical edges arise as complete intersections of tropical half-spaces. In this sense Corollary 3.6 is a signed version of [SS04, Proposition 6.3].

A point x in a tropically convex set S is called an *extreme point* of S if, for any $y, z \in S$, $x \in \operatorname{tconv}(\{y,z\})$ implies $x = y$ or $x = z$.

PROPOSITION 3.11. *Suppose that* $(A\ b)$ *is tropically sign generic and that Assumption* 3 *holds. Then the extreme points of* $\mathcal{P}(A,b)$ *are tropical basic points.*

Proof. Consider any lift $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$. Then $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}) \subset \mathbb{K}_+^n$ by Assumption 3, and $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})) = \mathcal{P}(A,b)$ by Corollary 3.6. The basic points of $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ are precisely its extreme points. As a result, we have $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}) = \operatorname{conv}(\boldsymbol{P}) + \operatorname{pos}(\boldsymbol{R})$, where $\boldsymbol{P}$ is the set of basic points, $\operatorname{conv}(\boldsymbol{P})$ its convex hull, and $\operatorname{pos}(\boldsymbol{R})$ is a pointed polyhedral cone generated by a finite set $\boldsymbol{R} \subset \mathbb{K}^n$. Note that $\boldsymbol{R} \subset \mathbb{K}_+^n$ as $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}) \subset \mathbb{K}_+^n$. Thus, by [DY07, Proposition 2.1], we know that $\mathcal{P}(A,b) = \operatorname{tconv}(\operatorname{val}(\boldsymbol{P})) \oplus \operatorname{tpos}(\operatorname{val}(\boldsymbol{R}))$. By Proposition-Definition 3.8, $\operatorname{val}(\boldsymbol{P})$ is precisely the set of tropical basic points of $\mathcal{P}(A,b)$. The tropical analogue of Milman's converse of the Krein–Milman theorem, which is proved, for instance, in Theorem 2 of [GK11] in the case of polyhedra, implies that the set of extreme points of $\mathcal{P}(A,b)$ is included in $\operatorname{val}(\boldsymbol{P})$. It follows that every extreme point is basic. ◻

Remark 3.12. While extreme points are basic points, the converse does not hold. For example, $(1,1)$ is a basic point of the tropical polyhedron $\{x \in \mathbb{T}^2 \mid x_1 \leqslant 1 \text{ and } x_2 \leqslant 1\}$, but it is not extreme.

The polars of sign cyclic polyhedral cones studied in [AGK11] are also examples in which not every basic point is extreme. Actually, Theorems 2 and 6 in that reference provide combinatorial characterizations of extreme and basic points in terms of lattice paths. A comparison of both characterizations shows that the lattice paths are more constrained in the case of extreme points.

4. Pivoting between two tropical basic points. In this section, we show how to pivot from a tropical basic point to another, i.e., to move along a tropical edge between the two points, in a tropical polyhedron $\mathcal{P}(A,b)$, where $A \in \mathbb{T}_\pm^{m\times n}$ and $b \in \mathbb{T}_\pm^m$. Under some genericity conditions, by Corollary 3.6 this is equivalent to the same pivoting operation in an arbitrary lift $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ over Puiseux series of the considered tropical polyhedron, where $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$. However, our method relies only on the tropical matrix A and the tropical vector b. The complexity of this tropical pivot operation will be shown to be $O(n(m+n))$, which is analogous to the classical pivot operation.

Pivoting is more easily described in homogeneous terms. For $W = (A\ b)$ we consider the tropical cone $\mathcal{C} = \mathcal{P}(W, \mathbb{0})$, seen as a subset of the tropical projective space $\mathbb{TP}^n$. This cone is defined as the intersection of the half-spaces $\mathcal{H}_i^{\geqslant} := \{x \in \mathbb{TP}^n \mid W_i^+ \odot x \geqslant W_i^- \odot x\}$ for $i \in [m]$. Similarly, we denote by $\mathcal{H}_i$ the s-hyperplane $\{x \in \mathbb{TP}^n \mid W_i^+ \odot x = W_i^- \odot x\}$. We also let $\mathcal{C}_I := \mathcal{P}_I(W, \mathbb{0})$ for any subset $I \subset [m]$.

Throughout this section, we make the following assumptions.

Assumption 4. The matrix W is tropically generic.

Assumption 5. Every point in $\mathcal{C} \setminus \{(\mathbb{0}, \dots, \mathbb{0})\}$ has finite coordinates.

Assumption 4 is a tropical version of primal nondegeneracy. It is strictly stronger than the condition that $W = (A\ b)$ is tropically sign generic used in the previous section, and hence, in particular, we can make use of Corollary 3.6. Under Assumption 5, which is strictly stronger than Assumption 3, the tropical polyhedron $\mathcal{P}(A, b)$ is a bounded subset of $\mathbb{R}^n$. To see this, consider $\mathcal{C}$ as a subset of $\mathbb{T}^{n+1}$. As $\mathcal{C}$ is a closed set, Assumption 5 implies that there exists a vector $\ell \in \mathbb{R}^{n+1}$ such that $x \geqslant \ell$ for all $x \in \mathcal{C}$. Let $\operatorname{tconv}(P) \oplus \operatorname{tpos}(R)$ be the internal description of $\mathcal{P}(A, b)$ provided by Theorem 2.4. If R contains a point r, then it is easy to verify that $(r, \mathbb{0})$ lies in $\mathcal{C}$, which contradicts Assumption 5. Since every $p \in P$ belongs to $\mathcal{P}(A, b)$, the point $(p, \mathbb{1})$ belongs to $\mathcal{C}$, and thus $p_j \geqslant l_j$ for all $j \in [n]$. It follows that $\mathcal{P}(A, b) = \operatorname{tconv}(P)$ is a bounded subset of $\mathbb{R}^n$.

Assumptions 4 and 5 are two of the three conditions required for a tropical linear program to be standard in the sense of Theorem 1.1.

As a consequence, the Puiseux polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is also bounded and contained in the interior of $\mathbb{K}_+^n$.

Through the bijection given in (2.4), the tropical basic point associated with a suitable subset $I \subset [m]$ is identified with the unique projective point $x^I \in \mathbb{TP}^n$ in the intersection $\mathcal{C}_I$. In addition, when pivoting from the basic point x^I, we move along a tropical edge $\mathcal{E}_K := \mathcal{C}_K$ defined by a set $K = I \setminus \{i_{\mathsf{out}}\}$ for some $i_{\mathsf{out}} \in I$.

By Proposition-Definitions 3.8 and 3.10, a tropical edge $\mathcal{E}_K$ is a tropical line segment $\operatorname{tconv}(x^I, x^{I'})$. The other endpoint $x^{I'} \in \mathbb{TP}^n$ is a basic point for $I' = K \cup \{i_{\mathsf{ent}}\}$, where $i_{\mathsf{ent}} \in [m] \setminus I$. So, the notation i_{out} and i_{ent} refers to the indices leaving and entering the set of active constraints I which is maintained by the algorithm. Notice that the latter set corresponds to the *nonbasic* indices in the classical primal simplex method, so that the indices entering/leaving I correspond to the indices leaving/entering the usual basis, respectively.

As a tropical line segment, $\mathcal{E}_K$ is known to be the concatenation of at most n ordinary line segments.

PROPOSITION 4.1 (see [DS04, Proposition 3]). *Let $\mathcal{E}_K = \operatorname{tconv}(x^I, x^{I'})$ be a tropical edge. Then there exists an integer $q \in [n]$ and $q+1$ points $\xi^1, \dots, \xi^{q+1} \in \mathcal{E}_K$ such that*

$$\mathcal{E}_K = [\xi^1, \xi^2] \cup \dots \cup [\xi^q, \xi^{q+1}], \qquad \textit{where } \xi^1 = x^I \textit{ and } \xi^{q+1} = x^{I'}.$$

Every ordinary segment is of the form

$$[\xi^p, \xi^{p+1}] = \{x^p + \lambda e^{J_p} \mid 0 \leqslant \lambda \leqslant \mu_p\}, \tag{4.1}$$

where the length of the segment μ_p is a positive real number, $J_p \subset [n+1]$, and the jth coordinate of the vector e^{J_p} is equal to 1 *if $j \in J_p$ and to* 0 *otherwise. Moreover, the sequence of subsets $J_1, \dots, J_q$ satisfies*

$$\emptyset \subsetneq J_1 \subsetneq \dots \subsetneq J_q \subsetneq [n+1].$$

The vector e^{J_p} is called the *direction* of the segment $[\xi^p, \xi^{p+1}]$. The intermediate points $\xi^2, \ldots, \xi^q$ are called *breakpoints*. In the tropical polyhedron depicted in Figure 3, breakpoints are represented by white dots.

Note that in the tropical projective space $\mathbb{TP}^n$, the directions e^J and $-e^{[n+1]\setminus J}$ coincide. Both correspond to the direction of $\mathbb{T}^n$ obtained by removing the $(n+1)$th coordinate of either $-e^{[n+1]\setminus J}$ if $(n+1) \in J$ or e^J otherwise.

4.1. Overview of the pivoting algorithm. We now provide a sketch of the pivoting operation along a tropical edge $\mathcal{E}_K$. Geometrically, the idea is to traverse the ordinary segments $[\xi^1, \xi^2], \ldots, [\xi^q, \xi^{q+1}]$ of $\mathcal{E}_K$. At each point ξ^p, for $p \in [q]$, we first determine the direction vector e^{J_q}, then move along this direction until the point ξ^{p+1} is reached. As the tangent digraph at a point $x \in \mathcal{C}$ encodes the local geometry of the tropical cone $\mathcal{C}$ around x, the direction vectors can be read from the tangent digraphs. Moreover, the tangent digraphs are acyclic under Assumption 4. This imposes strong combinatorial conditions on the tangent digraph, which, in turn, allows us to easily determine the feasible directions.

For the sake of simplicity, let us suppose that the tropical edge consists of two consecutive segments $[\xi, \xi']$ and $[\xi', \xi'']$, with direction vectors e^J and $e^{J'}$, respectively.

Starting at the basic point $\xi = x^{K \cup \{i_{\mathsf{out}}\}}$, we shall prove below that, at every basic point, the tangent digraph is a spanning tree where every hyperplane node has exactly one incoming arc and one outgoing arc. In other words, for every $i \in K \cup \{i_{\mathsf{out}}\}$, the sets $\arg\max(W_i^+ \odot \xi)$ and $\arg\max(W_i^- \odot \xi)$ are both reduced to a singleton, say, $\{j_i^+\}$ and $\{j_i^-\}$. We want to "get away" from the s-hyperplane $\mathcal{H}_{i_{\mathsf{out}}}$. Since the direction vector e^J is a 0/1 vector, the only way to do so is to increase the variable indexed by $j^+_{i_{\mathsf{out}}}$ while not increasing the component indexed by $j^-_{i_{\mathsf{out}}}$. Hence, we must have $j^+_{i_{\mathsf{out}}} \in J$ and $j^-_{i_{\mathsf{out}}} \notin J$. While moving along e^J, we also want to stay inside the s-hyperplane $\mathcal{H}_i$ for $i \in K$. Hence, if $j_i^+ \in J$ for some $i \in K$, we must also have $j_i^- \in J$. Similarly, if $j_i^+ \notin J$, then we must also have $j_i^- \notin J$. Removing the hyperplane node i_{out} from the tangent digraph $\vec{\mathcal{G}}_\xi$ provides two connected components; the first one, C_+, contains $j^+_{i_{\mathsf{out}}}$, and the second one, C_-, contains $j^-_{i_{\mathsf{out}}}$. From the discussion above, it follows that the set J consists of the coordinate nodes in C^+.

When moving along e^J from ξ, we leave the s-hyperplane $\mathcal{H}_{i_{\mathsf{out}}}$. Consequently, the hyperplane node i_{out} "disappears" from the tangent digraph. It turns out that this is the only modification that happens to the tangent digraph. More precisely, at every point in the open segment $]\xi, \xi'[$, the tangent digraph is the graph obtained from $\vec{\mathcal{G}}_\xi$ by removing the hyperplane node i_{out} and its two incident arcs. We shall denote this digraph by $\vec{\mathcal{G}}_{]\xi,\xi'[}$. By construction, $\vec{\mathcal{G}}_{]\xi,\xi'[}$ is acyclic and consists of two connected components, and every hyperplane node has one incoming and one outgoing arc.

We shall move from ξ along e^J until "something" happens to the tangent digraph. In fact only two things can happen, depending on whether ξ' is a breakpoint or a basic point. As we supposed ξ' to be a breakpoint, a new arc a_{new} will "appear" in the tangent digraph, i.e., $\vec{\mathcal{G}}_{\xi'} = \vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{a_{\mathsf{new}}\}$. Let us sketch how the arc a_{new} can be found. We denote $a_{\mathsf{new}} = (j_{\mathsf{new}}, k)$, where j_{new} is a coordinate node and $k \in K$ is a hyperplane node. We shall see that j_{new} must belongs to J (i.e., the component C_+), while k must belongs to the component C_-.

Hence, the arc a_{new} "reconnects" the two components C_+ and C_- (see Figure 8). Since k had one incoming and one outgoing arc in $\vec{\mathcal{G}}_{]\xi,\xi'[}$, it has exactly three incident arcs in $\vec{\mathcal{G}}_{\xi'}$. One of them is $a_{\mathsf{new}} = (j_{\mathsf{new}}, k)$; the second one, $a_{\mathsf{old}} = (j_{\mathsf{old}}, k)$, has the same orientation as a_{new}; and the third one, $a' = (k, l)$, has an orientation opposite to a_{new} and a_{old}.

Let us now find the direction vector $e^{J'}$ of the second segment $[\xi', \xi'']$. Consider the hyperplane node k with the three incident arcs $a_{\mathsf{new}}, a_{\mathsf{old}}$, and a'. By Proposition 4.1, we know that $J \subset J'$; hence we must increase the variable j_{new}. Since we want to stay inside the hyperplane $\mathcal{H}_k$, it follows that we must also increase the variable indexed by ℓ. On the other hand, we do not increase the variable j_{old}. As before, all hyperplane nodes $i \in K \setminus \{k\}$ have exactly one incoming and one outgoing arc. Removing the arc a_{old} from the graph provides two connected components; the first one, C'_+, contains the coordinate nodes j_{new} and ℓ as well as the hyperplane node k, while the second one, C'_-, contains j_{old}. The new direction set J' is given by the coordinate nodes in C'_+.

The tangent digraph in the open segment $]\xi', \xi''[$ is again constant and is defined by $\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$. Hence, $\vec{\mathcal{G}}_{]\xi',\xi''[}$ is an acyclic graph, with two connected components C'_+ and C'_-, where every hyperplane node has one incoming and one outgoing arc.

The basic point ξ'' is reached when a new s-hyperplane $i_{\mathsf{ent}} \notin K$ is hit. This happens when the hyperplane node i_{ent} "appears" in the tangent digraph, along with one incoming (j^+, i_{ent}) and one outgoing arc (i_{ent}, j^-). Observe that we must have $j^- \in J$ and $j^+ \notin J$. It follows that the two components C'_+ and C'_- are reconnected by adding i_{ent} and its two incident arcs.

In section 4.2, we prove that the tangent digraphs satisfy the above-mentioned characterization and that they provide the feasible directions. In section 4.3, we characterize the lengths of the ordinary segments, that is, we deduce when a arc or a hyperplane node "appears" in the tangent digraphs. In section 4.4, we prove that the tangent digraphs evolve as described above. It allows us to incrementally update the information needed to find the directions and lengths of the segments. This will finally provide an efficient implementation of the pivoting operation.

4.2. Directions of ordinary segments. Given a point x in a tropical cone $\mathcal{D}$, we say that the direction e^J, with $\emptyset \subsetneq J \subsetneq [n+1]$, is *feasible* from x in $\mathcal{D}$ if there exists $\mu > 0$ such that the ordinary segment $\{x + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ is included in $\mathcal{D}$. The following lemma will be helpful to prove the feasibility of a direction.

LEMMA 4.2. *Let $x \in \mathbb{R}^{n+1}$. Then, the following properties hold:*
- (i) *If x belongs to $\mathcal{H}_i^{\geqslant} \setminus \mathcal{H}_i$, every direction is feasible from x in $\mathcal{H}_i^{\geqslant}$.*
- (ii) *If x belongs to $\mathcal{H}_i$, the direction e^J is feasible from x in the half-space $\mathcal{H}_i^{\geqslant}$ if and only if $\arg\max(W_i^+ \odot x) \cap J \neq \emptyset$ or $\arg\max(W_i^- \odot x) \cap J = \emptyset$.*
- (iii) *If x belongs to $\mathcal{H}_i$, the direction e^J is feasible from x in the s-hyperplane $\mathcal{H}_i$ if and only if the sets $\arg\max(W_i^+ \odot x) \cap J$ and $\arg\max(W_i^- \odot x) \cap J$ are both empty or both nonempty.*

Proof. The first point is immediate. To prove the last two points, observe that if $x \in \mathcal{H}_i$, then $W_i^+ \odot x = W_i^- \odot x > \mathbb{0}$, thanks to $x \in \mathbb{R}^{n+1}$ and Assumption 1. Then, for $\lambda > 0$ sufficiently small, we have

$$W_i^+ \odot (x + \lambda e^J) = \begin{cases} (W_i^+ \odot x) + \lambda & \text{if } \arg\max(W_i^+ \odot x) \cap J \neq \emptyset \,, \\ W_i^+ \odot x & \text{otherwise} \,, \end{cases}$$

and the same property holds for $W_i^- \odot x$. ∎

We propose to determine feasible directions with tangent graphs. It turns out that tangent graphs in a tropical edge have a very special structure. Indeed, under Assumption 4, these graphs do not contain any cycle by Lemma 3.3. In other words, they are forests: each connected component is a tree. Hence we have the equality

$$\text{number of nodes} = \text{number of edges} + \text{number of connected components}\,. \tag{4.2}$$

Note that [DS04, Proposition 17] determines that number of connected components.

We introduce some additional basic notions and notation on directed graphs. Two nodes of a digraph are said to be *weakly connected* if they are connected in the underlying undirected graph. Given a directed graph $\vec{\mathcal{G}}$ and a set $\mathcal{A}$ of arcs between some nodes of $\vec{\mathcal{G}}$, we denote by $\vec{\mathcal{G}} \cup \mathcal{A}$ the digraph obtained by adding the arcs of $\mathcal{A}$. Similarly, if $\mathcal{A}$ is a subset of arcs of $\vec{\mathcal{G}}$, we denote by $\vec{\mathcal{G}} \setminus \mathcal{A}$ the digraph where the arcs of $\mathcal{A}$ have been removed. By extension, if $\mathcal{N}$ is a subset of nodes of $\vec{\mathcal{G}}$, then $\vec{\mathcal{G}} \setminus \mathcal{N}$ is defined as the digraph obtained by removing the nodes in $\mathcal{N}$ and their incident arcs. The *degree* of a node of $\vec{\mathcal{G}}$ is defined as the pair (p_1, p_2), where p_1 and p_2 are the numbers of incoming and outgoing arcs incident to the node.

PROPOSITION 4.3. *Let x be a point in a tropical edge $\mathcal{E}_K$. Then, exactly one of the following cases arises:*

(C1) *x is a basic point for the basis $K \cup \{i_{\mathsf{out}}\}$, where $i_{\mathsf{out}} \in [m] \setminus K$. The tangent graph $\mathcal{G}_x$ at x is a spanning tree, and the set of hyperplane nodes is $K \cup \{i_{\mathsf{out}}\}$. In the tangent digraph $\vec{\mathcal{G}}_x$, every hyperplane node has degree $(1,1)$. Let J be the set of coordinate nodes weakly connected to the unique node in $\arg\max(W^+_{i_{\mathsf{out}}} \odot x)$ in the digraph $\vec{\mathcal{G}}_x \setminus \{i_{\mathsf{out}}\}$. The only feasible direction from x in $\mathcal{E}_K$ is e^J.*

(C2) *x is in the relative interior of an ordinary segment. The tangent graph $\mathcal{G}_x$ is a forest with two connected components, and the set of hyperplane nodes is K. In the tangent digraph $\vec{\mathcal{G}}_x$, every hyperplane node has degree $(1,1)$. Let J be the set of coordinate nodes in one of the components. The two feasible directions from x in $\mathcal{E}_K$ are e^J and $-e^J = e^{[n+1]\setminus J}$.*

(C3) *x is a breakpoint. The tangent graph $\mathcal{G}_x$ is a spanning tree, and the set of hyperplane nodes is K. In the tangent digraph $\vec{\mathcal{G}}_x$, there is exactly one hyperplane node k with degree $(2,1)$ or $(1,2)$, while all other hyperplane nodes have degree $(1,1)$. Let a and a' be the two arcs incident to k with same orientation. Let J and J' be the set of coordinate nodes weakly connected to k in $\vec{\mathcal{G}}_x \setminus \{a\}$ and $\vec{\mathcal{G}}_x \setminus \{a'\}$, respectively. The two feasible directions from x in $\mathcal{E}_K$ are e^J and $e^{J'}$.*

Proof. Since x has finite entries, the graph $\mathcal{G}_x$ contains exactly $n+1$ coordinate nodes. Let n' be the number of hyperplane nodes in $\mathcal{G}_x$. Consider any $i \in K$. Since x is contained in the s-hyperplane $\mathcal{H}_i$ and $x \in \mathbb{R}^{n+1}$, we have $W^+_i \odot x = W^-_i \odot x > \mathbb{0}$. Thus K is contained in the set of hyperplane nodes. Therefore $n' \geqslant n-1$. As there is at least one connected component, there is at most $n+n'$ edges by (4.2). In addition, each hyperplane node is incident to at least two edges, so that there is at least $2n'$ edges in $\mathcal{G}_x$. We deduce that $n' \leqslant n$. As a result, by using (4.2), we can distinguish three cases:

- (i) $n' = n$, in which case there is only one connected component in $\mathcal{G}_x$ and exactly $2n$ edges. In addition, all the hyperplane nodes have degree $(1,1)$ in $\vec{\mathcal{G}}_x$.
- (ii) $n' = n-1$, the graph $\mathcal{G}_x$ contains precisely two connected components and $2n'-2$ edges. As in the previous case, every hyperplane node has degree $(1,1)$ in $\vec{\mathcal{G}}_x$.
- (iii) $n' = n-1$ and $\mathcal{G}_x$ has one connected component. In this case, there are $2n'-1$ edges. In $\vec{\mathcal{G}}_x$, there is exactly one hyperplane node with degree $(2,1)$ or $(1,2)$, and all the other hyperplane nodes have degree $(1,1)$.

We next show that these cases correspond to the ones in our claim.

Case (i). Since $n' = n$, the set of hyperplane nodes is of the form $K \cup \{i_{\mathsf{out}}\}$ for some $i_{\mathsf{out}} \notin K$. Moreover, $\mathcal{G}_x$ is a spanning tree. As a consequence, it contains a matching between the coordinate nodes $[n]$ and the hyperplane nodes $K \cup \{i_{\mathsf{out}}\}$. Such a matching can be constructed as follows. Let $\vec{\mathcal{G}}'$ be the digraph obtained by directing the edges of $\mathcal{G}_x$ toward the coordinate node $n+1$. In this digraph, every coordinate node $j \in [n]$ has exactly one outgoing arc to a hyperplane node $\sigma(j)$, as there is exactly one path from j to $n+1$ in the spanning tree $\mathcal{G}_x$. Moreover, every hyperplane node i has exactly one incoming arc and one outgoing arc in $\vec{\mathcal{G}}'$. Indeed, i is incident to two arcs in $\vec{\mathcal{G}}'$, and exactly one of them leads to the path to coordinate node $n+1$. We conclude that $\sigma(j) \neq \sigma(j')$ when $j \neq j'$. Thus the set of edges $\{(j, \sigma(j)) \mid j \in [n]\}$ forms the desired matching. Then, by Lemma 3.2, the submatrix W' of W formed from the columns in $[n]$ and the rows in $K \cup \{i_{\mathsf{out}}\}$ satisfy $\operatorname{tper}(|W'|) > \mathbb{0}$. Furthermore, $W' = A_{K \cup \{i_{\mathsf{out}}\}}$. As a consequence, x is a basic point for the set $K \cup \{i_{\mathsf{out}}\}$.

Since the graph $\mathcal{G}_x$ is a spanning tree where the hyperplane node i_{out} is not a leaf, removing i_{out} from $\mathcal{G}_x$ provides two connected components C^+ and C^-, containing the coordinate nodes in $\arg\max(W^+_{i_{\mathsf{out}}} \odot x)$ and in $\arg\max(W^-_{i_{\mathsf{out}}} \odot x)$, respectively. Let J be the set of the coordinate nodes in C^+.

We claim that the direction e^J is feasible from x in $\mathcal{E}_K$. Indeed, if the hyperplane node $i \in K$ belongs to C^+, then $\arg\max(W^+_i \odot x) \subset J$ and $\arg\max(W^-_i \odot x) \subset J$. In contrast, if the node $i \in K$ belongs to C^-, we have $\arg\max(W^+_i \odot x) \cap J = \arg\max(W^-_i \odot x) \cap J = \emptyset$. By Lemma 4.2, this shows that the direction e^J is feasible in all s-hyperplanes $\mathcal{H}_i$ with $i \in K$. It is also feasible in the half-space $\mathcal{H}^{\geqslant}_{i_{\mathsf{out}}}$, since $x \in \mathcal{H}_{i_{\mathsf{out}}}$ and $\arg\max(W^+_{i_{\mathsf{out}}} \odot x) \subset J$. Finally, for all $i \notin K \cup \{i_{\mathsf{out}}\}$, the point x belongs to $\mathcal{H}^{\geqslant}_i \setminus \mathcal{H}_i$. Indeed, if $x \in \mathcal{H}_i$, then i would be a hyperplane node. Thus, by Lemma 4.2, the direction e^J is feasible in $\mathcal{H}^{\geqslant}_i$. As $\mathcal{E}_K = (\cap_{i \in K} \mathcal{H}_i) \cap (\cap_{i \notin K} \mathcal{H}^{\geqslant}_i)$, this proves the claim.

Since x is a basic point it admits exactly one feasible direction in $\mathcal{E}_K$. Thus e^J is the only feasible direction from x in $\mathcal{E}_K$.

Case (ii). In this case, $\mathcal{G}_x$ is a forest with two components C_1 and C_2, and K is precisely the set of hyperplane nodes. Let J be the set of coordinate nodes in C_1. Then Lemma 4.2 shows that the direction e^J is feasible from x in $\mathcal{E}_K$. Indeed, the point x belongs to $\mathcal{H}^{\geqslant}_i \setminus \mathcal{H}_i$ for $i \notin K$. In addition, for all $i \in K$, the sets $\arg\max(W^+_i \odot x) \cap J$ and $\arg\max(W^-_i \odot x) \cap J$ are both nonempty if i belongs to C_1, and both are empty otherwise.

Symmetrically, the direction $e^{[n+1] \setminus J} = -e^J$ is also feasible in $\mathcal{E}_K$, as $[n+1] \setminus J$ is the set of coordinate nodes in the component C_2. It follows that x is in the relative interior of an ordinary segment.

Case (iii). The graph $\mathcal{G}_x$ is a spanning tree. Let k be the unique half-space node of degree $(2,1)$ or $(1,2)$ in $\vec{\mathcal{G}}_x$ and a, a' the two arcs incident to k with the same orientation.

Then $\vec{\mathcal{G}}_x \setminus \{a\}$ consists of two weakly connected components C_1 and C_2. Without loss of generality, we assume that k belongs to C_1. Let J be the set of coordinate nodes in C_1. We now prove that e^J is feasible from x in $\mathcal{E}_K$, thanks to Lemma 4.2. Indeed, $x \in \mathcal{H}^{\geqslant}_i \setminus \mathcal{H}_i$ for $i \notin K$. In addition, if $i \in K$, the sets $\arg\max(W^+_i \odot x) \cap J$ and $\arg\max(W^-_i \odot x) \cap J$ are both nonempty if $i \in C_1$, and both empty if $i \in C_2$. Thus, e^J is feasible in the s-hyperplane $\mathcal{H}_i$.

Similarly, let J' be the set of coordinate nodes weakly connected to k in $\vec{\mathcal{G}}_x \setminus \{a'\}$. Then the direction $e^{J'}$ is also feasible. Note that J and J' are neither equal nor complementary. Thus, there are two distinct and nonopposite directions which are feasible from x in $\mathcal{E}_K$, which implies that x is a breakpoint. $\square$

Example 4.4. Figure 4 depicts the tangent digraphs at every point of the tropical edge $\mathcal{E}_K$ for $K = \{\mathcal{H}_1,\mathcal{H}_2\}$, and this illustrates Proposition 4.3. The set $I = \{\mathcal{H}_1,\mathcal{H}_2,\mathcal{H}_3\}$ of constraints determines the basic point $x^I = (1,0,0)$. From its tangent digraph, we deduce that the initial ordinary segment of the edge $\mathcal{E}_K$ is directed by $e^{\{2\}}$.

The tangent digraph at a point in $](1,1,0),(1,0,0)[$ has exactly two weakly connected components. They yield the feasible directions $e^{\{2\}}$ and $e^{\{1,3,4\}}$, which correspond to the vectors $(0,1,0)$ and $(0,-1,0)$ of $\mathbb{T}^3$.

At the breakpoint $(1,1,0)$, the tangent digraph is weakly connected, and the hyperplane node $\mathcal{H}_1$ has degree $(2,1)$. Removing the arc from coordinate node 4 to $\mathcal{H}_1$ provides two weakly connected components, respectively, $\{1,2\} \cup \{\mathcal{H}_1\}$ and $\{3,4\} \cup \{\mathcal{H}_2\}$. The coordinate nodes of the component containing $\mathcal{H}_1$ yields the feasible direction $e^{\{1,2\}}$. Similarly, it can be verified that the other feasible direction, obtained by removing the arc from coordinate node 2, is the vector $e^{\{1,3,4\}}$.

4.3. Moving along an ordinary segment. In this section we provide exact details on how to obtain the next point ξ' from a given point ξ and a direction given in terms of the set J. We determine whether ξ' is a basic point or a breakpoint, and we determine the length μ of the resulting segment $[\xi,\xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ of the tropical edge $\mathcal{E}_K$. The metric results from this section will be interpreted in terms of tangent digraphs in the next section.

For all $i \in [m]$, we define

$$\begin{aligned} \lambda_i^+(\xi,J) &:= (W_i^+ \odot \xi) - \max_{j\in J}(w_{ij}^+ + \xi_j) \ , \\ \lambda_i^-(\xi,J) &:= (W_i^+ \odot \xi) - \max_{j\in J}(w_{ij}^- + \xi_j) \ , \end{aligned}$$

where $W = (w_{ij})$. When it is clear from the context, $\lambda_i^+(\xi,J)$ and $\lambda_i^-(\xi,J)$ will be simply denoted by λ_i^+ and λ_i^-. By Assumptions 1 and 5, we have $W_i^+ \odot \xi > \mathbb{0}$. In contrast, $\max_{j\in J}(w_{ij}^+ + \xi_j)$ and $\max_{j\in J}(w_{ij}^- + \xi_j)$ may be equal to $-\infty$, in which case we use the convention $-(-\infty) = +\infty$, and so $\lambda_i^+ = +\infty$ and $\lambda_i^- = +\infty$, respectively. When $\max_{j\in J}(w_{ij}^+ + \xi_j)$ and $\max_{j\in J}(w_{ij}^- + \xi_j)$ are finite, the scalars λ_i^+ and λ_i^- are nonnegative real numbers.

Let $x^\lambda := \xi + \lambda e^J$. Observe that λ_i^+ is the smallest $\lambda \geqslant 0$ such that $W_i^+ \odot \xi = w_{ij}^+ + x_j^\lambda$ for some $j \in J$. Similarly, λ^- is the smallest $\lambda \geqslant 0$ such that $W_i^+ \odot \xi = w_{ij}^- + x_j^\lambda$ for some $j \in J$. More precisely, we have

$$\begin{aligned} W_i^+ \odot x^\lambda &= \begin{cases} W_i^+ \odot \xi & \text{if } 0 \leqslant \lambda \leqslant \lambda_i^+, \\ (W_i^+ \odot \xi) + \lambda - \lambda_i^+ & \text{if } \lambda \geqslant \lambda_i^+, \end{cases} \\ W_i^- \odot x^\lambda &= \begin{cases} W_i^- \odot \xi & \text{if } 0 \leqslant \lambda \leqslant \beta_i^-, \\ (W_i^+ \odot \xi) + \lambda - \lambda_i^- & \text{if } \lambda \geqslant \beta_i^- \ , \end{cases} \end{aligned}$$

where $\beta_i^- = \lambda_i^- + (W_i^- \odot \xi) - (W_i^+ \odot \xi)$. In particular $\beta_i^- \leqslant \lambda_i^-$ and equality holds when $i \in K$. The evolution of $W_i^+ \odot (\xi + \lambda e^J)$ versus $W_i^- \odot (\xi + \lambda e^J)$ is visualized in Figure 7.

The endpoint ξ' is either a breakpoint or a basic point. We will prove that it is a basic point if a new hyperplane node $i_{\mathsf{ent}} \notin K$ "appears" in the tangent digraph. In that case the index i_{ent} must belong to the following set:

$$\mathsf{Ent}(\xi,J) := \{i \in [m] \setminus K \mid \arg\max(W_i^+ \odot \xi) \cap J = \emptyset\} \ .$$

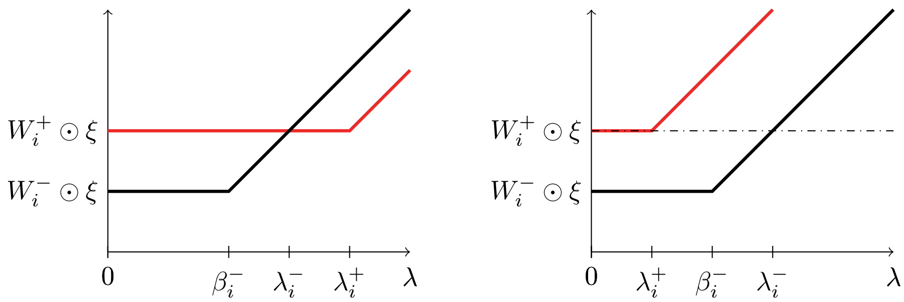

FIG. 7. *Evolution of $W_i^+ \odot (\xi + \lambda e^J)$ (in red) and $W_i^- \odot (\xi + \lambda e^J)$ (in black) with $\lambda \geqslant 0$ for $i \in \mathsf{Ent}$ and $\lambda_i^- < \lambda_i^+$ (left) or $\lambda_i^- > \lambda_i^+$ (right).*

We shall see that ξ' is a breakpoint if a hyperplane node $k \in K$ "acquires" a new arc and thus becomes of degree $(2, 1)$ or $(1, 2)$. Such a node k must be an element of the following set:

$$\mathsf{Br}(\xi, J) := \{i \in K \mid \arg\max(W_i^+ \odot \xi) \cap J = \emptyset \text{ and } \arg\max(W_i^- \odot \xi) \cap J = \emptyset\} .$$

We already mentioned that the notation i_{ent} (and so, $\mathsf{Ent}(\xi, J)$) and i_{out} is chosen by analogy with the entering or leaving indices in the classical simplex method. Note that the set $\mathsf{Br}(\xi, J)$ does not have any classical analogue. It represents intermediate indices which shall be examined before a leaving index is found.

When this does not bear the risk of confusion, we simply use the notation Br and Ent.

PROPOSITION 4.5. *Let $\{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ be an ordinary segment of a tropical edge $\mathcal{E}_K$. The following properties hold:*

(i) *The length μ of the segment is the greatest scalar $\lambda \geqslant 0$ satisfying the following conditions:*

$$\begin{aligned} &\lambda \leqslant \min(\lambda_i^+, \lambda_i^-) && \text{for all } i \in \mathsf{Br}\,, \\ &\lambda \leqslant \lambda_i^- && \text{for all } i \in \mathsf{Ent} \text{ such that } \lambda_i^- \leqslant \lambda_i^+. \end{aligned} \tag{4.3}$$

(ii) *If $\mu = \lambda_{i_{\mathsf{ent}}}^-$ for $i_{\mathsf{ent}} \in \mathsf{Ent}$, then $\xi + \mu e^J$ is a basic point for the basis $K \cup \{i_{\mathsf{ent}}\}$.*

(iii) *If $\mu = \min(\lambda_k^+, \lambda_k^-)$ for $k \in \mathsf{Br}$, then $\xi + \mu e^J$ is a breakpoint.*

Proof. Let $x^\lambda := \xi + \lambda e^J$ for all $\lambda \geqslant 0$.

We claim that x^λ belongs to $\mathcal{E}_K$ if λ satisfies (4.3). To that end, we first show that $x^\lambda \in \mathcal{H}_i$ for $i \in K$. Consider an $i \in \mathsf{Br}$. Then $\beta_i^- = \lambda_i^-$. Therefore, for all $0 \leqslant \lambda \leqslant \min(\lambda_i^+, \lambda_i^-)$ we have $x^\lambda \in \mathcal{H}_i$ since

$$W_i^+ \odot x^\lambda = W_i^+ \odot \xi = W_i^- \odot \xi = W_i^- \odot x^\lambda .$$

Let $i \in K \setminus \mathsf{Br}$. Then by Lemma 4.2, $\arg\max(W_i^+ \odot \xi) \cap J$ and $\arg\max(W_i^- \odot \xi) \cap J$ are both nonempty. Thus $\lambda_i^+ = \lambda_i^- = \beta_i^- = 0$. Therefore, $x^\lambda \in \mathcal{H}_i$ for all $\lambda \geqslant 0$ since in this case

$$W_i^+ \odot x^\lambda = (W_i^+ \odot \xi) + \lambda = W_i^- \odot x^\lambda .$$

We now examine the half-spaces $\mathcal{H}_i^{\geqslant}$, where $i \in [m] \setminus K$. If $i \notin \mathsf{Ent}$, then $\arg\max(W_i^+ \odot \xi) \cap J \neq \emptyset$. Consequently, $\lambda_i^+ = 0$. Thus $x^\lambda \in \mathcal{H}_i^{\geqslant}$ for all $\lambda \geqslant 0$ as we have

$$W_i^+ \odot x^\lambda = (W_i^+ \odot \xi) + \lambda \geqslant \max(W_i^- \odot \xi, (W_i^+ \odot \xi) + \lambda - \lambda_i^-) = W_i^- \odot x^\lambda .$$

If $i \in \mathsf{Ent}$ and $0 \leqslant \lambda \leqslant \min(\lambda_i^+, \lambda_i^-)$, then $x^\lambda \in \mathcal{H}_i^{\geqslant}$. Indeed,

$$W_i^+ \odot x^\lambda = W_i^+ \odot \xi \geqslant \max(W_i^- \odot \xi, (W_i^+ \odot \xi) + \lambda - \lambda_i^-) = W_i^- \odot x^\lambda .$$

Now if further $\lambda_i^+ < \lambda_i^-$, then, for $\lambda \geqslant \lambda_i^+$, we have

$$W_i^+ \odot x^\lambda = (W_i^+ \odot \xi) + \lambda - \lambda_i^+ \geqslant \max(W_i^- \odot \xi, (W_i^+ \odot \xi) + \lambda - \lambda_i^-) = W_i^- \odot x^\lambda .$$

We conclude that if $i \in \mathsf{Ent}$ and $\lambda_i^+ < \lambda_i^-$, then $x^\lambda \in \mathcal{H}_i^{\geqslant}$ for all $\lambda \geqslant 0$.

Second, we claim that the solution set of the inequalities (4.3) admits a greatest element $\lambda^* \in \mathbb{R}$. By contradiction, suppose that $x^\lambda \in \mathcal{E}_K$ for all $\lambda \geqslant 0$. Recall that e^J and $-e^{[n+1]\setminus J}$ coincide as elements of $\mathbb{TP}^n$. Consequently the half-ray $\{\xi - \lambda e^{[n+1]\setminus J} \mid \lambda \geqslant 0\}$ is contained in $\mathcal{E}_K$ and thus in $\mathcal{C}$. Since $\mathcal{C}$ is closed, it contains the point $y \in \mathbb{T}^{n+1}$ defined by $y_j = \xi_j$ if $j \in J$ and $y_j = \mathbb{0}$ otherwise. As $J \subsetneq [n+1]$, this contradicts Assumption 5.

Third, we claim that $\lambda^* = \mu$. To prove the claim it is sufficient to show that x^{λ^*} is either a breakpoint or a basic point of $\mathcal{E}_K$. We distinguish three cases:

(a) $\lambda^* = \lambda_i^-$ for some $i \in \mathsf{Ent}$. Then $W_i^- \odot x^{\lambda^*} = W_i^+ \odot \xi$. Moreover $\lambda^* \leqslant \lambda_i^+$ and thus $W_i^+ \odot x^{\lambda^*} = W_i^+ \odot \xi$. This implies that $W_i^+ \odot x^{\lambda^*} = W_i^- \odot x^{\lambda^*} > \mathbb{0}$. As a consequence, $i \notin K$ is a hyperplane node in the tangent graph $\mathcal{G}_{x^{\lambda^*}}$. By Proposition 4.3, we conclude that x^{λ^*} is a basic point for the set $K \cup \{i\}$.

(b) $\lambda^* = \lambda_i^+ \leqslant \lambda_i^-$ for some $i \in \mathsf{Br}$. Then, observe that

$$\arg\max(W_i^+ \odot x^{\lambda^*}) = \arg\max(W_i^+ \odot \xi) \cup \arg\max_{j \in J}(w_{ij}^+ + \xi_j) . \tag{4.4}$$

The two sets on the right-hand side of (4.4) are nonempty and disjoint, since $i \in \mathsf{Br}$. Thus $\arg\max(W_i^+ \odot x^{\lambda^*})$ contains at least two distinct elements. Moreover, $x^{\lambda^*} \in \mathcal{E}_K$ by the discussion above, and thus $i \in K$ appears as a hyperplane node in $\vec{\mathcal{G}}_{x^{\lambda^*}}$. Consequently, the hyperplane node i has at least two incoming arcs in $\vec{\mathcal{G}}_{x^{\lambda^*}}$. We deduce by Proposition 4.3 that the degree of the hyperplane node i in $\vec{\mathcal{G}}_{x^{\lambda^*}}$ is $(2,1)$ and that x^{λ^*} is a breakpoint.

(c) $\lambda^* = \lambda_i^- \leqslant \lambda_i^+$ for some $i \in \mathsf{Br}$. By the same argument as above, $\arg\max(W_i^- \odot x^{\lambda^*})$ contains at least two distinct elements. This implies that x^{λ^*} is a breakpoint and that the hyperplane node i has degree $(1,2)$.

Note that the arguments above also prove Proposition 4.5(ii) and (iii). ∎

Remark 4.6. When $\xi + \mu e^J$ is a breakpoint, the proof of Proposition 4.5 ensures that the hyperplane node k in the tangent digraph $\vec{\mathcal{G}}_{\xi+\mu e^J}$ has degree $(2,1)$ if $\mu = \lambda_k^+$ or $(1,2)$ if $\mu = \lambda_k^-$. In particular, this proves that μ is equal to only one scalar among the λ_i^-, λ_k^+ and λ_k^-, where $i \in \mathsf{Ent}$ and $k \in \mathsf{Br}$.

Example 4.7. We now have all the ingredients required to perform a tropical pivot. Feasible directions are given by Proposition 4.3, while Proposition 4.5 provides the lengths of ordinary segments and the stopping criterion.

Let us illustrate this on our running example. We start from the basic point $(4,4,2)$ (i.e., the point $(4,4,2,0)$ in $\mathbb{TP}^3$) given by $I = \{\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_5\}$, and we move along the edge $\mathcal{E}_K$, where $K = \{\mathcal{H}_1, \mathcal{H}_2\}$. The tangent digraph at $(4,4,2)$ is depicted in the bottom right of Figure 4. By Proposition 4.3(C1), the initial direction is $-e^{\{1,2,3\}}$, i.e., $J = \{4\}$. By definition, Br is formed by the hyperplane nodes which are not adjacent to the coordinate node 4 in the tangent digraph. Hence, $\mathsf{Br} = \{\mathcal{H}_1, \mathcal{H}_2\}$. Moreover, in the homogeneous setting, the inequalities $\mathcal{H}_3$ and $\mathcal{H}_4$ read

$$\begin{aligned} &x_2 \geqslant x_4, \\ &x_1 \geqslant \max(x_4, x_2 - 3). \end{aligned}$$

In both of them, the maximum in the left-hand side is reduced to one term, and it does not involve x_4. Thus, $\mathsf{Ent} = \{\mathcal{H}_3, \mathcal{H}_4\}$. The reader can verify that

$$\begin{aligned} \lambda^+_{\mathcal{H}_1} &= 3 - 0 = 3, & \lambda^-_{\mathcal{H}_1} &= 3 - (-\infty) = +\infty, \\ \lambda^+_{\mathcal{H}_2} &= 2 - (-\infty) = +\infty, & \lambda^-_{\mathcal{H}_2} &= 2 - 0 = 2, \\ \lambda^+_{\mathcal{H}_3} &= 4 - (-\infty) = +\infty, & \lambda^-_{\mathcal{H}_3} &= 4 - 0 = 4, \\ \lambda^+_{\mathcal{H}_4} &= 4 - (-\infty) = +\infty, & \lambda^-_{\mathcal{H}_4} &= 4 - 0 = 4. \end{aligned}$$

As a result, the length of the initial ordinary segment is $\mu = 2$, given by $\mu = \lambda^-_{\mathcal{H}_2} \leqslant \lambda^+_{\mathcal{H}_2}$. As $\mathcal{H}_2 \in \mathsf{Br}$, the point $(4, 4, 2) - 2e^{\{1,2,3\}} = (2, 2, 0)$ is a breakpoint.

The next feasible direction is $-e^{\{1,2\}}$ as $J = \{3, 4\}$. We still have $\mathsf{Ent} = \{\mathcal{H}_3, \mathcal{H}_4\}$ but now $\mathsf{Br} = \{\mathcal{H}_1\}$. The length of this ordinary segment is $\mu = 1 = \lambda^+_{\mathcal{H}_1}$. Consequently, we reach the breakpoint $(1, 1, 0) = (2, 2, 0) - 1e^{\{1,2\}}$, where the next feasible direction, $-e^{\{2\}}$, is given by $J = \{1, 3, 4\}$. The set Br is now empty and $\mathsf{Ent} = \{\mathcal{H}_4\}$. Clearly, $\mu = 1 = \lambda^-_{\mathcal{H}_4}$. As $\mathcal{H}_4 \in \mathsf{Ent}$, the next endpoint $(1, 0, 0) = (1, 1, 0) - 1e^{\{2\}}$ is a basic point.

4.4. Efficient implementation of the pivoting operation. Our implementation of the pivoting operation relies on the incremental update of the tangent digraph along the tropical edge. This avoids computing from scratch the tangent digraph at each breakpoint, in which case the time complexity of the pivoting operation would be naively in $O(n^2 m)$.

In the previous section we described the "travel" from a given point ξ into the direction given by J to the next point, called ξ'. Our key observation is that the tangent digraph is constant in the open segment $]\xi, \xi'[$ and that it "acquires" a new arc or a new hyperplane node when the endpoint ξ' is reached. This is made precise in the lemma below and the subsequent proposition.

LEMMA 4.8. *Let $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ be an ordinary segment of $\mathcal{E}_K$. Every point in $]\xi, \xi'[$ has the same tangent digraph $\vec{\mathcal{G}}_{]\xi,\xi'[}$, which is equal to the intersection of $\vec{\mathcal{G}}_\xi$ and $\vec{\mathcal{G}}_{\xi'}$.*

Proof. Let $x^\lambda := \xi + \lambda e^J$. By Proposition 4.3, the hyperplane node set of $\vec{\mathcal{G}}_{x^\lambda}$ for $\lambda \in]0, \mu[$ is equal to K. If ξ and ξ' are both basic points, the sets of hyperplane nodes in their tangent digraphs are respectively of the form $K \cup \{i_{\mathsf{out}}\}$ and $K \cup \{i_{\mathsf{ent}}\}$, where $i_{\mathsf{out}}, i_{\mathsf{ent}} \notin K$, and $i_{\mathsf{out}} \neq i_{\mathsf{ent}}$. If one of the two endpoints, say, ξ, is a breakpoint, the hyperplane node set of its tangent digraph is K, while the hyperplane node set of $\vec{\mathcal{G}}_{\xi'}$ contains K. In all cases, the intersection of the hyperplane node sets of $\vec{\mathcal{G}}_\xi$ and $\vec{\mathcal{G}}_{\xi'}$ is equal to K. Moreover, the coordinate node set of $\vec{\mathcal{G}}_x$ is equal to $[n + 1]$ for all $x \in [\xi, \xi']$.

Let $i \in \mathsf{Br}$. If $0 < \lambda < \mu$, then in particular $\lambda < \min(\lambda^+_i, \lambda^-_i)$ by Proposition 4.5. Hence,

$$\arg\max(W^\pm_i \odot x^\lambda) = \arg\max(W^\pm_i \odot \xi) \ . \tag{4.5}$$

Besides, $\arg\max(W^\pm_i \odot \xi') = \arg\max(W^\pm_i \odot x^\mu)$ is a superset of $\arg\max(W^\pm_i \odot \xi)$, and the inclusion is strict when μ is equal to the corresponding scalar λ^+_i or λ^-_i.

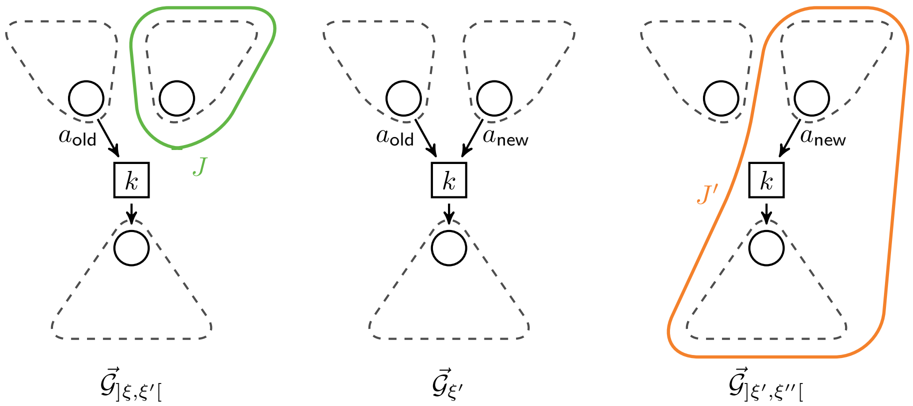

FIG. 8. *Illustration of Proposition* 4.9(ii) *and Remark* 4.10, *with a sequence of tangent digraphs around a breakpoint ξ' between two consecutive segments $[\xi,\xi'] \cup [\xi',\xi'']$. The direction of $[\xi,\xi']$, from ξ to ξ', is given by the set of coordinate nodes J, indicated in green. The direction of the second segment, from ξ' to ξ'', is governed by J' depicted in orange.*

Similarly, let $i \in K \setminus \mathsf{Br}$. By Lemma 4.2, $\arg\max(W_i^+ \odot \xi) \cap J$ and $\arg\max(W_i^- \odot \xi) \cap J$ are both nonempty. Moreover, for all $\lambda > 0$, we have

$$\arg\max(W_i^\pm \odot x^\lambda) = \arg\max(W_i^\pm \odot \xi) \cap J \ . \tag{4.6}$$

In particular, $\arg\max(W_i^\pm \odot \xi') = \arg\max(W_i^\pm \odot \xi) \cap J$.

Equations (4.5) and (4.6) ensure that $\arg\max(W_i^\pm \odot x^\lambda) = \arg\max(W_i^\pm \odot \xi) \cap \arg\max(W_i^\pm \odot \xi')$ for all $i \in K$ and $\lambda \in]0,\mu[$. This shows that the arc set of $\vec{\mathcal{G}}_{x^\lambda}$ is precisely the intersection of the arc sets of $\vec{\mathcal{G}}_\xi$ and $\vec{\mathcal{G}}_{\xi'}$. $\square$

PROPOSITION 4.9. *Let $[\xi,\xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ be an ordinary segment of $\mathcal{E}_K$.*

(i) *If ξ is a basic point,* i.e., $\xi = x^{K\cup\{i_{\mathsf{out}}\}}$ for a given $i_{\mathsf{out}} \notin K$, then

$$\vec{\mathcal{G}}_{]\xi,\xi'[} = \vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\} \ .$$

(ii) *If* ξ' is a breakpoint and k the hyperplane node of $\vec{\mathcal{G}}_{\xi'}$ with degree $(2,1)$ or $(1,2)$, then

$$\vec{\mathcal{G}}_{\xi'} = \vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{a_{\mathsf{new}}\} \ ,$$

where a_{new} is an arc between k and the unique element of $\arg\max_{j\in J}(|w_{kj}| + \xi_j)$.

Moreover, if $[\xi',\xi'']$ is the next ordinary segment in $\mathcal{E}_K$, then

$$\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\},$$

where a_{old} is the unique arc incident to k with the same orientation as a_{new} in $\vec{\mathcal{G}}_{\xi'}$.

An illustration of (ii) is given in Figure 8.

Proof. Let $x^\lambda := \xi + \lambda e^J$.

(i) By Proposition 4.3(C2), the tangent digraph $\vec{\mathcal{G}}_{]\xi,\xi'[}$ does not contain the hyperplane node i_{out}. As $\vec{\mathcal{G}}_{]\xi,\xi'[}$ is a subdigraph of $\vec{\mathcal{G}}_\xi$ by Lemma 4.8, we deduce that it is also a subdigraph of $\vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\}$. Since $\vec{\mathcal{G}}_{]\xi,\xi'[}$ and $\vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\}$ have the same number of nodes and arcs by Proposition 4.3, we conclude that they are equal.

(ii) We assume that k has degree $(2,1)$ in $\vec{\mathcal{G}}_{\xi'}$, the proof being similar when k has degree $(1,2)$. To begin with, we know that $\mu = \lambda_k^+(\xi, J)$ thanks to Remark 4.6. Let $l \in \arg\max_{j \in J}(w_{kj}^+ + \xi_j)$. Then for all $0 < \lambda < \lambda_k^+(\xi, J)$, we have $w_{kl}^+ + x_l^\lambda < W_k^+ \odot x^\lambda$, while $w_{kl}^+ + x_l^\mu = W_k^+ \odot x^\mu$. It follows that the arc (l,k) does not belong to $\vec{\mathcal{G}}_{]\xi,\xi'[}$, whereas it appears in $\vec{\mathcal{G}}_{\xi'}$. We deduce that $\vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{(l,k)\}$ is a subgraph of $\vec{\mathcal{G}}_{\xi'}$ thanks to Lemma 4.8. Both are equal by Proposition 4.3. Note that $\arg\max_{j \in J}(w_{kj}^+ + \xi_j)$ is reduced to $\{l\}$ as k has two incoming arcs in $\vec{\mathcal{G}}_{\xi'}$. Due to Remark 4.6 we have $\lambda_k^+(\xi, J) < \lambda_k^-(\xi, J)$. It follows that

$$\arg\max_{j \in J}(|w_{kj}| + \xi_j) = \arg\max_{j \in J}(w_{kj}^+ + \xi_j) = \{l\} \ .$$

In the second place, by applying Lemma 4.8 to the segment $[\xi', \xi'']$, we know that $\vec{\mathcal{G}}_{]\xi',\xi''[}$ is a subdigraph of $\vec{\mathcal{G}}_{\xi'}$. By Proposition 4.3, the hyperplane node k has degree $(1,1)$ in $\vec{\mathcal{G}}_{]\xi',\xi''[}$. Thus, the digraph $\vec{\mathcal{G}}_{]\xi',\xi''[}$ is either equal to $\vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{new}}\}$ or $\vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$. As the former corresponds to the tangent digraph $\vec{\mathcal{G}}_{]\xi,\xi'[}$, we deduce that $\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$. ∎

Remark 4.10. We point out that in Proposition 4.9(ii), the set J' corresponding to the direction of the next segment $[\xi', \xi'']$ is precisely given by the set of coordinate nodes weakly connected to k in the digraph $\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$; see Figure 8 for an illustration.

Indeed, according to Proposition 4.3(C2), the digraph $\vec{\mathcal{G}}_{]\xi',\xi''[}$ consists of two weakly connected components. Let $\mathfrak{J}$ be the set of coordinate nodes of the component containing the hyperplane node k. From any point in $]\xi', \xi''[$, the two feasible directions are $\pm e^{\mathfrak{J}}$. As a result, $J' = \mathfrak{J}$ or $J' = [n+1] \setminus \mathfrak{J}$. Let l be the coordinate node incident to a_{new}. Then $l \in J$ by Proposition 4.9(ii), and so $l \in J'$ as $J \subset J'$. In addition, since a_{new} still appears in $\vec{\mathcal{G}}_{]\xi',\xi''[}$, the coordinate node l is weakly connected to k. Therefore, $l \in \mathfrak{J}$. We conclude that $J' = \mathfrak{J}$, as expected.

The following proposition allows us to incrementally maintain the sets Ent, Br and the associated scalars $\lambda_i^{\pm}$ along the tropical edge $\mathcal{E}_K$.

PROPOSITION 4.11. *Let $[\xi, \xi'] \cup [\xi', \xi'']$ be two consecutive ordinary segments of $\mathcal{E}_K$, where $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ and $[\xi', \xi''] = \{\xi' + \lambda e^{J'} \mid 0 \leqslant \lambda \leqslant \mu'\}$. Then,*

(i) $\mathsf{Br}(\xi', J') \subset \mathsf{Br}(\xi, J)$;

(ii) $\arg\max(W_i^+ \odot \xi') = \arg\max(W_i^+ \odot \xi)$ *for all* $i \in \mathsf{Ent}(\xi', J')$;

(iii) $\mathsf{Ent}(\xi', J') = \{i \in \mathsf{Ent}(\xi, J) \mid \mu < \lambda_i^+(\xi, J)$ *and* $\arg\max(W_i^+ \odot \xi) \cap (J' \setminus J) = \emptyset\}$;

(iv) *for all* $i \in \mathsf{Ent}(\xi', J') \cup \mathsf{Br}(\xi', J')$, *we have*

$$W_i^+ \odot \xi' = W_i^+ \odot \xi$$

$$\lambda_i^+(\xi', J') = \min\left(\lambda_i^+(\xi, J) - \mu \ , \ (W_i^+ \odot \xi) - \max_{j \in J' \setminus J}(w_{ij}^+ + \xi_j)\right),$$

$$\lambda_i^-(\xi', J') = \min\left(\lambda_i^-(\xi, J) - \mu \ , \ (W_i^+ \odot \xi) - \max_{j \in J' \setminus J}(w_{ij}^- + \xi_j)\right).$$

Proof.

(i) Suppose by contradiction that $i \notin \mathsf{Br}(\xi, J)$. Then, by Lemma 4.2 the intersections $\arg\max(W_i^+ \odot \xi) \cap J$ and $\arg\max(W_i^- \odot \xi) \cap J$ are both nonempty. As a consequence, $\arg\max(W_i^+ \odot \xi')$ and $\arg\max(W_i^- \odot \xi')$ are included in J. Since $J \subset J'$ by Proposition 4.1, we conclude that $i \notin \mathsf{Br}(\xi', J')$.

(ii) First observe that $\mathsf{Ent}(\xi', J') \subset \mathsf{Ent}(\xi, J)$. Indeed, consider an $i \in K \setminus \mathsf{Ent}(\xi, J)$. Then $\arg\max(W_i^+ \odot \xi) \cap J \neq \emptyset$, which implies $\arg\max(W_i^+ \odot \xi') \subset J$. Using the inclusion $J \subset J'$, we obtain that $\arg\max(W_i^+ \odot \xi') \cap J' \neq \emptyset$, and therefore $i \notin \mathsf{Ent}(\xi', J')$.

Second if $i \in \mathsf{Ent}(\xi, J)$ satisfies $\mu \geqslant \lambda_i^+(\xi, J)$, then $\arg\max(W_i^+ \odot \xi')$ intersects $J \subset J'$, thus $i \notin \mathsf{Ent}(\xi', J')$. As a consequence,

$$\mathsf{Ent}(\xi', J') \subset \{i \in \mathsf{Ent}(\xi, J) \mid \mu < \lambda_i^+(\xi, J)\}\,. \tag{4.7}$$

Finally for any $i \in \mathsf{Ent}(\xi', J')$, we have $\mu < \lambda_i^+(\xi, J)$ and therefore $\arg\max(W_i^+ \odot \xi') = \arg\max(W_i^+ \odot \xi)$.

(iii) Using (4.7) let us consider an $i \in \mathsf{Ent}(\xi, J)$ such that $\mu < \lambda_i^+(\xi, J)$. Then, as above, $\arg\max(W_i^+ \odot \xi') = \arg\max(W_i^+ \odot \xi)$. Moreover, $i \in \mathsf{Ent}(\xi, J)$ implies $\arg\max(W_i^+ \odot \xi) \cap J = \emptyset$. Thus $\arg\max(W_i^+ \odot \xi') \cap J' = \emptyset$ if and only if $\arg\max(W_i^+ \odot \xi) \cap (J' \setminus J) = \emptyset$.

(iv) Consider $i \in \mathsf{Ent}(\xi', J') \cup \mathsf{Br}(\xi', J')$. If $i \in \mathsf{Ent}(\xi', J')$, then $\mu < \lambda_i^+(\xi, J)$ by (4.7). Otherwise, if $i \in \mathsf{Br}(\xi', J')$, then $i \in \mathsf{Br}(\xi, J)$ by Proposition (4.11)(i) and thus $\mu \leqslant \lambda_i^+(\xi, J)$ by (4.3). In both cases, we obtain $W_i^+ \odot \xi' = W_i^+ \odot \xi$.

Let us rewrite $\lambda_i^+(\xi', J')$ as follows:

$$\lambda_i^+(\xi', J') = \min\left((W_i^+ \odot \xi') - \max_{j \in J}(w_{ij}^+ + \xi'_j)\ ,\ (W_i^+ \odot \xi') - \max_{j \in J' \setminus J}(w_{ij}^+ + \xi'_j)\right).$$

We saw that $W_i^+ \odot \xi' = W_i^+ \odot \xi$. Furthermore, $\xi'_j = \xi_j + \mu$ if $j \in J$ and $\xi'_j = \xi_j$ otherwise. Thus the first term of the minimum above is equal to

$$(W_i^+ \odot \xi) - \max_{j \in J}(w_{ij}^+ + \xi + \mu) = \lambda_i^+(\xi, J) - \mu.$$

The second term satisfies

$$(W_i^+ \odot \xi') - \max_{j \in J' \setminus J}(w_{ij}^+ + \xi'_j) = (W_i^+ \odot \xi) - \max_{j \in J' \setminus J}(w_{ij}^+ + \xi_j)\,.$$

The same argument holds for $\lambda_i^-(\xi', J')$. ∎

We now present an algorithm (Algorithm 2) allowing us to move along an ordinary segment $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$ of the tropical edge $\mathcal{E}_K$. This algorithm takes as input the initial endpoint ξ, together with some auxiliary data, including the set J encoding the direction of the segment $[\xi, \xi']$, the tangent digraph in $]\xi, \xi'[$, and the sets $\mathsf{Ent}(\xi, J)$ and $\mathsf{Br}(\xi, J)$. It also uses an auxiliary function Ω, which is defined for the pairs $(i, j) \in \mathsf{Ent}(\xi, J) \times [n+1]$ and which returns in time $O(1)$ whether $j \in \arg\max(W_i^+ \odot \xi)$. We shall see in the main pivoting algorithm that this function is defined once and for all when pivoting over the whole tropical edge.

Algorithm 2 returns the other endpoint ξ'. On top of that, if ξ' is a breakpoint of $\mathcal{E}_K$, it provides the set J' corresponding to the direction of the next ordinary segment $[\xi', \xi'']$ of $\mathcal{E}_K$, some additional data corresponding to ξ', J' (for instance the sets $\mathsf{Ent}(\xi', J')$ and $\mathsf{Br}(\xi', J')$), and the digraph $\vec{\mathcal{G}}_{]\xi', \xi''[}$.

Several kinds of data structures are manipulated in Algorithm 2, and we need to specify the complexity of the underlying operations. Arithmetic operations over $\mathbb{T}$ are supposed to be done in time $O(1)$. Tangent digraphs are represented by adjacency lists. They are of size $O(n)$, and so they can be visited in time $O(n)$. Matrices are stored as two-dimensional arrays, so an arbitrary entry can be accessed in $O(1)$.

Vectors and the values $W_i^+ \odot \xi$, $\lambda_i^+(\xi, J)$ and $\lambda_i^-(\xi, J)$ for $i \in [m]$ are stored as arrays of scalars. Apart from $\Delta = J' \setminus J$, sets are represented as Boolean arrays, so that testing membership takes $O(1)$. The set Δ is stored as a list, and thus iterating over its elements can be done in $O(|\Delta|)$.

Algorithm 2. Traversal of an ordinary segment of an tropical edge.

Input: An endpoint ξ of an ordinary segment $[\xi, \xi']$ of a tropical edge $\mathcal{E}_K$ and
- the set J encoding the direction e^J of $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leqslant \lambda \leqslant \mu\}$
- the tangent digraph $\vec{\mathcal{G}}_{]\xi,\xi'[}$ in the relative interior of $[\xi, \xi']$
- the sets $\mathsf{Ent}(\xi, J)$ and $\mathsf{Br}(\xi, J)$
- the scalars $W_i^+ \odot \xi$, $\lambda_i^+(\xi, J)$ and $\lambda_i^-(\xi, J)$ for $i \in \mathsf{Br}(\xi, J) \cup \mathsf{Ent}(\xi, J)$
- an auxiliary function Ω determining in time $O(1)$ if $j \in \arg\max(W_i^+ \odot \xi)$ for all $i \in \mathsf{Ent}(\xi, J)$ and $j \in [n+1]$

Output: The other endpoint ξ' and,
if ξ' is a basic point, the integer $i_{\mathsf{ent}} \notin K$ such that $\xi' = x^{K \cup \{i_{\mathsf{ent}}\}}$;
if ξ' is a breakpoint:
- the set J' encoding the direction $e^{J'}$ of the next ordinary segment $[\xi', \xi''] = \{\xi' + \lambda e^{J'} \mid 0 \leqslant \lambda \leqslant \mu'\}$
- the tangent digraph $\vec{\mathcal{G}}_{]\xi',\xi''[}$
- the sets $\mathsf{Ent}(\xi', J')$ and $\mathsf{Br}(\xi', J')$
- the scalars $W_i^+ \odot \xi'$, $\lambda_i^+(\xi', J')$ and $\lambda_i^-(\xi', J')$ for $i \in \mathsf{Br}(\xi', J') \cup \mathsf{Ent}(\xi', J')$

Proposition 4.5(i)–(ii)

$\mu \leftarrow \min\{\min(\lambda_i^+(\xi, J), \lambda_i^-(\xi, J)) \mid i \in \mathsf{Br}(\xi, J)$ or $(i \in \mathsf{Ent}(\xi, J)$ and $\lambda_i^-(\xi, J) \leqslant \lambda_i^+(\xi, J))\}$ $O(m)$
$\xi' \leftarrow \xi + \mu e^J$ $O(n)$
if $\mu = \lambda_{i_{\mathsf{ent}}}^-(\xi, J)$ for some $i_{\mathsf{ent}} \in \mathsf{Ent}(\xi, J)$ **then**
  return (ξ', i_{ent}) `(`ξ' `is a basic point)`

Proposition 4.9(ii), Remark 4.6

$k \leftarrow$ the unique element of $\mathsf{Br}(\xi, J)$ such that $\mu = \min(\lambda_k^+(\xi, J), \lambda_k^-(\xi, J))$ `(`ξ' `is a breakpoint)`
$\ell \leftarrow$ the unique element in $\arg\max_{j \in J} |w_{kj}| + \xi_j$ $O(n)$
$a_{\mathsf{new}} \leftarrow$ the arc from ℓ to k if $\lambda_k^+(\xi, J) < \lambda_k^-(\xi, J)$, the arc from k to ℓ otherwise $O(1)$
$\vec{\mathcal{G}}_{\xi'} \leftarrow \vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{a_{\mathsf{new}}\}$ $O(n)$
compute $\mathsf{Br}(\xi', J')$ by visiting $\vec{\mathcal{G}}_{\xi'}$ $O(n)$

Remark 4.10

$a_{\mathsf{old}} \leftarrow$ the only arc incident to k in $\vec{\mathcal{G}}_{\xi'}$ with the same orientation as a_{new} $O(1)$
$\vec{\mathcal{G}}_{]\xi',\xi''[} \leftarrow \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$ $O(n)$
$J' \leftarrow$ coordinate nodes of $\vec{\mathcal{G}}_{]\xi',\xi''[}$ weakly connected to k $O(n)$

Proposition 4.11(iii)–(iv)

$\Delta \leftarrow$ the list of elements of $J' \setminus J$ $O(n)$
$\mathsf{Ent}(\xi', J') \leftarrow \{i \in \mathsf{Ent}(\xi, J) \mid \mu < \lambda_i^+(\xi, J)$ and $\arg\max(W_i^+ \odot \xi) \cap \Delta = \emptyset\}$ `using the function` Ω $O(m|J' \setminus J|)$
for $i \in \mathsf{Ent}(\xi', J') \cup \mathsf{Br}(\xi', J')$ **do** $O(m)$ `iterations`
  $W_i^+ \odot \xi' := W_i^+ \odot \xi$ $O(1)$
  $\lambda_i^+(\xi', J') := \min\Big(\lambda_i^+(\xi, J) - \mu\ ,\ (W_i^+ \odot \xi) - \max_{j \in \Delta}(w_{ij}^+ + \xi_j)\Big)$ $O(|J' \setminus J|)$
  $\lambda_i^-(\xi', J') := \min\Big(\lambda_i^+(\xi, J) - \mu\ ,\ (W_i^+ \odot \xi) - \max_{j \in \Delta}(w_{ij}^- + \xi_j)\Big)$ $O(|J' \setminus J|)$
return $(\xi', J', \vec{\mathcal{G}}_{]\xi',\xi''[}, \mathsf{Ent}(\xi', J'), \mathsf{Br}(\xi', J'), (W_i^+ \odot \xi')_i, (\lambda_i^+(\xi', J'))_i, (\lambda_i^-(\xi', J'))_i)$

Proposition 4.12. *Algorithm 2 is correct, and its time complexity is bounded by* $O(n + m|J' \setminus J|)$.

Proof.

Correctness. The correctness of the highlighted parts of the algorithm straightforwardly follows from the corresponding results given as annotations.

At line 9, the set $\mathsf{Br}(\xi', J')$ is built by iterating over the nodes of $\vec{\mathcal{G}}_{\xi'}$ and collecting the hyperplane nodes i with no neighbor in J'. This is correct since the set of hyperplane nodes is precisely K (by Proposition 4.3 and the fact that ξ' is a breakpoint), and because the adjacent nodes of each $i \in K$ are precisely the elements of $\arg\max(W_i^+ \odot \xi') \cup \arg\max(W_i^- \odot \xi')$ by construction of $\vec{\mathcal{G}}_{\xi'}$.

Complexity. At lines 8 and 11, the operations of removing or adding an arc can be performed in $O(n)$ by visiting the digraphs. Identifying the arc a_{old} at line 11 amounts to iterate over the arcs incident to k, and there are exactly three such arcs by Proposition 4.3.

Testing whether a hyperplane node i of $\vec{\mathcal{G}}_{\xi'}$ satisfies $\arg\max(W_i^+ \odot \xi') \cap J' = \arg\max(W_i^- \odot \xi') \cap J' = \emptyset$ can be done in $O(1)$, by determining whether the adjacent coordinate nodes (at most 3) in $\vec{\mathcal{G}}_{\xi'}$ belong to J'. Thus the set $\mathsf{Br}(\xi', J')$ can be built in time $O(n)$ by iterating over the hyperplane nodes of $\vec{\mathcal{G}}_{\xi'}$.

Given $i \in \mathsf{Ent}(\xi, J)$, determining whether $\arg\max(W_i^+ \odot \xi) \cap \varDelta = \emptyset$ can be performed by calling the auxiliary function $\varOmega$ for every element $j \in \varDelta$. It follows that $\mathsf{Ent}(\xi', J')$ can be computed at line 14 in time $O(m|J' \setminus J|)$.

Computations at lines 17 and 18 are done by iterating over elements $j \in \varDelta$ and then retrieving the values of $W_i^+ \odot \xi$, w_{ij}^+, w_{ij}^-, and ξ_j. Since these values are stored in arrays, they can be accessed in constant time. Therefore, $\lambda_i^+(\xi', J')$ and $\lambda_i^-(\xi', J')$ are computed in time $O(|\varDelta|) = O(|J' \setminus J|)$. The complexity of other operations is easily obtained. In total, the complexity of the algorithm is $O(n + m|J' \setminus J|)$.

ALGORITHM 3. LINEAR-TIME PIVOTING ALGORITHM.

Input: A basic point x^I of $\mathcal{P}(A, b)$, the associated set I, and an integer $i_{\mathsf{out}} \in I$
Output: The other basic point $x^{I'}$ of the edge $\mathcal{E}_{I \setminus \{i_{\mathsf{out}}\}}$, and the integer $i_{\mathsf{ent}} \in I \setminus \{i_{\mathsf{out}}\}$ such that $I' = (I \setminus \{i_{\mathsf{out}}\}) \cup \{i_{\mathsf{ent}}\}$

1 compute $\vec{\mathcal{G}}_{x^I}$ — $O(mn)$
2 $\vec{\mathcal{G}}_{]\xi^1,\xi^2[} \leftarrow \vec{\mathcal{G}}_{x^I} \setminus \{i_{\mathsf{out}}\}$ — $O(n)$
3 $J \leftarrow$ coordinate nodes weakly connected to the element of $\arg\max(W_{i_{\mathsf{out}}}^+ \odot x^I)$ in $\vec{\mathcal{G}}_{]\xi,\xi'[}$ — $O(n)$
4 compute $E \leftarrow \mathsf{Ent}(x^I, J)$ and $B \leftarrow \mathsf{Br}(x^I, J)$ — $O(mn)$
5 compute $W_i^+ \odot x^I$, $\lambda_i^+(x^I, J)$ and $\lambda_i^-(x^I, J)$ for all $i \in E \cup B$ — $O(mn)$
6 $\varOmega \leftarrow$ function defined on the set $E \times [n+1]$ by $\varOmega(i,j) = \begin{cases} \textbf{true} & \text{if } j \in \arg\max(W_i^+ \odot x^I) \\ \textbf{false} & \text{otherwise} \end{cases}$ — $O(mn)$
7 $\mathit{input} \leftarrow x^I, J, \vec{\mathcal{G}}_{]\xi^1,\xi^2[}, E, B, (W_i^+ \odot x)_{i \in E \cup B}, (\lambda_i^+(x^I, J))_{i \in E \cup B}, (\lambda_i^-(x^I, J))_{i \in E \cup B}$
8 **while true do** — `at most` n `iterations`
9 call Algorithm 2 on $(\mathit{input}, \varOmega)$ and stores the result in *output*
10 **if** *output* is of the form (ξ', i_{ent}) **then return** (ξ', i_{ent})
11 **else** $\mathit{input} \leftarrow \mathit{output}$

THEOREM 4.13. *Algorithm* 3 *allows us to pivot from a basic point along a tropical edge in time $O(n(m+n))$ and space $O(nm)$.*

Proof. First observe that the function $\varOmega$ initially defined at line 6 does not need to be updated during the iterations of the loop from lines 8 to 11. Indeed, let $[\xi, \xi']$

and $[\xi', \xi'']$ be two consecutive ordinary segments of direction e^J and $e^{J'}$, respectively. By Proposition 4.11, we have the inclusion $\mathsf{Ent}(\xi', J') \subset \mathsf{Ent}(\xi, J)$ and the equality $\arg\max(W_i^+ \odot \xi') = \arg\max(W_i^+ \odot \xi)$ for all $i \in \mathsf{Ent}(\xi', J')$. It follows that if Ω is a function determining whether $j \in \arg\max(W_i^+ \odot \xi)$ for all $i \in \mathsf{Ent}(\xi, J)$, it can be used as well to determine whether $j \in \arg\max(W_i^+ \odot \xi')$ for all $i \in \mathsf{Ent}(\xi', J')$.

Then, the correctness of the algorithm straightforwardly follows from Proposition 4.9(i) (for the computation of $\vec{\mathcal{G}}_{]\xi^1,\xi^2[}$ at line 2), Proposition 4.3 (for the computation of J at line 3), and Proposition 4.12.

The complexity of the operations from lines 1 to 7 can easily be verified to be in $O(mn)$. Let $q \leqslant n$ be the number of iterations of the loop from lines 8 and 11, and let $e^{J_1}, e^{J_2}, \dots, e^{J_q}$ be the directions of the ordinary segments followed during the successive calls to Algorithm 2. By Proposition 4.12, the total complexity of the loop is

$$O(nq + m|J_2 \setminus J_1| + m|J_3 \setminus J_2| + \cdots + m|J_q \setminus J_{q-1}|)\,,$$

which can be bounded by $O(n(m+n))$. Finally, the space complexity is obviously bounded by $O(nm)$. $\square$

5. Reduced costs. In this section, we introduce the concept of tropical reduced costs, which are merely the signed valuations of the reduced costs over Puiseux series. Then, pivots improving the objective function and optimality over Puiseux series can be determined only by the signs of the tropical reduced costs. We show that, under some genericity assumptions, the tropical reduced costs can be computed using only the tropical entries A and c in time $O(n(m+n))$. This complexity is similar to the classical simplex algorithm, as this operation corresponds to the update of the inverse of the basic matrix A_I.

5.1. Symmetrized tropical semiring. Until now our coordinate domain was the set of signed tropical numbers $\mathbb{T}_\pm$. As noted in section 2.1.1, this has the drawback of not being a semiring since, in general, $a \oplus (\ominus a)$ is not defined. This can be remedied by extending $\mathbb{T}_\pm$ to the *symmetrized tropical semiring* from [Plu90], which we denote here as $\mathbb{S}$. We shall see in particular that the computation of *tropical reduced costs* reduces to the resolution of the analogue of a Cramer system over the symmetrized tropical semiring.

As a set $\mathbb{S}$ is the union of $\mathbb{T}_\pm$ and a third copy of $\mathbb{T}$, denoted $\mathbb{T}_\bullet$. The members of the latter, written as $a^\bullet$ for $a \in \mathbb{T}$, are the *balanced tropical numbers.* The numbers a, $\ominus a$ and $a^\bullet$ are pairwise distinct unless $a = \mathbb{0}$. Sign and modulus are extended to $\mathbb{S}$ by setting $\operatorname{sign}(a^\bullet) = 0$ and $|a^\bullet| = a$.

The addition of two elements $x, y \in \mathbb{S}$, denoted by $x \oplus y$, is defined to be $\max(|x|, |y|)$ if the maximum is attained only by elements of positive sign, $\ominus \max(|x|, |y|)$ if it is attained only by elements of negative sign, and $\max(|x|, |y|)^\bullet$ otherwise. For instance, $(\ominus 1) \oplus 1 \oplus (\ominus 3) = 1^\bullet \oplus (\ominus 3) = \ominus 3$. The multiplication $x \odot y$ of two elements $x, y \in \mathbb{S}$ yields the element with modulus $|x| + |y|$ and with sign $\operatorname{sign}(x)\operatorname{sign}(y)$. For example, $(\ominus 1) \odot 2 = \ominus 3$ and $(\ominus 1) \odot (\ominus 2) = 3$ but $1^\bullet \odot (\ominus 2) = 3^\bullet$. An element $x \in \mathbb{T}_\pm$ not equal to $\mathbb{0}$ has a multiplicative inverse x^{-1} which is the element of modulus $-|x|$ and with the same sign as x. The addition $A \oplus B$ and multiplication $A \odot B$ of two matrices $A = (a_{ij})$ and $B = (b_{ij})$ are the matrices with entries $a_{ij} \oplus b_{ij}$ and $\bigoplus_k a_{ik} \odot b_{kj}$, respectively.

The set $\mathbb{S}$ also comes with the *reflection map* $x \mapsto \ominus x$ which sends a balanced number to itself, a positive number a to $\ominus a$, and a negative number $\ominus a$ to a. We will write $x \ominus y$ for $x \oplus (\ominus y)$. Two numbers $x, y \in \mathbb{S}$ satisfy the *balance relation* $x \nabla y$ when $x \ominus y$ is a balanced number. Note that

$$x \nabla y \implies x = y \qquad \text{for all } x, y \in \mathbb{T}_\pm \ . \tag{5.1}$$

The balance relation is extended entrywise to vectors in $\mathbb{S}^n$. In the semiring $\mathbb{S}$, the relation ∇ plays the role of the equality relation; in particular the next result shows that a version of Cramer's theorem is valid in the tropical setting, up to replacing equalities by balances.

The tropical determinant of the square matrix $M = (m_{ij}) \in \mathbb{S}^{n \times n}$ is given by

$$\operatorname{tdet}(M) = \bigoplus_{\sigma \in \operatorname{Sym}(n)} \operatorname{tsign}(\sigma) \odot m_{1\sigma(1)} \odot \cdots \odot m_{n\sigma(n)}. \tag{5.2}$$

Also observe that a square matrix of $\mathbb{T}_\pm^{n \times n}$ is tropically sign singular if and only if its tropical determinant is a balanced number.

THEOREM 5.1 (signed tropical Cramer theorem [Plu90]). *Let $M \in \mathbb{S}^{n\times n}$ and $d \in \mathbb{S}^n$. Every solution $y \in \mathbb{T}_\pm^n$ of the system of balances*

$$M \odot y \nabla d \tag{5.3}$$

satisfies

$$\operatorname{tdet}(M) \odot y_j \nabla \operatorname{tdet}(M_{j\leftarrow d}) \qquad \textit{for all } j \in [n], \tag{5.4}$$

where $M_{j\leftarrow d}$ is the matrix obtained by replacing the jth column of M by d.

Conversely, if the tropical determinants $\operatorname{tdet}(M)$ *and* $\operatorname{tdet}(M_{j\leftarrow d})$ *for $j \in [n]$ are not balanced elements, then the vector with entries* $y_j = \operatorname{tdet}(M)^{-1} \odot \operatorname{tdet}(M_{j\leftarrow d})$ *is the unique solution of* (5.3) *in $\mathbb{T}_\pm^n$.*

This result was proved in [Plu90]; see also [AGG09, AGG14] for more recent discussions. A different tropical Cramer theorem (without signs) was proved by Richter-Gebert, Sturmfels, and Theobald [RGST05]; their proof relies on the notion of a coherent matching field introduced by Sturmfels and Zelevinsky [SZ93].

Remark 5.2. The quintuple $(\mathbb{T}_\pm, \max, +, \ominus 0, \mathbb{T}_\bullet)$ is an example of a "fuzzy ring" in the sense of [Dre86, Definition 1.1]. In the notation of that reference, $\mathbb{T}_\pm$ is "the group of units" and $\mathbb{T}_\bullet$ is the set denoted "K_0."

5.2. Computing solutions of tropical Cramer systems. The Jacobi iterative algorithm of [Plu90] allows one to compute a signed solution y of the system $M \odot y \nabla d$; see also [AGG14] for more information. We next present a combinatorial version of this algorithm for the special case where the entries of M and d are in $\mathbb{T}_\pm$.

Suppose that $\operatorname{tdet}(M) \neq \mathbb{0}$, and let σ be a maximizing permutation in $\operatorname{tdet}(M)$ (or equivalently, in $\operatorname{tper}(|M|)$). The *Cramer digraph* of the system associated with σ is the weighted bipartite directed graph over the "column nodes" $\{1, \dots, n+1\}$ (the index $n+1$ represents the affine component) and "row nodes" $\{1, \dots, n\}$ defined as follows: every row node $i \in [n]$ has an outgoing arc to the column node $\sigma(i)$ with weight $m_{i\sigma(i)}^{-1}$, and an incoming arc from every column node $j \neq \sigma(i)$ with weight $\ominus m_{ij}$ when $j \in [n]$, and weight d_i when $j = n+1$.

Example 5.3. The maximizing permutation for the system of balances (5.5) below is $\sigma(1) = 1, \sigma(2) = 3$ and $\sigma(3) = 2$. The Cramer digraph is represented in Figure 9.

$$\begin{pmatrix} \ominus(-1) & -\infty & -\infty \\ -1 & \ominus(-2) & 0 \\ \ominus(-1) & 0 & -\infty \end{pmatrix} \odot \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} \nabla \begin{pmatrix} -2 \\ 0 \\ -1 \end{pmatrix}. \tag{5.5}$$

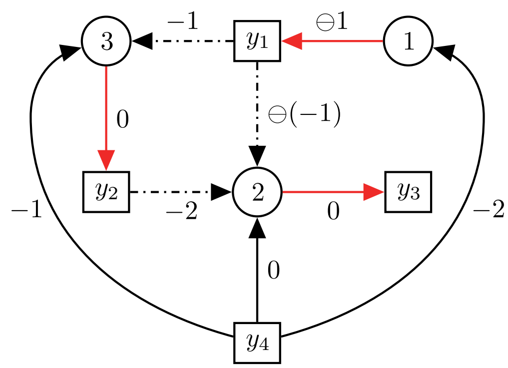

FIG. 9. *The Cramer digraph for the system of balances in* (5.5)*. Column nodes are squares and row nodes are circles. Arcs with weight $-\infty$ are omitted. The maximizing permutation σ is given by the red arcs. The coordinate y_j of the signed solution y of* (5.5) *is obtained by the multiplication (in $\mathbb{S}$) of the weight on the longest path from y_4 to y_j.*

Note that all the coefficients $m_{i\sigma(i)}$ are different from $\mathbb{0}$. In what follows, it will be convenient to consider the longest path problem in the weighted digraph obtained from the Cramer digraph associated with σ by forgetting the tropical signs, i.e., by taking the modulus of each weight. Note in particular that there is no directed cycle the weight of which has a positive modulus (otherwise σ would not be a maximizing permutation in the tropical determinant of M). Consequently, the latter longest path problem is well-defined (longest weights being either finite or $-\infty$, but not $+\infty$).

The *digraph of longest paths* from a node v refers to the subgraph of the Cramer digraph formed by the arcs belonging to a longest path from node v. This digraph is acyclic and each of its nodes is reachable from the node v (possibly with a path of length $\mathbb{0}$). As a result, it always contains a directed tree rooted at v. Such a directed tree can be described by a map which sends every node (except the root) to its parent node. Note that by construction of the Cramer digraph, a column node j has only one possible parent node $\sigma^{-1}(j)$. Consequently, we will describe a directed tree of longest paths by a map γ that sends every row node to its parent column node.

PROPOSITION 5.4. *Let $M \in \mathbb{T}_\pm^{n\times n}$ such that* $\operatorname{tdet}(M) \neq \mathbb{0}$ *and $d \in \mathbb{T}_\pm^n$. Let σ be a maximizing permutation in the tropical determinant of M. In the Cramer digraph of the system $M \odot y \,\nabla\, d$ associated with σ, consider the digraph of longest paths from the column node $n+1$. In this digraph of longest paths, choose any directed subtree γ rooted at the column node $n+1$. Then, the recursive relations*

$$y_{\sigma(i)} = \begin{cases} d_i \odot m_{i\sigma(i)}^{-1} & \textit{when } \gamma(i) = n+1 \ , \\ \ominus m_{i\gamma(i)} \odot m_{i\sigma(i)}^{-1} \odot y_{\gamma(i)} & \textit{otherwise} \end{cases} \tag{5.6}$$

provide a solution in $\mathbb{T}_\pm^n$ of the system $M \odot y \,\nabla\, d$.

Proof. Since the column node $n+1$ reaches all column nodes in the directed tree defined by γ, (5.6) defines a point y in $\mathbb{T}_\pm^n$. The modulus $|y_j|$ is the weight of a longest path from the column node $n+1$ to the column node j. By the optimality conditions of the longest paths problem, for any $i \in [n]$, we have

$$|m_{i\sigma(i)}| + |y_{\sigma(i)}| \geqslant |d_i| \ ,$$
$$|m_{i\sigma(i)}| + |y_{\sigma(i)}| \geqslant |m_{ij}| + |y_j| \quad \text{for all } j \in [n] \ .$$

Furthermore, we have $|m_{i\sigma(i)}| + |y_{\sigma(i)}| = |m_{i\gamma(i)}| + |y_{\gamma(i)}|$ when $\gamma(i) \neq n+1$ and $|m_{i\sigma(i)}| + |y_{\sigma(i)}| = |d_i|$ otherwise.

Thus, if $\gamma(i) \neq n+1$, the terms $m_{i\sigma(i)} \odot y_{\sigma(i)}$ and $m_{i\gamma(i)} \odot y_{\gamma(i)}$ have maximal modulus among the terms of the sum $m_{i1} \odot y_1 \oplus \cdots \oplus m_{in} \odot y_n \ominus d_i$. Moreover, (5.6) ensures that $m_{i\sigma(i)} \odot y_{\sigma(i)} \oplus m_{i\gamma(i)} \odot y_{\gamma(i)}$ is balanced. Similarly, if $\gamma(i) = n+1$, then $m_{i\sigma(i)} \odot y_{\sigma(i)} \ominus d_i$ is balanced and the terms $m_{i\sigma(i)} \odot y_{\sigma(i)}$ and d_i have maximal modulus in $m_{i1} \odot y_1 \oplus \cdots \oplus m_{in} \odot y_n \ominus d_i$. In both cases, we conclude that $M_i \odot y \mathrel{\nabla} d_i$. ∎

A digraph of longest paths for Example 5.3 is shown in Figure 9. From the relations (5.6), we obtain the signed solution $y = (\ominus(-1), -1, 0)$.

5.2.1. Complexity analysis. We now discuss the complexity of the method provided by Proposition 5.4. First, a maximizing permutation σ can be found in time $O(n^3)$ by the Hungarian method; see [Sch03, section 17.2]. Second, the digraph of longest paths, as well as a directed tree of longest paths, can be determined in time $O(n^3)$ using the Bellman–Ford algorithm; see [Sch03, section 8.3]. Last, the solution x can be computed in time $O(n)$.

However, we claim that the complexity of the second step can be decreased to $O(n^2)$. The idea is to consider a variant of the Cramer digraph with nonpositive weights, and then to apply Dijkstra's algorithm to solve the longest paths problem. We exploit the fact that the Hungarian method is a primal-dual algorithm, which returns, along with a maximizing permutation σ, an optimal solution (u, v) to the dual assignment problem:

$$
\begin{aligned}
&\min_{u,v\in\mathbb{R}^n} \quad \sum_{i=1}^{n} u_i + \sum_{j=1}^{n} v_j \\
&\qquad\qquad |m_{ij}| - u_i - v_j \leqslant 0 \quad \text{for all } i,j \in [n] \ .
\end{aligned}
\tag{5.7}
$$

By complementary slackness, we have

$$
|m_{i\sigma(i)}| = u_i + v_{\sigma(i)} \quad \text{for all } i \in [n] \ . \tag{5.8}
$$

Since $\operatorname{tdet}(M) \neq \mathbb{0}$, the assignment problem has a solution with a finite cost. Therefore, the dual problem (5.7) is feasible and bounded. Thus it admits a solution $u, v \in \mathbb{R}^n$.

We make the diagonal change of variables $y_j = v_j \odot z_j$, for all $j \in [n]$, where the z_j are the new variables. We consider the matrix $M' = (m'_{ij})$ obtained from M by the diagonal scaling $m'_{ij} = \mu^{-1} \odot u_i^{-1} \odot m_{ij} \odot v_j^{-1}$, where μ is a real number to be fixed soon, together with the vector d' with entries $d'_i = \mu^{-1} \odot u_i^{-1} \odot d_i$ for all $i \in [n]$. Then, dividing (tropically) every row i of the system $M \odot y \mathrel{\nabla} d$ by μ and by u_i, and performing the above change of variables, we arrive at the equivalent system $M' \odot z \mathrel{\nabla} d'$. By choosing $\mu := \max(\max_i(|d_i| - u_i), 0)$, we get that $|d'_i| \leqslant 0$, and $|m'_{ij}| \leqslant 0$ for all $i, j \in [n]$. The longest path problem to be solved in order to apply the construction of Proposition 5.4 to $M' \odot z \mathrel{\nabla} d'$ now involves a digraph with nonpositive weights.

It follows that the latter problem can be solved by applying Dijkstra's algorithm to the digraph with modified costs. Moreover, the directed tree provided by Dijkstra's algorithm is also valid in the original problem.

Algorithm 4. Computing tropical reduced costs.

Input: A basic point x^I of $\mathcal{P}(A, b)$, the associated set I, the objective function c
Output: The tropical reduced costs y^I

$\mathcal{G}_{x^I} \leftarrow$ tangent graph at x^I	$O(mn)$
$\sigma \leftarrow$ maximizing permutation in $\mathrm{tdet}(A_I)$ obtained by a traversal of $\mathcal{G}_{x^I}$	$O(n)$
$u \leftarrow -x^I$	$O(n)$
$v \leftarrow A_I^+ \odot x^I$	$O(mn)$
$\mu \leftarrow \max(\max_{j\in[n]}(c_j - u_j), 0)$	$O(n)$
$M' \leftarrow$ tropically signed matrix with entries $m'_{ij} = \mu^{-1} \odot u_i^{-1} \odot a_{ji} \odot v_j^{-1}$	$O(n^2)$
$d' \leftarrow$ tropically signed vector with entries $d_i = \mu^{-1} \odot u_i^{-1} \odot c_i$	$O(n)$
$\vec{C} \leftarrow$ Cramer digraph of the system $M' \odot y \nabla d'$ for the permutation σ	$O(n^2)$
apply Dijkstra's algorithm to $\vec{C}$ from column node $n+1$	$O(n^2 + n\log(n))$
$\gamma \leftarrow$ the tree of longest paths returned by Dijkstra's algorithm	
$z \leftarrow$ signed vector obtained by applying (5.6) to the tree γ	$O(n)$
return y^I the signed vector with entries $y^I_j = v_j \odot z_j$	$O(n)$

5.3. Tropical reduced costs as a solution of a tropical Cramer system. In the rest of this section, we suppose that Assumption 5 holds, so we only consider basic points x^I with finite entries. We also make the following assumption, which is a tropical version of dual nondegeneracy.

Assumption 6. The matrix $(A^T \ c^T)$ is tropically sign generic.

We can now define the vector of *tropical reduced costs* of a set $I \subset [m]$ of cardinality n such that $\mathrm{tdet}(A_I) \neq \mathbb{0}$ to be the unique solution $y^I \in \mathbb{T}_\pm^m$ of the system of m balances

$$\begin{cases} A^\top \odot y \ \nabla \ c^\top, \\ \qquad y_i \ \nabla \ \mathbb{0} \quad \text{for all } i \in [m] \setminus I. \end{cases} \tag{5.9}$$

Proposition 5.5. *Let x^I be a tropical basic point of $\mathcal{P}(A, b)$ for a suitable $I \subset [m]$. Then there is a unique solution $y^I \in \mathbb{T}_\pm^m$ of the system of balances* (5.9).

Let $(\boldsymbol{A}\ \boldsymbol{b})$ be any lift of $(A\ b)$. Pivoting from the basic point $\boldsymbol{x}^I$ of the Puiseux polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ along the edge $\boldsymbol{\mathcal{E}}_{I\setminus\{k\}}$ (for $k \in I$) improves the objective function if and only if the tropical reduced cost y^I_k is tropically negative. The basic point $\boldsymbol{x}^I$ is an optimum of the Puiseux linear program if and only if the tropical reduced costs y^I are tropically nonnegative.

Proof. First, the signed valuation of the Puiseux reduced costs $\boldsymbol{y}^I$ yields a signed solution of (5.9). Let us show that this solution is unique. We apply Theorem 5.1 with $M = A_I^\top$ and $d = c^\top$. Since I yields a basic point, the matrix $\boldsymbol{A}_I$ is not singular, thus $\mathrm{tdet}(M) \neq \mathbb{0}$. By Assumption 6, the tropical determinants of the matrices M and $M_{j\leftarrow d}$ for $j \in [n]$ belong to $\mathbb{T}_\pm$. Then by (5.4), the vector y^I with entries $y^I_j = \mathrm{tdet}(M)^{-1} \odot \mathrm{tdet}(M_{j\leftarrow d})$ is the unique solution of (5.9).

We have shown that the tropical signs of the tropical reduced costs are exactly the signs of the Puiseux reduced costs, which proves the second part of the proposition. ∎

Example 5.6. In Example 2.8, the tropical reduced costs associated with $I = \{\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3\}$ are given by $y^I = (\ominus(-1), -1, 0)$, which is the signed solution of (5.5). It follows that the only edge with negative reduced cost is $\mathcal{E}_{\{\mathcal{H}_2,\mathcal{H}_3\}}$.

Theorem 5.7. *Algorithm* 4 *computes the tropical reduced costs. Its time complexity is bounded by $O(n(m + n))$.*

Proof. The maximizing permutation σ is computed from $\mathcal{G}_{x^I}$ in line 2 as follows. We first determine a matching between the coordinate nodes $1,\dots,n$ and the set I of hyperplane nodes using the technique described in the proof of Proposition 4.3, case (i). By Lemma 3.2, this matching provides a maximizing permutation in $\operatorname{tdet}(|A_I|)$. It can be obviously computed by a traversal of $\mathcal{G}_{x^I}$ starting from coordinate node $n+1$. Since $\mathcal{G}_{x^I}$ contains $2n+1$ nodes and $2n$ edges (see the proof of Proposition 4.3), this traversal requires $O(n)$ operations. The complexity of the other operations of this algorithm is straightforward and given in annotations. We conclude that the overall time complexity is $O(m(n+n))$.

Let $v = A_I^+ \odot x^I$. For any hyperplane node $j \in I$ and any $i \in [n]$, we have $v_j \geqslant |a_{ji}| + x^I_i$, where $A = (a_{ij})$. Moreover, equality holds for every edge (j,i) in the tangent graph. In particular with the permutation σ, we have $v_{\sigma(i)} = |a_{\sigma(i)i}| + x^I_i$. By Assumptions 1 and 5, we have $v \in \mathbb{R}^n$ and $x^I \in \mathbb{R}^n$. Thus $u = -x^I$ and v form an optimal solution to the dual assignment problem (5.7) for the matrix $M = A_I^\top$. It follows from the discussion in section 5.2 that the operations between lines 3 and 12 compute the tropical reduced costs. $\square$

We conclude this section by applying Algorithm 1 to the tropical linear program of Example 2.8.

Example 5.8. We start from the tropical basic point $(4,4,2)$ associated with $I = \{\mathcal{H}_1,\mathcal{H}_2,\mathcal{H}_5\}$. For this set, tropical reduced costs are $y_{\mathcal{H}_1} = \ominus(-1)$, $y_{\mathcal{H}_2} = -1$ and $y_{\mathcal{H}_5} = \ominus 4$. We choose $i_{\mathsf{out}} = \mathcal{H}_5$ and pivot along the tropical edge $\mathcal{E}_{\{\mathcal{H}_1,\mathcal{H}_2\}}$.

We arrive at the basic point $(1,0,0)$, associated with $I = \{\mathcal{H}_1,\mathcal{H}_2,\mathcal{H}_3\}$. The reduced costs are $y_{\mathcal{H}_1} = \ominus(-1)$, $y_{\mathcal{H}_2} = -1$ and $y_{\mathcal{H}_3} = 0$. The only tropically negative reduced cost is $y_{\mathcal{H}_1}$, thus we pivot along $\mathcal{E}_{\{\mathcal{H}_2,\mathcal{H}_3\}}$.

The new basic point is $(0,0,0)$, corresponding to the set $\{\mathcal{H}_2,\mathcal{H}_3,\mathcal{H}_4\}$. The reduced costs are tropically positive: $y_{\mathcal{H}_2} = -1$, $y_{\mathcal{H}_3} = 0$, and $y_{\mathcal{H}_4} = -2$. Thus $(0,0,0)$ is optimal.

6. Proof of the main theorem and generalization to Hahn series. We now have all the tools needed to prove Theorem 1.1 under the assumptions of primal nondegeneracy (Assumption 4), finiteness (Assumption 5), and dual nondegeneracy (Assumption 6). If a tropical linear program satisfies all three conditions we call it *standard*.

Proof of Theorem 1.1. The time complexity of one iteration of the tropical simplex algorithm follows from the complexity of the tropical pivoting operation (Theorem 4.13) and of the computation of tropical reduced costs (Theorem 5.7).

Proposition-Definition 3.8 and 3.10 ensure that the tropical pivoting operation traces the image by the valuation map of the pivoting operation over Puiseux series. By Proposition 5.5, choosing the pivot according to the signs of the tropical reduced costs amounts to choosing a pivot according the signs of the Puiseux reduced costs.

We claim that under our assumptions, the edges of the Puiseux polyhedron have a positive length (i.e., as a set, they are not reduced to a point). By contradiction, suppose that an edge $\boldsymbol{\mathcal{E}}_K$ between the basic points $\boldsymbol{x}^{K\cup\{k\}}$ and $\boldsymbol{x}^{K\cup\{k'\}}$ has zero length, where $k \neq k'$ and $k,k' \notin K$. Then $\boldsymbol{x}^{K\cup\{k\}} = \boldsymbol{x}^{K\cup\{k'\}}$. Thus the tropical basic point $x = \operatorname{val}(\boldsymbol{x}^{K\cup\{k\}}) = \operatorname{val}(\boldsymbol{x}^{K\cup\{k'\}})$ is contained in the $n+1$ tropical s-hyperplanes $\mathcal{H}(A_i,b_i)$ for $i \in K\cup\{k,k'\}$. Since x has finite entries by Assumption 5, the $n+1$ elements of $K\cup\{k,k'\}$ appear as hyperplane nodes in the tangent graph at x. This contradicts Proposition 4.3 and proves the claim.

The basic points $\boldsymbol{x}^{K\cup\{k\}}$ and $\boldsymbol{x}^{K\cup\{k'\}}$ of an edge $\boldsymbol{\mathcal{E}}_K$ are related by $\boldsymbol{x}^{K\cup\{k'\}} = \boldsymbol{x}^{K\cup\{k\}} + \boldsymbol{\mu}\boldsymbol{d}^k$, where $\boldsymbol{\mu} > 0$ is the length of $\boldsymbol{\mathcal{E}}_K$ and $\boldsymbol{d}^k$ its direction defined in (2.10).

When pivoting from $\boldsymbol{x}^{K\cup\{k\}}$ to $\boldsymbol{x}^{K\cup\{k'\}}$, the objective value increases by $\boldsymbol{\mu}(\boldsymbol{c}\boldsymbol{d}^k)$. Furthermore, $\boldsymbol{y}_k = \boldsymbol{c}\boldsymbol{d}^k$ is the reduced cost of the pivot along $\boldsymbol{\mathcal{E}}_K$ from the basic point $\boldsymbol{x}^{K\cup\{k\}}$. As a consequence, as long as a pivot with a negative reduced cost is chosen, each iteration improves the objective function over Puiseux series. By Assumption 5, the Puiseux polyhedron is bounded, and thus the value of the Puiseux linear program is finite. Therefore, Algorithm 1 does terminate.

Finally, the output of Algorithm 1 is a tropical basic point with tropically nonnegative reduced costs. By Proposition 5.5, the corresponding Puiseux basic point is an optimum of the Puiseux linear program. Then by Proposition 2.6, the tropical basic point is an optimum of the tropical linear program. ∎

We described tropical linear programming in the max-plus version of the tropical semiring. However, the proofs of our results also hold in any semiring $(\mathbb{T}_G, \max, +)$ which arises from an abelian totally ordered group $(G, +, \geqslant)$, i.e., the semiring is defined on the set $\mathbb{T}_G = G \cup \{\mathbb{0}\}$, the order on G is extended to $\mathbb{T}_G$ by setting $\mathbb{0} \leqslant x$ for any $x \in G$, and the maximum is defined with respect to the order on G. In this setting, the notion of "tropical general position" still makes sense. Puiseux series are then replaced by the ordered field $\mathbb{R}[\![t^G]\!]$ of (formal) Hahn series with real coefficients and with value group $(G, +)$; recall that Hahn series are required to have a well-ordered support. The analysis of section 3 relies only on the fact that the coefficients of the series are real numbers (Theorem 3.5). In section 4 the description of a tropical edge as the concatenation of ordinary segments still holds. Finally, in section 5, the tropical Cramer theorem (Theorem 5.1) is still valid in this generalized setting.

THEOREM 6.1. *The assertions of Theorem* 1.1 *remain valid if the tropical semiring is replaced by* $\mathbb{T}_G$ *and if the field of real Puiseux series is replaced by the field of real Hahn series* $\mathbb{R}[\![t^G]\!]$*, the execution time being now evaluated in a model in which every arithmetic operation in the group* G *takes a time* $O(1)$.

We end this paper by mentioning some simple extensions of the present results. Our version of the tropical simplex algorithm can readily be adapted to the maximization of a tropical linear form over a tropical polyhedron, instead of the minimization. Indeed, the former problem can be handled by taking as a cost vector a vector of negative tropical numbers and lifting it to a cost vector of real Puiseux series with negative leading coefficients.

More generally, one may consider a tropical cost vector with both negative and positive coordinates. The present tropical simplex algorithm can still be defined in this setting; however, its interpretation in terms of tropical optimization problem turns out to be less satisfactory. Indeed, the cost function $c \odot x$ may be a balanced tropical number for some feasible vectors x, whereas there is no total order on the symmetrized tropical semiring with a natural interpretation in terms of lift to real Puiseux series. Hence, the tropical minimization problem appears to be somehow ill defined. However, for an input in general position, the cost function evaluated at any tropical basic point will always be unbalanced. Then, the tropical simplex algorithm, with some straightforward modifications, can be used to return a tropical basic point whose cost is minimal among all tropical basic points but which may be incomparable with respect to some nonbasic feasible points.

We did not study the overall complexity of the tropical simplex algorithm. But we expect, as in the classical case, an exponential behavior on some particular examples (such as the Klee–Minty cubes [KM72]).

The results of this paper should allow the construction of more general tropical pivoting algorithms. In particular, the criss-cross method [FT97], which pivots between unfeasible basic points of the arrangement of s-hyperplanes, should also tropicalize. Pivots can be handled with Algorithm 3. The selection of pivots would involve the tropical signs of the basic point and of the reduced costs.

Finally, we briefly comment on the complexity of deciding if the tropical linear program (2.5) given by $A \in \mathbb{T}_{\pm}^{m\times n}$, $b \in \mathbb{T}_{\pm}^{m}$, and $c \in \mathbb{T}^{1\times n}$ satisfies our standard conditions. It is always safe to assume that Assumptions 1 and 2 are satisfied, for it takes at most $O(mn)$ time to simplify the input if this is not the case [GK11, Lemma 1]. Verifying the finiteness condition in Assumption 5 requires solving n tropical linear feasibility problems to check for a nontrivial intersection with the boundary of the tropical projective space, which amounts to solving mean-payoff games [AGG12]. So it is unclear whether this can be done in polynomial time. Checking for nondegeneracity takes exponential time in the classical case [Eri96], and hence this should also hold for the tropical analogue, Assumption 4.

REFERENCES

[AGG09] M. Akian, S. Gaubert, and A. Guterman, *Linear independence over tropical semirings and beyond*, in Proceedings of the International Conference on Tropical and Idempotent Mathematics, G. L. Litvinov and S. N. Sergeev, eds., Contemp. Math. 495, AMS, Providence, RI, 2009, pp. 1–38.

[AGG12] M. Akian, S. Gaubert, and A. Guterman, *Tropical polyhedra are equivalent to mean payoff games*, Internat. Algebra Comput., 22 (2012), 125001.

[AGG13] X. Allamigeon, S. Gaubert, and E. Goubault, *Computing the vertices of tropical polyhedra using directed hypergraphs*, Discrete Comput. Geom., 49 (2013), pp. 247–279.

[AGG14] M. Akian, S. Gaubert, and A. Guterman, *Tropical Cramer determinants revisited*, in Proceedings of the International Conference on Tropical and Idempotent Mathematics, G. L. Litvinov and S.N. Sergeev, eds., Contemp. Math. 616, AMS, Providence, RI, 2014, pp. 1–45.

[AGK11] X. Allamigeon, S. Gaubert, and R. D. Katz, *The number of extreme points of tropical polyhedra*, J. Combin. Theory Ser. A, 118 (2011), pp. 162–189.

[AGNS11] M. Akian, S. Gaubert, V. Nitica, and I. Singer, *Best approximation in max-plus semimodules*, Linear Algebra Appl., 435 (2011), pp. 3261–3296.

[AK13] X. Allamigeon and R. D. Katz, *Minimal external representations of tropical polyhedra*, J. Combin. Theory Ser. A, 120 (2013), pp. 907–940.

[Ale13] D. Alessandrini, *Logarithmic limit sets of real semi-algebraic sets*, Adv. Geom., 13 (2013), pp. 155–190.

[BA08] P. Butkovič and A. Aminu, *Introduction to max-linear programming*, IMA J. Management Math., 20 (2008), pp. 233–249.

[BH04] W. Briec and C. Horvath, *$\mathbb{B}$-convexity*, Optimization, 53 (2004), pp. 103–127.

[BPCR06] S. Basu, R. Pollack, and M. F. Coste-Roy, *Algorithms in Real Algebraic Geometry*, Algorithm Comput. Math. 10, Springer, New York, 2006.

[BV07] H. Bjorklund and S. Vorobyov, *A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games*, Discrete Appl. Math., 155 (2007), pp. 210–229.

[CG79] R. A. Cuninghame-Green, *Minimax Algebra*, Lecture Notes in Econom. and Math. Systems 166, Springer-Verlag, Berlin, 1979.

[CGB03] R. A. Cuninghame-Green and P. Butkovič, *The equation $a\otimes x = b\otimes y$ over (max,+)*, Theoret. Comput. Sci., 293 (2003), pp. 3–12.

[CGQ04] G. Cohen, S. Gaubert, and J. P. Quadrat, *Duality and separation theorem in idempotent semimodules*, Linear Algebra Appl., 379 (2004), pp. 395–422.

[Cha09] J. Chaloupka, *Parallel algorithms for mean-payoff games: An experimental evaluation*, In Algorithms—ESA 2009, Lecture Notes in Comput. Sci. 5757, Springer, Berlin, 2009, pp. 599–610.

[DG06] V. Dhingra and S. Gaubert, *How to solve large scale deterministic games with mean payoff by policy iteration*, in Proceedings of the *1st* international conference on Performance evaluation methodolgies and tools (VALUETOOLS), Vol. 180, Pisa, Italy, 2006, article 12.

[Dre86] A. Dress, *Duality theory for finite and infinite matroids with coefficients*, Adv. Math., 59 (1986), pp. 97–123.

[DS04] M. Develin and B. Sturmfels, *Tropical convexity*, Doc. Math., (2004).

[DSS05] M. Develin, F. Santos, and B. Sturmfels, *On the rank of a tropical matrix*, in Combinatorial and Computational Geometry, Math. Sci. Res. Inst. Publ. 52, Cambridge University Press, Cambridge, UK, 2005, pp. 213–242.

[DY07] M. Develin and J. Yu, *Tropical polytopes and cellular resolutions*, Experiment. Math., 16 (2007), pp. 277–291.

[Eri96] J. Erickson, *New lower bounds for convex hull problems in odd dimensions*, in Proceedings of the 12th Annual ACM Symposium on Computational Geometry, 1996.

[FAA02] J. A. Filar, E. Altman, and K. E. Avrachenkov, *An asymptotic simplex method for singularly perturbed linear programs*, Oper. Res. Lett., 30 (2002), pp. 295–307.

[FT97] K. Fukuda and T. Terlaky, *Criss-cross methods: A fresh view on pivot algorithms*, Math. Program., 79 (1997), pp. 369–395.

[GG98] S. Gaubert and J. Gunawardena, *The duality theorem for min-max functions*, C.R. Acad. Sci., 326 (1998), pp. 43–48.

[GK11] S. Gaubert and R. D. Katz, *Minimal half-spaces and external representation of tropical polyhedra*, J. Algebraic Combin., 33 (2011), pp. 325–348.

[GKK88] V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan, *Cyclic games and an algorithm to find minimax cycle means in directed graphs*, USSR Computat. Math. Math. Phys., 28 (1988), pp. 85–91.

[GKS12] S. Gaubert, R. D. Katz, and S. Sergeev, *Tropical linear-fractional programming and parametric mean payoff games*, J. Symbolic Comput., 47 (2012), pp. 1447–1478, December.

[GM84] M. Gondran and M. Minoux, *Linear algebra in dioids: A survey of recent results*, Ann. Discrete Math., 19 (1984), pp. 147–164.

[GS07] S. Gaubert and S. Sergeev, *Cyclic projectors and separation theorems in idempotent convex geometry*, J. Math. Sci. 155 (2008), pp. 815–829.

[Jer73] R. G. Jeroslow, *Asymptotic linear programming*, Oper. Res., 21 (1973), pp. 1128–1141.

[Jos05] M. Joswig, *Tropical halfspaces*, in Combinatorial and Computational Geometry, Math. Sci. Res. Inst. Publ. 52, Cambridge University Press, Cambridge, UK, 2005, pp. 409–431.

[KM72] V. Klee and G. J. Minty, *How good is the simplex algorithm?* in Inequalities III, Academic Press, New York, 1972, pp. 159–175.

[LMS01] G. L. Litvinov, V. P. Maslov, and G. B. Shpiz, *Idempotent functional analysis: An algebraic approach*, Math. Notes, 69 (2001), pp. 696–729.

[Mar10] T. Markwig, *A field of generalised Puiseux series for tropical geometry*, Rend. Semin. Mat. Univ. Politec. Torino, 68 (2010), pp. 79–92.

[Meg87] N. Megiddo, *On the complexity of linear programming*, in Advances in Economic Theory: Fifth World Congress, T. Bewley, ed., Cambridge University Press, Cambridge, UK, 1987, pp. 225–268.

[Meg89] N. Megiddo, *On the complexity of linear programming*, in Advances in Economic Theory (Cambridge, MA, 1985), Econom. Soc. Monogr. 12, Cambridge University Press, Cambridge, UK, 1989, pp. 225–268.

[Plu90] M. Plus, *Linear systems in* (max,+) *algebra*, in Proceedings of the 29th IEEE Conference on Decision and Control, IEEE, 1990, pp. 151–156.

[RGST05] J. Richter-Gebert, B. Sturmfels, and T. Theobald, *First steps in tropical geometry*, in Idempotent Mathematics and Mathematical Physics, Contemp. Math. 377, AMS, Providence, RI, 2005, pp. 289–317.

[Sch03] A. Schrijver, *Combinatorial Optimization. Polyhedra and Efficiency*, *Vol. A*, Algorithms Combin. 24, Springer-Verlag, Berlin, 2003.

[Sch09] S. Schewe, *From parity and payoff games to linear programming*, in Mathematical Foundations of Computer Science 2009, Lecture Notes in Comput. Sci. 5734, Springer, Berlin, 2009, pp. 675–686.

[Sei54] A. Seidenberg, *A new decision method for elementary algebra*, Ann. Math., (1954), pp. 365–374.

[SS04] D. Speyer and B. Sturmfels, *The tropical Grassmannian*, Adv. Geom., 4 (2004), pp. 389–411.

[SZ93] B. Sturmfels and A. Zelevinsky, *Maximal minors and their leading terms*, Adv. Math., 98 (1993), pp. 65–112.

[Tar51] A. Tarski, *A Decision Method for Elementary Algebra and Geometry*, Research report R-109, RAND Corporation, Santa Monica, CA, 1951.

[Vir01] O. Viro, *Dequantization of real algebraic geometry on logarithmic paper*, in European Congress of Mathematics, Vol. I (Barcelona, 2000), Progr. Math. 201, Birkhäuser, Basel, 2001, pp. 135–146.

[Zim76] K. Zimmermann, *Extremální Algebra*, Ekonomický ùstav ČSAV, Praha, 1976.

[Zim77] K. Zimmermann, *A general separation theorem in extremal algebras*, Ekonom.-Mat. Obzor, 13 (1977), pp. 179–201.

[Zim81] U. Zimmermann, *Linear and Combinatorial Optimization in Ordered Algebraic Structures*, North-Holland, Amsterdam, 1981.

[ZP96] U. Zwick and M. Paterson, *The complexity of mean payoff games on graphs*, Theoret. Comput. Sci., 158 (1996), pp. 343–359.