Tropical aspects of linear programming

Pascal Benchimol

▶ **To cite this version:**

Pascal Benchimol. Tropical aspects of linear programming. Optimization and Control [math.OC]. Ecole Polytechnique, 2014. English. NNT : . tel-01198482

HAL Id: tel-01198482

https://polytechnique.hal.science/tel-01198482

Submitted on 13 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Thèse présentée pour obtenir le grade de

DOCTEUR DE L'ÉCOLE POLYTECHNIQUE

Spécialité : Mathématiques appliquées

par

Pascal Benchimol

Tropical aspects of linear programming

soutenue le 2 décembre 2014 devant le jury composé de :

Jérôme Bolte	Université Toulouse 1 Capitole	président du jury
Francisco Santos	Universidad de Cantabria	rapporteur
Thorsten Theobald	Goethe Universität	rapporteur
Stéphane Gaubert	INRIA Saclay – École Polytechnique	directeur de thèse
Xavier Allamigeon	INRIA Saclay – École Polytechnique	co-directeur de thèse
Michael Joswig	Technische Universität Berlin	examinateur
Ilia Itenberg	Université Paris 6	examinateur
Antoine Deza	McMaster University	examinateur

Contents

Abstract

In this thesis, we present new results on the complexity of classical linear programming on the one hand, and of tropical linear programming and mean payoff games on the other hand. Our contributions lie in the study of the interplay between these two problems provided by the dequantization process. This process tranforms classical linear programs into linear programs over tropical semirings, such as the $\mathbb{R} \cup \{-\infty\}$ endowed with max as addition and $+$ as muliplication.

Concerning classical linear programming, our first contribution is a tropicalization of the simplex method. More precisely, we describe an implementation of the simplex method that, under genericity conditions, solves a linear program over an ordered field. Our implementation uses only the restricted information provided by the valuation map, which corresponds to the "orders of magnitude" of the input. Using this approach, we exhibit a class of classical linear programs over the real numbers on which the simplex method, with any pivoting rule, performs a number of iterations which is polynomial in the input size of the problem. In particular, this implies that the corresponding polyhedra have a diameter which is polynomial in the input size.

Our second contribution concerns interior point methods for classical linear programming. We disprove the continuous analog of the Hirsch conjecture proposed by Deza, Terlaky and Zinchenko, by constructing a family of linear programs with $3r+4$ inequalities in dimension $2r+2$ where the central path has a total curvature which is exponential in r. We also point out suprising features of the tropicalization of the central path. For example it has a purely geometric description, while the classical central path depends on the algebraic representation of a linear program. Moreover, the tropical central path may lie on the boundary of the tropicalization of the feasible set, and may even coincide with a path of the tropical simplex method.

Concerning tropical linear programming and mean payoff games, our main result is a new procedure to solve these problems based on the tropicalization of the simplex method. The latter readily applies to tropical linear programs satisfying genericity conditions. In order to solve arbitrary problems, we devise a new perturbation scheme. Our key tool is to use tropical semirings based on additive groups of vectors ordered lexicographically.

Then, we transfer complexity results from classical to tropical linear programming. We show that the existence of a polynomial-time pivoting rule for the classical simplex method, satisfying additional assumptions, would provide a polynomial algorithm for

tropical linear programming and thus for mean payoff games. By transferring the analysis of the shadow-vertex rule of Adler, Karp and Shamir, we also obtain the first algorithm that solves mean payoff games in polynomial time on average, assuming the distribution of the games satisfies a symmetry property.

We establish tropical counterparts of the notions of basic points and edges of a polyhedron. This yields a geometric interpretation of the tropicalization of the simplex method. As in the classical case, the tropical algorithm pivots on the graph of an arrangement of hyperplanes associated to a tropical polyhedron. This interpretation is based on a geometric connection between the cells of an arrangement of classical hyperplanes and their tropicalization. Building up on this geometric interpretation, we present algorithmic refinements of the tropical pivoting operation. We show that pivoting along an edge of a tropical polyhedron defined by m inequalities in dimension n can be done in time $O(n(m + n))$, a complexity similar to the classical pivoting operation. We also show that the computation of reduced costs can be done tropically in time $O(n(m + n))$.

Résumé

Cette thèse présente de nouveaux résultats de complexité concernant d'un côté la programmation linéaire classique, et de l'autre la programmation linéaire tropicale, cette dernière étant reliée aux jeux répétés. Les contributions proviennent de l'étude du processus de déquantisation qui relie ces deux problèmes. La déquantisation transforme les programmes linéaires classiques en programmes linéaires sur des semi-anneaux tropicaux, comme l'ensemble $\mathbb{R} \cup \{-\infty\}$ muni de max en tant qu'addition, et de + en tant que multiplication.

Concernant la complexité de la programmation linéaire, notre première contribution est la tropicalisation de la méthode du simplexe. Plus précisément, nous décrivons une implémentation de la méthode du simplexe qui, sous des conditions de généricité, résoud un programme linéaire sur un corps ordonné. Cette implémentation utilise seulement l'information partielle donnée par la valuation, ce qui correspond aux "ordres de grandeur" des coefficients du problème. Cette approche permet de construire une classe de programmes linéaires réels sur lesquels la méthode du simplexe termine en un nombre d'itérations qui est polynomial en la taille de l'encodage binaire du problème, et ce indépendamment du choix de la régle de pivotage.

Notre deuxième contribution concerne les méthodes de points intérieurs pour la programmation linéaire classique. Nous réfutons l'analogue continu de la conjecture de Hirsch proposé par Deza, Terlaky et Zinchenko, en construisant une famille de programmes linéaires décrits par $3r + 4$ inégalités sur $2r + 2$ variables pour lesquels le chemin central a une courbure totale qui est exponentielle en r. La tropicalisation du chemin central présente des propriétés inattendues. Par exemple, le chemin central tropical peut être décrit de manière purement géométrique, alors que de manière classique le chemin central dépend de la représentation des contraintes. De plus, le chemin central tropical peut rencontrer la frontière de la tropicalisation de l'ensemble réalisable, et peut même coïncider avec un chemin suivi par la méthode du simplexe tropical.

Concernant la programmation linéaire tropicale et les jeux répétés, notre résultat principal est une nouvelle méthode pour résoudre ces problèmes, basée sur la tropicalisation de la méthode du simplexe. Cette dernière résoud directement les programmes linéaires tropicaux satisfaisant des conditions de généricités. Afin de résoudre les problèmes ne satisfaisant pas ces conditions, une technique de perturbation est utilisée. L'idée principale est d'utiliser des semi-anneaux tropicaux basés sur des groupes de vecteurs ordonnées lexicographiquement.

Nous transférons des résultats de complexité de la programmation linéaire classique vers la programmation linéaire tropicale. Nous montrons que l'existence d'une règle de pivotage polynomiale pour la méthode du simplexe classique fournirai un algorithme polynomial pour la programmation linéaire tropicale, et donc pour les jeux répétés. En transférant l'analyse de Adler, Karp et Shamir de la régle de pivotage dite du "shadow-vertex", nous obtenons le premier algorithme qui résoud les jeux répétés en temps polynomial en moyenne, en supposant que la distribution des jeux satisfait une propriété d'invariance.

Nous établissons une correspondance géométrique entre les cellules d'un arrangement d'hyperplans classiques et leur tropicalisation. Ceci donne une interprétation géométrique à la tropicalisation de la méthode du simplexe. Comme dans le cas classique, l'algorithme tropical pivote sur le graphe d'un arrangement d'hyperplans associé au polyèdre. Ce point de vue géométrique nous permet d'établir des raffinements algorithmiques de l'opération de pivotage tropical. Nous présentons un algorithme qui pivote le long d'une arête d'un polyèdre tropical défini par m inégalités en dimension n en temps $O(n(m+n))$. Nous montrons aussi que le calcul des signes des coûts réduits peut se faire tropicalement en temps $O(n(m+n))$

Remerciements

First and foremost, I am grateful to my advisors, Stéhane Gaubert, Xavier Allamigeon and Michael Joswig, for their wonderful guidance. This work has greatly benefited from Stéphane's numerous ideas, and his exaltation in sharing his impressively broad mathematical knowledge. I am indebted to Xavier for his unwavering support, abundant suggestions, and his availability for my random appearances in his office. I am really thankful for Michael's sharp advices, his dedication to drive this research forward, and his warm hospitality during my visits in Darmstadt and Berlin.

I want to thank Thorsten Theobald and Paco Santos for the great honor they have conferred upon me by reviewing this manuscript. I sincerely appreciate the participation of Ilia Itenberg, Antoine Deza and Jérôme Bolte to my thesis committee.

This work has been co-funded by the "Délégation Générale de l'Armement" (DGA) and a Monge fellowship from the École Polytechnique. I also had the chance to attend several conferences thanks to funds from INRIA and its MaxPlus team. I am grateful to these institutions for their financial support.

This work has unfolded at CMAP, a very enjoyable work place. The administrative team (Nasséra, Alexandra, Jessica Gameiro, Wallis) has always been cheerful and helpful. I really appreciated the nice coffee breaks, meals, beers, and moments I spent with my fellow PhD students. Thank you Laurent, Lætitia, Manon, Gwenael, Michael, Etienne, Antoine, Hélène, Charline, Matthieu, and all the others that I forgot to cite. I also spent a wonderful time with people from the adjacents laboratory. I am thinking of Claire, Cécile and Victor from the LIX, Marine from LMD, Pascale from the CMLS and Pascaline and Jérôme from LPP.

I have many friends to thank for the great moments we have shared during the last three years. Thank you Cyril, Géraldine, Emmanuel, Vincent, Guillaume, Sabrina or Adeline, Pierre-Alain, Sébastien, Bastien, Fabien, for the various evenings and holidays we shared. I always enjoy seeing the friends I made in Montreal, Cécile and Cécile, Sylvain, Mélanie, Julien, Chloé, Hubert, Raphaël, with a special thanks to my best roommate, Ileana.

I am deeply and truly grateful to my parents and my brothers for their unconditional love and support, as well as to my family.

Chapter 1

Introduction

1.1 Context

1.1.1 Linear programming and its complexity

Linear programming is a foundation of mathematical optimization, in both its theoretical and practical aspects. A linear program seeks a minimizer of a linear form satisfying linear constraints (see Figure 1.1, left for an illustration). Several kinds of problems in operations research can be modelled within this framework. The ability to solve linear programs also serves as a building block for more general optimization problems, such as convex programming, integer programming or non-linear programming. From a more theoretical point of view, linear programming is related to the geometry and combinatorics of polyhedra.

One of the main open questions concerns the precise complexity of linear programming. The well-known simplex method, introduced by Dantzig [Dan98], moves on the vertex/edge graph (Figure 1.1, middle) of the feasible set until an optimal solution is reached. At each iteration, the next vertex is chosen by a *pivoting rule*. The number of iterations depends on the choice of the pivoting rule. The method is extremely effective in practice. However, pathological examples show that, for most known pivoting rules, the method can be compelled to visit exponentially many vertices.

The ellipsoid method of Khachiyan [Kha80] was a theoretical breakthrough. It proved that linear programs can be solved in polynomial-time. More precisely, the ellipsoid method solves a linear program within a time bounded by a polynomial in L, where L is the number of bits required to describe the problem. In a nutshell, the method determines the emptyness of a polyhedron using a sequence of ellipsoids, whose volumes shrink exponentially fast. The ellipsoid method extends to arbitrary convex problems, provided that a separation oracle is known [GLS88]. Despite its theoretical appeal, the ellipsoid method is not efficient in practice.

The interior-point methods, initiated by Karmakar [Kar84], combine good practical performances with a polynomial-time worst-case complexity. These methods are driven to an optimal solution by a trajectory, called the central path, that goes through the

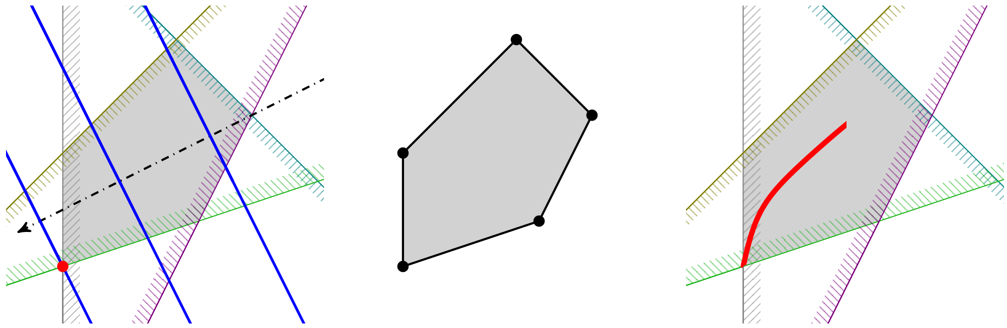

Figure 1.1: Left: a linear program; the feasible set (the gray shaded area) is a polyhedron defined by the halfspaces represented by hashed lines; the objective function is displayed by the dotted arrow; three level sets are depicted in blue; the red dot is the unique optimal solution. Middle: the vertex/edge graph of this polyhedron. Right: the central path for this linear program.

interior of the feasible set (see Figure 1.1, right).

Yet, it is unknown whether linear programs can be solved in *strongly* polynomial time. An algorithm is strongly polynomial if, given a problem described by n rational numbers, it peforms a number of arithmetic operations which is polynomial in n, and the space used by the algorithm is polynomial in the bit length of the input. The existence of a strongly polynomial algorithm for linear programming has been recognized by Smale as one of the mathematical problems of the 21st century [Sma98].

Since the invention of the simplex method, linear programming has been an active field of research. We mention a few significant results, and we refer to [DL11] for an overview of recent advances.

The simplex method and the diameter of polyhedra

Klee and Minty [KM72] showed that the simplex method with the pivoting rule originally proposed by Dantzig visits all vertices of a "tilted" cube, and thus performs a number of iterations which is exponential in the dimension. The same behavior occurs with the "steepest" edge rule [GS79], the "best improvement" rule [Jer73b] or Bland's rule [AC78]. These worst-case examples are subsumed by the deformed products of Amenta and Ziegler [AZ96]. More recently, superpolynomial behavior was also proved for randomized pivoting rule [FHZ11], or "history-based" rules that take into account the previously visited vertices [AF13, Fri11].

On the other hand, for linear programs with special properties, several positive results are known. The simplex method is strongly polynomial for linear programs that arise from network flow problems [Orl97], from Markov Decision Problems (MDP) with a fixed discount rate [Ye11], or from deterministic MDP with any discount rate [HKZ14]. Kitahara and Mizuno showed that, with any pivoting rule that selects improving pivots, the number of iterations is bounded by a polynomial in the value of entries of the vertices of the problem [KM13a, KM13b]. In particular, this proves strong polynomiality for linear programs on polyhedra with 0/1 vertices.

The shadow-vertex pivoting rule, introduced by Gass and Saaty [GS55], was used in several noteworthy results. First, it proved that the simplex method is polynomial *on average*, for certain distributions of instances [Sma83, Bor87, AKS87]. Second, Spielman and Teng [ST04] proved that the simplex method has polynomial *smoothed* complexity when the shadow vertex rule is used. Third, it was used by Kelner and Spielman [KS06] to obtain a *randomized* algorithm with polynomial expected running time. Note however that superpolynomial worst-case examples are known for this rule [Gol94, Mur80].

The complexity of the simplex method is tightly linked to the combinatorics of polyhedra, in particular, to the diameter of their vertex/edge graph. Hirsch conjectured that the diameter of a polytope described by m inequalities in dimension n does not exceed $m-n$. In a recent breakthrough, Santos disproved this conjecture [San12]. Yet, whether the diameter is bounded by a polynomial in m and n remains an open question. Kalai and Kleitman obtained a general bound of $m^{\log(n)+2}$ [KK92], that was improved recently by Todd [Tod14] to $(m-n)^{\log(n)}$. Bonifas *et al.* obtained a bound that depends on the value of the subdeterminants of the input matrix [BDSE$^+$12] (see also [BR14] for a constructive version). The Hirsch conjecture holds in special cases, such as 0/1 polytopes [Nad89] or transportation polytopes [DLKOS09]. For a thorough survey on the diameter of polyhedra, we refer to [KS10].

Interior point methods, and the curvature of the central path

Interior point methods performs a piece-wise linear approximation of the central path to reach an optimal solution. The curvature measures how far a path differs from a straight line. Intuitively, a central path with high curvature should be harder to approximate with line segments, and thus this suggests more iterations of the interior point methods. Dedieu and Shub [DS05] conjectured that the total curvature of a linear program in dimension n is bounded by $O(n)$. This conjecture holds when averaged over all regions of an arrangement of hyperplanes. It was proved by Dedieu, Malajovich and Shub [DMS05] via the multihomogeneous Bézout Theorem and by De Loera, Sturmfels and Vinzant [DLSV10] using matroid theory. However, the redundant Klee-Minty cube of [DTZ09] and the "snake" in [DTZ08] are instances which show that the total curvature can be in $\Omega(m)$ for a linear program described by m inequalities. By analogy with the classical Hirsch conjecture, Deza, Terlaky and Zichencko made the following conjecture.

Conjecture 1.1 (Continuous Hirsch conjecture [DTZ08])**.** *The total curvature of the central path of a linear program defined by m inequalites is bounded by $O(m)$.*

Besides the redundant Klee-Minty cube [DTZ09] and the "snake" [DTZ08], Gilbert, Gonzaga and Karas [GGK04] also exhibited ill-behaved central paths. They showed that the central path can have a "zig-zag" shape with infinitely many turns, on a problem defined in $\mathbb{R}^2$ by non-linear but convex functions. In terms of iteration-complexity of interior-point methods, several worst-case results have been proposed [Ans91, KY91, JY94, Pow93, TY96, BL97]. In particular, Stoer and Zhao [ZS93] showed that the iteration-complexity of a certain class of path-following methods is governed by an integral along the central path. This quantity, called the *Sonnevend's curvature* was introduced in [SSZ91]. The tight relationship between the total Sonnevend's curvature and the iteration-complexity of interior-points methods have been extended to semi-definite and symmetric cone programs [KOT13]. Monteiro and Tsuchiya [MT08] proved that a central path in dimension n consists of $O(n^2)$ "long" parts where the Sonnevend's curvature is small, while the remaining part of the path is relatively small. This was also observed by Vavasis and Ye [VY96] using a notion of *crossover events*.

Note that Sonnevend's curvature is a different notion than the *geometric* curvature we study in this manuscript. To the best of our knowledge, there is no explicit relation between the geometric curvature and the iteration-complexity of interior-point methods. However, these two notions of curvature share similar properties. In particular, the total geometric curvature and the total Sonnevend's curvature are maximal when the number of inequalities is twice the dimension [DTZ08, MT13b]. On the redundant Klee-Minty cube, both the total geometric curvature and the Sonnevend's curvature are large [MT13a, DTZ09].

Sonnevend's curvature relates to another iteration-complexity bound expressed in terms of a condition number associated with the matrix describing a linear program, see [MT08].

We also mention that Megiddo and Shub [MS89], as well as Powell [Pow93], showed that interior point methods may exhibit a simplex-like behavior. For more litterature on interior points methods, we refer to [Wri05, Gon12] and the references therein.

1.1.2 Tropical geometry

Tropical geometry is the (algebraic) geometry on the *max-plus* semiring $(\mathbb{R}_{\max}; \oplus, \odot)$ where the set $\mathbb{R}_{\max} = \mathbb{R} \cup \{-\infty\}$ is endowed with the operations $a \oplus b = \max(a, b)$ and $a \odot b = a + b$. The max-plus, or min-plus semirings, are now dubbed *tropical semirings* in honor of pioneering work of the mathematician and computer scientist Imre Simon. Tropical semirings were studied under various names in relation with optimization [CG79], graph algorithms [GM84], or discrete event systems [CMQV89, BCOQ92, CGQ99, HOvdW06]. Tropical algebra has a strong combinatorial flavor. For example, determinants correspond to optimal assignments, and eigenvalues corresponds to cycles of maximum mean in a graph [But03].

The set $\mathbb{R}_{\max}$ can be seen as the set of "orders of magnitude". If one think of tropical numbers $a, b \in \mathbb{T}$ as exponents of usual numbers, *e.g.*, 10^a and 10^b, then, the tropical operations max and $+$ reflect the usual addition and multiplication on the exponents, *i.e.*, $10^a + 10^b \approx 10^{\max(a,b)}$ and $10^a \cdot 10^b = 10^{a+b}$. More formally, we can identify $a \in \mathbb{R}_{\max}$ with a class $\Theta(t^a)$ of real valued functions, where $f \in \Theta(t^a)$ when f satisfy $ct^a \leq f(t) < c't^a$ for some positive constants $c, c' \in \mathbb{R}$, and for any t large enough. Then, $f \in \Theta(t^a)$ and $g \in \Theta(t^b)$ satisfy $f + g \in \Theta(t^{\max(a,b)})$ and $f \cdot g \in \Theta(t^{a+b})$. Thus, the valuation map

$$\mathrm{val} : f \mapsto \lim_{t \to \infty} \log_t(f(t)) \ ,$$

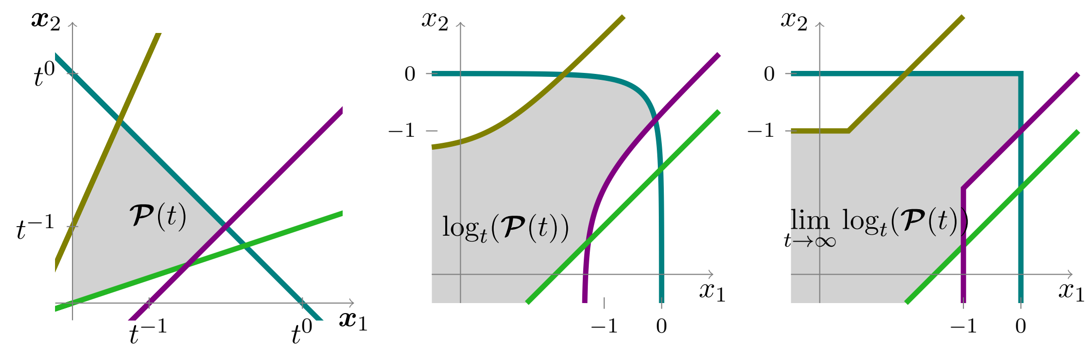

Figure 1.2: Illustration of the dequantization process on the family of polyhedra $(\boldsymbol{\mathcal{P}}(t))_t$ of Example 1.2. Left: a polyhedron $\boldsymbol{\mathcal{P}}(t)$. Middle: its image under the logarithmic map in base t. Right: the dequantization of the family $(\boldsymbol{\mathcal{P}}(t))_t$.

where $\log_t(x) = \log(x)/\log(t)$, is a semiring homomorphism from a certain set of positive real valued functions, equipped with addition and multiplication, to the max-plus semiring. This logarithmic limit process is known as Maslov's *dequantization* [Lit07], or Viro's method [Vir01]. It can be traced back to the work of Bergman [Ber71]. More generally, the tropical semiring can also be thought of as the image of a non-Archimedean field under its valuation map. The non-Archimedean fields typically used are the field of rational functions, formal Puiseux series [EKL06, DY07, RGST05] or the field of generalized Puiseux series with real exponents [Mar10].

Example 1.2. Consider for example the field $\mathbb{R}(t)$ of rational functions in the variable t. We can order $\mathbb{R}(t)$ by setting $f \in \mathbb{R}(t)$ to be positive when $f(t) > 0$ for all t large enough. Now consider the following polyhedron $\boldsymbol{\mathcal{P}}$ over the ordered field $\mathbb{R}(t)$:

$$
\begin{aligned}
\boldsymbol{x}_1 + \boldsymbol{x}_2 &\leq 1 \\
\boldsymbol{x}_1 &\leq t^{-1} + t\boldsymbol{x}_2 \\
\boldsymbol{x}_2 &\leq t^{-1} + t^2\boldsymbol{x}_1 \\
\boldsymbol{x}_1 &\leq t^2\boldsymbol{x}_2 \\
\boldsymbol{x}_1 &\geq 0,\ \boldsymbol{x}_2 \geq 0 \\
\boldsymbol{x}_1 &\in \mathbb{R}(t),\ \boldsymbol{x}_2 \in \mathbb{R}(t)\ .
\end{aligned}
$$

When t is replaced by real numbers, we obtain a family of polyhedra $(\boldsymbol{\mathcal{P}}(t))_t$ in $\mathbb{R}^2$. One of these polyhedra is depicted in Figure 1.2 (left). Applying the map $x \mapsto \log_t(x)$ point-wise to $\boldsymbol{\mathcal{P}}(t)$ provides the set displayed in Figure 1.2 (middle). The dequantization of the family $(\boldsymbol{\mathcal{P}}(t))_t$ is the logarithmic limit illustrated in Figure 1.2 (right).

Through the dequantization process, tropical geometry provides a piece-wise linear "shadow" of classical geometry. The tropicalization of an algebraic variety V, that is, the joint vanishing locus of finitely many polynomials in d indeterminates over a field with a non-archimedean valuation, is a polyhedral complex in $\mathbb{R}^d$, the tropical variety

$T(V)$, which is obtained by applying the valuation map coordinate-wise to all points in V. For instance, if C is a planar algebraic curve over an algebraically closed field, then $T(C)$ is a planar graph. Key features of V are visible in $T(V)$. For example, if V is irreductible, then $T(V)$ is connected [EKL06].

Tropicalization has been succesfully applied in enumerative geometry. The Gromov-Witten invariants count the number of planar complex algebraic curve passing through a generic configuration of points. Mikhalkin showed that these numbers coincide with their tropical counterparts (counted with multiplicities), that are easier to compute [Mik05]. The same technique also applies to real algebraic curves [IKS03]. Given a tropical curve, one can construct real algebraic curves whose topology coincides with the tropical one. This method, known as Viro's patchworking, produced curves with a rich topology. In particular, patchworking was used to disprove the Ragsdale conjecture [IV96]. For a more detailed description of tropical varieties, we refer to [RGST05, EKL06, IMS07, MS]. Computational aspects are developped in [BJS+07], and complexity issues are studied in [The06]. An enlightning introduction is given in [Bru12].

The dequantization of semi-algebraic sets is a more recent subject of research. Speyer and Williams [SW05], studied the tropicalization of the positive part of the Grassmannian. Tabera explored the bases of real tropical varieties [Tab13], and Vinzant investigated their real radical ideals [Vin12]. The tropicalization of polytopes, the most simple class of semi-algebraic sets, was studied by Develin and Yu [DY07] and by Ardila and Develin [AKW06]. Alessandrini [Ale13] devised a framework to study the logarithmic limit of families of semi-algebraic sets.

1.1.3 Tropical linear programming

We are interested in the tropical counterpart of linear programming. A tropical linear program asks for a minimizer $x \in \mathbb{R}^n_{\max}$ of a tropical linear form

$$x \mapsto \max(c_1 + x_1, \dots, c_n + x_n) \ ,$$

for some $c \in \mathbb{R}^n_{\max}$, that satisfy finitely many constraints of the form:

$$\max(\alpha_1 + x_1, \dots, \alpha_n + x_n, \beta) \geq \max(\delta_1 + x_1 \dots, \delta_n + x_n, \gamma) \ ,$$

where $\alpha, \delta \in \mathbb{R}^n_{\max}$ and $\beta, \gamma \in \mathbb{R}_{\max}$. An example is depicted in Figure 1.3. The feasible set of a tropical linear program forms a tropical polyhedra, the most basic example of tropical convex sets.

The tropical analogues of convex sets have appeared in the work of several authors. Motivated by discrete optimization problems, Zimmerman established a separation result [Zim77]. Max-plus analogues of linear spaces were studied by Cuninghame-Green [CG79]. They were also considered by Litvinov, Maslov, and Shpiz under the name of idempotent spaces [LMS01]. Cohen, Gaubert, and Quadrat [CGQ01, CGQ04] also studied them under the name of semimodules, for a geometric approach of discrete event systems [CGQ99], further developed in [Kat07, DLGKL10]. They were also considered by Singer for abstract convex analysis [Sin97]. Tropical convexity is similar to

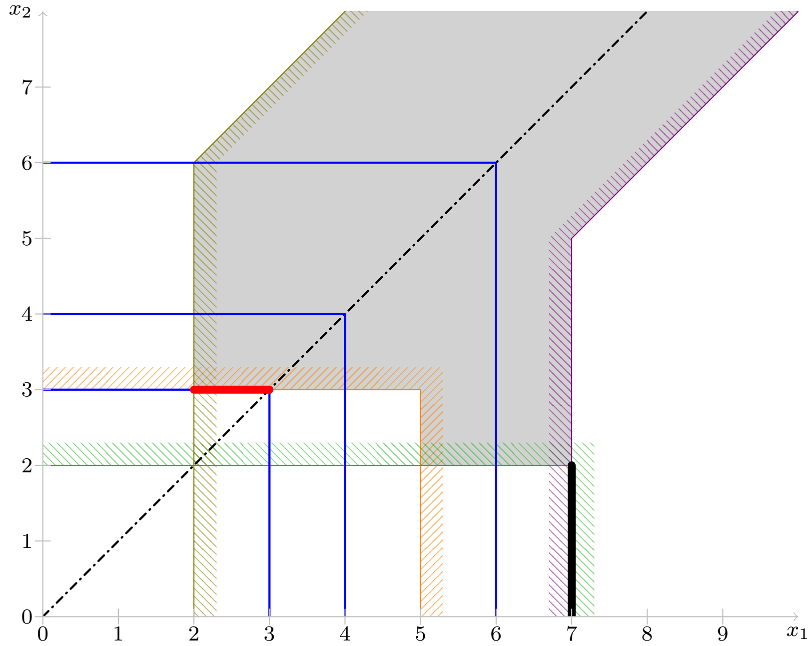

Figure 1.3: A tropical linear program. The feasible set is the union of the gray shaded area with the black halfline. Three level sets for the objective function $(x_1, x_2) \mapsto \max(x_1, x_2)$ are depicted in blue. The red segment is the set of optima.

the $\mathbb{B}$-convexity studied by Briec, Horvath, and Rubinov [BH04, BHR05]. Develin and Sturmfels gave in [DS04] another approach of tropical convexity. They studied tropical polyhedra as polyhedral complexes in the usual sense. Some more recent developments include [Jos05, JSY07, BY06, BSS07, GK09, Jos09, GM10, AGG10].

1.1.4 From tropical linear programming to mean payoff games

Akian, Gaubert and Guterman [AGG12] proved that mean payoff games are equivalent to tropical linear feasibility problems. We briefly recall this equivalence, and we refer the reader to [AGG12] for more information. We shall describe a mean payoff game by a pair of *payment matrices* $A, B \in \mathbb{R}_{\max}^{m \times n}$. We also fix an *initial state* $\bar{\jmath} \in [n]$. The corresponding game, with perfect information, is played by two players, called "Max" and "Min". Their moves alternate as follows. We start from state $j_0 := \bar{\jmath}$. Player Min chooses a state $i_1 \in [m]$ such that $B_{i_1 j_0} \neq -\infty$, and receives a payment of $B_{i_1 j_0}$ units from Player Max. Then, Player Max chooses a state $j_1 \in [n]$ such that $A_{i_1 j_1} \neq -\infty$, and receives a payment of $A_{i_1 j_1}$ from Player Min. Now Player Min again chooses a state $i_2 \in [m]$ such that $B_{i_2 j_1} \neq -\infty$, receives a payment of $B_{i_2 j_1}$ from Player Max, and so on. If $j_0, i_1, j_1, i_2, j_2, \ldots$ is the infinite sequence of states visited in this way, the *mean* *payoff* of Player Max is defined to be

$$\liminf_{p\to\infty} p^{-1}(-B_{i_1j_0} + A_{i_1j_1} - B_{i_2j_1} + A_{i_2j_2} + \cdots - B_{i_pj_{p-1}} + A_{i_pj_p}) \ . \tag{1.1}$$

Conversely, the *mean payoff* for Player Min is

$$\limsup_{p\to\infty} p^{-1}(-B_{i_1j_0} + A_{i_1j_1} - B_{i_2j_1} + A_{i_2j_2} + \cdots - B_{i_pj_{p-1}} + A_{i_pj_p}) \ . \tag{1.2}$$

It is assumed that A has no identically $-\infty$ row, and that B has no identically $-\infty$ column, so that at each stage, Players Min and Max have at least one available action with finite payment. Note that payments are algebraic, i.e., a negative payment is a positive payment in the reverse direction. A strategy is *positional* if the next state is selected as a deterministic function of the current state. A fundamental result established independently by Liggett and Lippman [LL69] and by Ehrenfeucht and Mycielski [EM79] shows that this game has a value and that it has optimal positional strategies. That is, there is a real number $\chi = \chi_{\bar{\jmath}}$, a positional strategy for Min, and a positional strategy for Max, such that the following properties hold:

- The mean payoff for Min is at most χ, if Min plays according to her positional strategy. This is independent of Max's play.
- The mean payoff for Max is at least χ, if Max plays according to his positional strategy. This is independent of Min's play.

Hence, with optimal play of both players the mean payoff for both players is exactly χ, and in this case the sequences in (1.1) and (1.2) converge to χ. We say that the initial state $\bar{\jmath}$ is *winning* for Player Max if $\chi_{\bar{\jmath}} \geq 0$. It should be noted that mean payoff games can be thought of as limits of discounted zero-sum games as the discount rate tends to 0. To decide whether or not a given state is winning is the natural decision problem MEAN-PAYOFF associated with a mean payoff game. Zwick and Paterson showed that MEAN-PAYOFF is in NP $\cap$ co-NP. It is a major open question in computational complexity whether there is a polynomial time algorithm for MEAN-PAYOFF. The following theorem characterizes the set of winning states in terms of a tropical version of a linear programming feasibility problem.

Theorem 1.3 ([AGG12, Theorem 3.2])**.** *The initial state $\bar{\jmath} \in [n]$ is winning for Player Max, in the mean payoff game with payment matrices A and B, if and only if there exists a solution $x \in \mathbb{R}^n_{\max}$ with $x_{\bar{\jmath}} = 0$, to the system:*

$$\max_{j\in[n]} A_{ij} + x_j \geq \max_{j\in[n]} B_{ij} + x_j \qquad \textit{for all} \ \ i \in [m] \ . \tag{1.3}$$

We next give more insight on this result as it is fundamental in the sequel. It relies on fixed point properties of the Shapley operator. The latter is a self-map F of $\mathbb{R}^n$, preserving the standard partial order of $\mathbb{R}^n$, which is such that $[F^k(0)]_{\bar{\jmath}}$ gives the value of the zero-sum game in *finite* horizon k with initial state $\bar{\jmath}$, with the same instantaneous

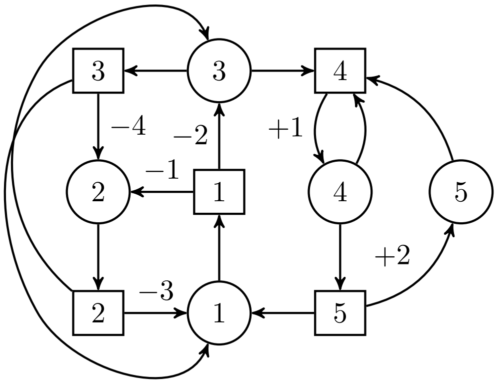

Figure 1.4: A mean payoff game. The states in which Max plays are depicted by squares, while the states in which Min plays are depicted by circles. Edges represents valid moves, and are weighted by payments. An edge with no weight indicate a 0 payment.

payments. (We denote by F^k the k-th iterate of F.) The limit of $[F^k(0)/k]_{\bar{\jmath}}$ as $k \to \infty$, i.e., the limit of the value per turn of the finite horizon game, is known to coincide with the value of the mean payoff game. The Shapley operator F does extend to a self-map of $\mathbb{R}^n_{\max}$. It is shown in [AGG12] that the value of the mean payoff game is nonnegative if and only if there exists a vector $x \in \mathbb{R}^n_{\max}$ such that $x_{\bar{\jmath}} \neq -\infty$ and $F(x) \geq x$, the latter inequality being equivalent to (1.3). Due to the homogeneity of the constraints in (1.3) there is a solution with $x_{\jmath} \neq -\infty$ if and only if there is a solution with $x_{\jmath} = 0$. A feasible point x serves as a certificate that all initial states j with $x_j \neq -\infty$ are winning. Also, if a feasible point x is known, a winning strategy for Player Max is obtained by moving from every state $i \in [m]$ to a state j achieving the maximum in $\max_{j \in [n]}(A_{ij} + x_j)$.

Example 1.4. The mean payoff game with the following payment matrices is depicted in Figure 1.4 (for the sake of readability, $-\infty$ entries are represented by the symbol "$\cdot$"):

$$A = \begin{pmatrix} \cdot & -1 & -2 & \cdot & \cdot \\ -3 & \cdot & 0 & \cdot & \cdot \\ 0 & -4 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & +1 & \cdot \\ 0 & \cdot & \cdot & \cdot & +2 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & \cdot & \cdot & \cdot & \cdot \\ \cdot & 0 & \cdot & \cdot & \cdot \\ \cdot & \cdot & 0 & \cdot & \cdot \\ \cdot & \cdot & 0 & 0 & 0 \\ \cdot & \cdot & \cdot & 0 & \cdot \end{pmatrix}.$$

In this game, the only winning initial states for Max are 4 and 5. Indeed, the point $(-\infty, -\infty, -\infty, 0, 0)$ is a solution of the system of tropical linear inequalities corresponding to the matrices A and B:

$$\begin{aligned} \max(x_2 - 1, x_3 - 2) &\geq x_1 \\ \max(x_1 - 3, x_3) &\geq x_2 \\ \max(x_1, x_2 - 4) &\geq x_3 \\ x_4 + 1 &\geq \max(x_3, x_4, x_5) \\ \max(x_1, x_5 + 2) &\geq x_4 \ . \end{aligned}$$

Each solution $x \in \mathbb{R}^5_{\max}$ satisfies $x_1 = x_2 = x_3 = -\infty$.

1.1.5 Algorithms for tropical linear programming

Several algorithms have been proposed for tropical linear programming or related problems. The alternating projection method introduce by Cuninghame-Green and Butkovič in [CGB03] determines the feasibility of a tropical polyhedra in pseudo-polynomial time [BA08], see also [But10, Chapter 10]. This was extended in [GS07] to cyclic projections (allowing one to determine a point in the intersection of more than two tropical convex sets), and applied in [AGNS11] to the situation in which a tropical convex set is given as an intersection of halfspaces. The algorithm proposed in [BZ06] also solves tropical linear feasibility problems, but exhibits an exponential behavior on a class of examples found by Bezem, Nieuwenhuis and Rodríguez-Carbonell [BNRC08]. Integers points of tropical polyhedra can be found in strongly polynomial time under genericity conditions [BM14a, BM14b]. The tropical double description method [AGG13] computes an internal representation of a tropical polyhedron described by inequalities. Hence, it provides an algorithm for tropical linear programming. However, the size of an internal representation grows exponentially with the dimension and the number of inequalities [AGK11].

Since, as we saw in Section 1.1.4, tropical linear feasibility problems are equivalent to mean payoff games, every algorithm solving mean payoff games can be applied to tropical linear programming. These include in particular value iteration algorithms [ZP96] and policy iteration algorithms [Pur95, CTGG99, JPZ06, DG06, BV07, Cha09]

A tropical linear program always arises as the tropicalization of a classical linear program over a non-Archimedean field. Hence, tropical linear programming can be thought of as an asymptotic version of linear programming [Jer73a], and the approach of Filar, Altman and Avrachenkov [FAA02] should also solve tropical linear programs.

The more general problem of tropical factional linear programming can be solved by the algorithms presented in [GKS12] and [GMH14].

1.2 Contributions

In this thesis, we present new results on the complexity of classical linear programming on the one hand, and of tropical linear programming and mean payoff games on the other hand. Our contributions lie in the study of the interplay between these two problems provided by the dequantization process.

Concerning classical linear programming, our first contribution is a tropicalization of the simplex method. More precisely, we describe an implementation of the simplex method that, under genericity conditions, solves a linear program over an ordered field. Our implementation uses only the restricted information provided by the valuation map, which corresponds to the "orders of magnitude" of the input. Consequently, the number of iterations of the simplex method can be measured in terms of the value of these "orders of magnitude". Using this approach, we exhibit a class of classical linear programs over the real numbers on which the simplex method, with any pivoting rule, performs a number of iterations which is polynomial in the input size of the problem. In particular, this implies that the corresponding polyhedra have a diameter which is polynomial in the input size.

Our second contribution to classical linear programming comes from the study of the dequantization of the central path. We disprove the continuous analog of the Hirsch conjecture proposed by Deza, Terlaky and Zinchenko, by constructing a family of linear programs with $3r + 4$ inequalities in dimension $2r + 2$ where the central path has a total curvature which is exponential in r. Our counter-example is obtained as a deformation of a family of tropical linear programs introduced by Bezem, Nieuwenhuis and Rodríguez-Carbonell. We also point out suprising features of the tropical central path. For example it has a purely geometric description, while the classical central path depends on the algebraic representation of a linear program. Moreover, the tropical central path may lie on the boundary of the tropicalization of the feasible set, and may even coincide with a path of the tropical simplex method.

Concerning tropical linear programming and mean payoff games, our main result is a new procedure to solve these problems based on the tropicalization of the simplex method. The latter readily applies to tropical linear programs satisfying genericity conditions. In order to solve arbitrary problems, we devise a new perturbation scheme. Our main tool is to use tropical semirings based on additive groups of vectors ordered lexicographically.

Then, we transfer complexity results from classical to tropical linear programming. We show that the existence of a polynomial-time pivoting rule for the classical simplex method, satisfying additional assumptions, would provide a polynomial algorithm for tropical linear programming and thus for mean payoff games. By transferring the analysis of the shadow-vertex rule of Adler, Karp and Shamir, we also obtain the first algorithm that solves mean payoff games in polynomial time on average, assuming the distribution of the games satisfies an symmetry property.

We establish tropical counterparts of the notions of basic points and edges of a polyhedron. This yields a geometric interpretation of the tropicalization of the simplex method. As in the classical case, the tropical algorithm pivots on the graph of an arrangement of hyperplanes associated to a tropical polyhedron. This interpretation is based on a geometric connection between the cells of an arrangement of classical hyperplanes and their tropicalization. Building up on this geometric interpretation, we present algorithmic refinements of the tropical pivoting operation. We show that pivoting along an edge of a tropical polyhedron defined by m inequalities in dimension n can be done in time $O(n(m+n))$, a complexity similar to the classical pivoting operation. We also show that the computation of reduced costs can be done tropically in time $O(n(m+n))$.

These algorithmics refinements, along with our perturbation scheme, have been implemented in the library Simplet [Ben14]. Hence, this library provides a solver for arbitrary tropical linear programs.

1.3 Organisation of the manuscript

This manuscript is organized as follows.

- Chapter 2 presents the framework used throughout this manuscipt. It recalls the definitions of tropical semirings, non-Archimedean fields, and related notions.
- Chapter 3 exposes the tropical implementation of the simplex method.
- In Chapter 4, we study tropical polyhedra and their relations with classical polyhedra. We also devise the perturbation scheme that allows to solve arbitrary tropical linear programs with the tropical simplex method
- In Chapter 5, we transfer complexity results based on the simplex method back and forth between tropical and classical linear programming.
- Chapter 6 concerns the tropicalization of the shadow-vertex rule, and of the average case analysis of Adler, Karp and Shamir.
- Chapter 7 exposes algorithmic refinements of the tropical simplex method.
- Chapter 8 deals with the tropical analysis of the central path.

Most of the notions needed to read this manuscript are given in Chapters 2, 3 and 4. The other chapters can mostly be read independently, even if the approach in Chapter 6 uses ideas already present in Chapter 5.

The tropicalization of simplex operations (pivoting and computing reduced costs) was presented in [ABGJ13b]. In [ABGJ13a], the tropicalization of *combinatorial* pivoting rules was presented (combinatorial pivoting rules rely only on the signs of the minors of the input matrix). The study of the tropical shadow-vertex rule [ABG14] led to the more general framework of *semi-algebraic* pivoting rules that we adopt in Chapter 3. Chapter 4 gathers results that appeared in [ABGJ13b] and [ABGJ13a]. Chapter 5 generalizes to semi-algebraic pivoting rules the transfer of complexity theorem presented in [ABGJ13a], and also contains new results. Chapter 6 is covered in [ABG14] and Chapter 7 is included in [ABGJ13b]. Chapter 8 is mostly covered in [ABGJ14], but includes a slight improvement of the curvature analysis.

Chapter 2

Preliminaries

In this chapter, we recall the definitions of (totally) ordered abelian groups and (totally) ordered fields. Note that the orders on the structures that we consider will always be total order. We also present basic notions of model theory. In particular, the notion of completeness of a theory will play an important role. Indeed, the completeness of the theory of ordered field will allow us to transfer results that holds on the field of real numbers to other ordered fields. We also present the framework we shall work with: tropical semirings and non-Archimedean fields. The tropical semirings we consider are constructed from arbitrary ordered groups. They arise as the image under the Archimedean valuation map of ordered fields, such as the field of formal Hahn series.

2.1 Model theory

We recall some definitions and results of model theory, referring the reader to [Mar02] for more background.

2.1.1 Languages and first-order formulæ

A *language* $\mathcal{L} = (\mathcal{R}, \mathcal{F}, \mathcal{C})$ consists of a set of relation symbols $\mathcal{R}$, a set of function symbols $\mathcal{F}$, and a set of constant symbols $\mathcal{C}$. Each relation symbol $R \in \mathcal{R}$ is equipped with an *arity*, n_R, which is a positive integer. Similarly, each function symbol $f \in \mathcal{F}$ also has an arity, denoted as n_f. For example, the language of ordered groups is $\mathcal{L}_{\mathrm{og}} = (<, +, 0)$, and the language of ordered rings is $\mathcal{L}_{\mathrm{or}} = (<, +, \cdot, 0, 1)$. In these two languages, the order relation symbol $<$ and the arithmetic function symbols $+, \cdot$ have arity two. We shall now describe the (first-order) formulæ of a language $\mathcal{L}$. Such a formula is a string of symbols built from the symbols of $\mathcal{L}$, a finite number of variable symbols $v_1, v_2, \dots,$ the equality symbol $=$, the Boolean symbols $\neg$, $\wedge$,$\vee$ ("not", "and", "or"), quantifiers $\forall$, $\exists$ and parentheses (,). An $\mathcal{L}$-*term* is either:

- a variable symbol v_i, for some $i \geq 1$
- a constant symbol $c \in \mathcal{C}$

- $f(t_1, \ldots, t_{n_f})$ where $f \in \mathcal{F}$ is a function symbol, and $t_1, \ldots t_{n_f}$ are $\mathcal{L}$-terms

An $\mathcal{L}$-*formula* is then defined inductively as follows:

- if t_1 and t_2 are $\mathcal{L}$-terms, then $t_1 = t_2$ is an $\mathcal{L}$-formula
- if $R \in \mathcal{R}$ is a relation symbol, and $t_1, \ldots t_{n_R}$ are terms, then $R(t_1, \ldots t_{n_R})$ is an $\mathcal{L}$-formula
- if ϕ and ψ are $\mathcal{L}$-formulæ, then $(\neg\phi)$, $(\phi \wedge \psi)$ and $(\phi \vee \psi)$ are $\mathcal{L}$-formulæ
- if ϕ is an $\mathcal{L}$-formula and v_i is a variable symbol, then $\exists v_i \phi$ and $\forall v_i \phi$ are $\mathcal{L}$-formulæ

A variable symbol v_i which occurs in a formula ϕ without being modified by a quantifier $\exists$ or $\forall$ is said to be a *free variable*. We shall emphasize the free variables $v_{i_1}, \ldots, v_{i_k}$ of a formula ϕ by writing $\phi(v_{i_1}, \ldots, v_{i_k})$. A formula without free variables is called a *sentence*. In the following, we shall use the usual abbreviations in first-order formulæ:

- $\phi \rightarrow \psi$ for $\neg\phi \vee \psi$
- $\phi \leftrightarrow \psi$ for $(\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$
- $\bigwedge_{i=1}^{n} \phi_i$ for $\phi_1 \wedge \phi_2 \wedge \cdots \wedge \phi_n$
- $\bigvee_{i=1}^{n} \phi_i$ for $\phi_1 \vee \phi_2 \vee \cdots \wedge \phi_n$
- $\exists! v\phi(v)$ for $\exists v(\phi(v) \wedge (\forall w(\phi(w) \rightarrow w = v))$.

2.1.2 Structures

Let $\mathcal{L} = (\mathcal{R}, \mathcal{F}, \mathcal{C})$ be a language. An $\mathcal{L}$-*structure* $\mathfrak{M}$ consists of a non-empty set M (called the *domain* of $\mathfrak{M}$) together with an *interpretation* $\mathcal{L}^{\mathfrak{M}}$ of the symbols of $\mathcal{L}$ in M. A relation symbol $R \in \mathcal{R}$ is interpreted by a subset $S_R \subseteq M^{n_R}$, where a tuple $(x_1, \ldots, x_{n_R})$ satisfies the relation R if $(x_1, \ldots, x_{n_R}) \in S_R$. A function symbol $f \in \mathcal{F}$ is interpreted by a map $M^{n_f} \rightarrow M$, and the interpretation of a constant symbol $c \in \mathcal{C}$ is an element of M. For example, the ordered field of real numbers $\bar{\mathbb{R}} = (\mathbb{R}, <, +, \cdot, 0, 1)$ is a $\mathcal{L}_{\text{or}}$-structure. In the following, we shall use the same notation $\mathcal{L}$ for a language and its interpretation. Hence, we will denote a $\mathcal{L}$-structure $\mathfrak{M}$ by $(M, \mathcal{L})$.

The interpretation of the language $\mathcal{L}$ induces an interpretation of the formulæ of $\mathcal{L}$ in the structure $\mathfrak{M}$. Every formula $\phi(v_{i_1}, \ldots, v_{i_k})$ defines a Boolean function $\phi^{\mathfrak{M}}$ on M^k. If $\phi^{\mathfrak{M}}$ is true at $a \in M^k$, we write $\mathfrak{M} \models \phi(a)$. In particular, if ϕ is a sentence of $\mathcal{L}$, the function $\phi^{\mathfrak{M}}$ is constant. Thus a sentence defines a statement on $\mathfrak{M}$ which is either true or false.

For example, consider the following formula in the language of ordered rings $\mathcal{L}_{\text{or}}$:

$$\phi(v) : \exists y \; (y \cdot y = v) \; .$$

When interpreted in a $\mathcal{L}_{\text{or}}$-structure $\mathfrak{M}$, the sentence $\phi(v)$ yields a boolean function which is true at $a \in M$ if a admits a square root in M. Now consider the following sentence:

$$\psi : \forall x\ (x > 0 \to \exists y\ (y \cdot y = x))\ .$$

A $\mathcal{L}_{\text{or}}$-structure $\mathfrak{M}$ satisfy $\mathfrak{M} \models \psi$ if every positive element of the domain of $\mathfrak{M}$ has a square root. This statement is false on $\mathcal{L}_{\text{or}}$-structure of the rational numbers $\bar{\mathbb{Q}} = (\mathbb{Q}, <, +, \cdot, 0, 1)$, but true on the ordered field of the real numbers $\bar{\mathbb{R}}$.

2.1.3 Complete theories

A $\mathcal{L}$-*theory* T is a set of sentences of the language $\mathcal{L}$, the *axioms* of the theory T. A $\mathcal{L}$-structure $\mathfrak{M}$ is a *model* of the theory T if $\mathfrak{M} \models \phi$ for all axioms $\phi \in T$. We say that a $\mathcal{L}$-theory is *complete* if, for any two models $\mathfrak{M}, \mathfrak{N}$ of the theory T, a $\mathcal{L}$-sentence ϕ is true in $\mathfrak{M}$ if and only if it is true in $\mathfrak{N}$.

We shall describe two complete theories: the theory of ordered divisible abelian groups and the theory of real closed fields.

Ordered abelian groups

The theory of abelian groups consists of the following sentences in the language $\mathcal{L}_{\text{g}} = \{+, 0\}$.

$$\begin{aligned}
&\forall x\ 0 + x = x + 0 = x && \text{(identity element)}\\
&\forall x \forall y \forall z\ x + (y + z) = (x + y) + z && \text{(associativity)}\\
&\forall x \exists y\ x + y = y + x = 0 && \text{(invertibility)}\\
&\forall x \forall y\ x + y = y + x && \text{(abelian)}
\end{aligned}$$

The theory of divisible abelian group is obtained by adding, for all integers $n \geq 2$, the axiom:

$$\forall x \exists y\ \underbrace{y + y + \cdots + y}_{n \text{ times}} = x\ . \qquad \text{(divisibility)}$$

Ordered (divisible) abelian groups are described in the language $\mathcal{L}_{\text{og}} = \{<, +, 0\}$ by the axioms of (divisible) abelian groups along with:

$$\begin{aligned}
&\forall x\ \neg(x < x) && \text{(irreflexivity)}\\
&\forall x \forall y \forall z\ ((x < y) \wedge y < z) \to x < z) && \text{(transitivity)}\\
&\forall x \forall y\ (x < y \vee x = y \vee y < x) && \text{(totality)}\\
&\forall x \forall y \forall z\ (x < y \to x + z < y + z) && \text{(translation-invariance)}
\end{aligned}$$

The non-strict order relation $\leq$ is then defined as $\forall x \forall y\ (x \leq y) \leftrightarrow (x < y \vee x = y)$.

Theorem 2.1 ([Mar02, Corollary 3.1.17])**.** *The theory of ordered divisible abelian groups is complete.*

This result can be traced back to Robinson [Rob77, Theorem 4.3.2].

Ordered fields

The theory of ordered fields is described in the language $\mathcal{L}_{\mathrm{or}} = (<, +, \cdot, 0, 1)$. It consists of the axioms of ordered abelian groups in the language $(<, +, 0)$, along with:

$$\begin{aligned}
&\forall x\ 1 \cdot x = x \cdot 1 = 1 && \text{(multiplicative identity element)}\\
&\forall x \forall y\ x \cdot y = y \cdot x && \text{(commutativity of multiplication)}\\
&\forall x \forall y \forall z\ x \cdot (y \cdot z) = (x \cdot y) \cdot z && \text{(multiplicative associativity)}\\
&\forall x\ (x = 0 \vee \exists y\ x \cdot y = 1) && \text{(multiplicative invertibility)}\\
&\forall x \forall y \forall z\ x \cdot (y + z) = (x \cdot y) + (x \cdot z) && \text{(distributivity)}\\
&\forall x \forall y \forall z\ ((x < y) \wedge z > 0) \to x \cdot z < y \cdot z) && \text{(scaling-invariance)}
\end{aligned}$$

The theory of real closed fields is obtained from the theory of ordered fields by also requiring every positive element to be a square root and every polynomial of odd degree to have at least one root. In symbols, the first requirement is described by:

$$\forall x\ (x > 0 \to \exists y\ ; (y \cdot y = x))\ ,$$

while the second requirement amounts to the set of sentences $\{\phi_n \mid n \geq 0\}$, defined by:

$$\phi_n : \forall a_0 \forall a_1 \dots \forall a_{2n+1} \exists x\ (a_{2n+1} \cdot x^{2n+1} + a_{2n} \cdot x^{2n} + \dots + a_1 \cdot x + a_0 = 0)\ ,$$

where we use x^k as an abbreviation for $\underbrace{x \cdot x \cdot \ldots \cdot x}_{k \text{ times}}$.

Theorem 2.2 ([Mar02, Corollary 3.3.16])**.** *The theory of real closed fields is complete.*

This theorem follows from the work of Tarski and Seidenberg [Tar51, Sei54] who proved that the theory of real closed fields admits elimination of quantifiers.

2.2 Tropical semirings and non-Archimedean ordered fields

The theory of ordered (commutative) semirings, in the language $(<, +, \cdot, 0, 1)$, is obtained from the theory of ordered fields by removing the axioms asserting the existence of inverse elements for the addition and the multiplication, and adding the axiom $\forall x\ 0 \cdot x = x \cdot 0 = 0$. The set of non-negative integers, or the set of non-negative real numbers, with their natural ordering and usual addition and multiplication, are examples of ordered semirings.

We are mainly interested in semirings formed with the non-negative elements of an ordered field. On these semirings, the notion of "order of magnitude" is captured by Archimedean classes. The map that sends an element to its Archimedean class is a homomorphism to another semiring, called a tropical semiring. We begin with the description of tropical semirings, and explain how they arise from ordered fields. We then present non-Archimedean fields, that is fields with a non-trivial set of Archimedean classes. First, the field of Hahn series consists of formal power series with exponents in an arbitrary ordered group. Second, we present fields of real-valued functions, called Hardy fields. Such a field consists of functions that are definable in a o-minimal structure. The advantage of an Hardy field over Hahn series is that it consists of real-valued functions, rather than formal objects.

2.2.1 Tropical semirings

We are interested in a specific class of ordered semirings, that we call *tropical semirings*. We shall describe them in the language $(<, \oplus, \odot, \mathbb{0}, \mathbb{1})$. Given an ordered abelian group $(G, <, +)$, the *tropical semiring* $(\mathbb{T}(G), <, \oplus, \odot, \mathbb{0}, \mathbb{1})$ is defined as follows. The base set is $\mathbb{T}(G) = G \cup \{\mathbb{0}_{\mathbb{T}(G)}\}$ where the new element $\mathbb{0}_{\mathbb{T}(G)} \notin G$ satisfies $\mathbb{0}_{\mathbb{T}(G)} > a$ for all $a \in G$. The additive law $\oplus$ is defined by $a \oplus b = \max(a, b)$, where the maximum is taken with respect to $<$. The multiplicative law $\odot$ is the group addition $+$ extended to $\mathbb{T}(G)$ by setting $a + \mathbb{0}_{\mathbb{T}(G)} = \mathbb{0}_{\mathbb{T}(G)} + a = \mathbb{0}_{\mathbb{T}(G)}$ for all $a \in \mathbb{T}(G)$. The zero and unit elements are $\mathbb{0}_{\mathbb{T}(G)}$ and $\mathbb{1}_{\mathbb{T}(G)} := 0_G$, the neutral element of G, respectively. In the following, we simply denote $\mathbb{T}(G)$ by $\mathbb{T}$ when this is clear from the context.

The operations are extended to matrices $A = (A_{ij}), B = (B_{ij})$ with entries in $\mathbb{T}$ by setting $A \oplus B = (A_{ij} \oplus B_{ij})$ and $A \odot B = (\bigoplus_k A_{ik} \oplus B_{kj})$. In the following, unless explicitly stated, the entries of a matrix A are denoted by A_{ij}. Moreover, we denote by A_I the submatrix of A obtained with rows indexed by I. By abuse of notation, we denote $A_{\{i\}}$ by A_i. We also denote by $A^\top$ the transpose of the matrix A. 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 $\arg(a \odot x)$ 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 $\leq$ on G extends to $\mathbb{T}$. This induces a partial ordering of tropical vectors by entry-wise 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 $(x, y) \mapsto x \oplus y := (\max(x_1, y_1), \dots, \max(x_n, y_n))$.

2.2.2 Non-Archimedean fields

Let K be an ordered field. Two elements $\boldsymbol{a}, \boldsymbol{b} \in K$ satisfy the *Archimedean relation* if they are within a rational multiple of each other, *i.e.*, if there exists a positive rational number r such that $r^{-1}|\boldsymbol{b}| < |\boldsymbol{a}| < r|\boldsymbol{b}|$, where $|\boldsymbol{a}| = \max(\boldsymbol{a}, -\boldsymbol{a})$ is the absolute value. The equivalence class of $\boldsymbol{a}$ for the Archimedean relation is called the *value* of $\boldsymbol{a}$ and is denoted by $\operatorname{val}(\boldsymbol{a})$. The map $\operatorname{val} : \boldsymbol{a} \mapsto \operatorname{val}(\boldsymbol{a})$ is called a *valuation map*.

We shall see below that $\operatorname{val}(K \setminus \{0\})$ is an ordered group, called the *value group* of K. A field K is *non-Archimedean* if its value group is not the trivial group, *i.e.*, if $K \setminus \{0\}$ has more than one Archimedean class.

A *lift* of a value $a \in \mathrm{val}(K)$ is an element $\boldsymbol{a} \in K$ such that $\mathrm{val}(\boldsymbol{a}) = a$. Clearly, such a lift is by no means unique. The set of all lifts will be denoted $\mathrm{val}^{-1}(\boldsymbol{a})$. The set of values $\mathrm{val}(K)$ inherits an order from the order on K, and an operation $\odot$ from the multiplication on K. They are defined for $a, b \in \mathrm{val}(K)$ by:

$$
\begin{aligned}
& a < b \Leftrightarrow a \neq b \text{ and } |\boldsymbol{a}| < |\boldsymbol{b}| && \text{for any } \boldsymbol{a} \in \mathrm{val}^{-1}(a), \boldsymbol{b} \in \mathrm{val}^{-1}(b) \\
& a \odot b = \mathrm{val}(\boldsymbol{ab}) && \text{for any } \boldsymbol{a} \in \mathrm{val}^{-1}(a), \boldsymbol{b} \in \mathrm{val}^{-1}(b).
\end{aligned}
$$

These definitions do not depend on the choice of the lifts $\boldsymbol{a}, \boldsymbol{b}$.

Proposition 2.3. *The ordering relation $a < b$ and the operation $a \odot b$ are well-defined for any $a, b \in \mathrm{val}(K)$. Moreover, the set $\mathrm{val}(K \setminus \{0\})$, equipped with the order $<$, the operation $\odot$, and the neutral element $\mathbb{1} := \mathrm{val}(1)$, forms an ordered abelian group.*

Proof. First, let us show that $a < b$ is well defined for any $a, b \in \mathrm{val}(K)$. Suppose that $a \neq b$. Consider any $\boldsymbol{a} \in \mathrm{val}^{-1}(a), \boldsymbol{b} \in \mathrm{val}^{-1}(b)$, and assume without loss of generality that $|\boldsymbol{a}| < |\boldsymbol{b}|$. Then, we claim that $|\boldsymbol{a}'| < |\boldsymbol{b}'|$ for all $\boldsymbol{a}' \in \mathrm{val}^{-1}(a)$ and $\boldsymbol{b}' \in \mathrm{val}^{-1}(b)$. By definition, there exists two rational numbers $p, q \in \mathbb{Q}$ such that $|\boldsymbol{a}'| < p|\boldsymbol{a}|$ and $|\boldsymbol{b}| < q|\boldsymbol{b}'|$. Since $|\boldsymbol{a}| < |\boldsymbol{b}|$, we obtain $|\boldsymbol{a}'| < pq|\boldsymbol{b}'|$. Suppose, by contradiction, that $|\boldsymbol{a}'| \geq |\boldsymbol{b}'|$. Then, $|\boldsymbol{a}'| > \frac{1}{2}|\boldsymbol{b}'|$. Consider the rational number $r = \max(pq, 2)$. We have $r^{-1}|\boldsymbol{b}'| < |\boldsymbol{a}'| < r|\boldsymbol{b}'|$. Hence, $\boldsymbol{a}'$ and $\boldsymbol{b}'$ belongs to the same Archimedean class. This contradicts the hypothesis $\mathrm{val}(\boldsymbol{a}) \neq \mathrm{val}(\boldsymbol{b})$, and proves our claim.

Second, let us show $a \odot b = \mathrm{val}(\boldsymbol{ab})$ is independent of the choice of $\boldsymbol{a} \in \mathrm{val}^{-1}(a)$ and $\boldsymbol{b} \in \mathrm{val}^{-1}(b)$. Consider any lifts $\boldsymbol{a}, \boldsymbol{a}' \in \mathrm{val}^{-1}(a)$ and$\boldsymbol{b}, \boldsymbol{b}' \in \mathrm{val}^{-1}(b)$. By definition, there exist two rational numbers $p, q \in \mathbb{Q}$ such that $p^{-1}|\boldsymbol{a}'| < |\boldsymbol{a}| < p|\boldsymbol{a}'|$ and $q^{-1}|\boldsymbol{b}'| < |\boldsymbol{b}| < q^{-1}|\boldsymbol{b}'|$. It follows that $(pq)^{-1}|\boldsymbol{a}'\boldsymbol{b}'| < |\boldsymbol{ab}| < (pq)|\boldsymbol{a}'\boldsymbol{b}'|$, and thus $\mathrm{val}(\boldsymbol{a}'\boldsymbol{b}') = \mathrm{val}(\boldsymbol{ab})$.

It follows immediately from the definition that $\odot$ is commutative, admits $\mathbb{1} = \mathrm{val}(1)$ as a neutral element, and that any element $\mathrm{val}(\boldsymbol{a}) \in \mathrm{val}(K \setminus \{0\})$ has an inverse $\mathrm{val}(\boldsymbol{a}^{-1})$. Moreover, it is also easy to see that $a > b$ implies $a \odot c > b \odot c$ for all $a, b \in \mathrm{val}(K)$ and $c \in \mathrm{val}(K \setminus \{0\})$. Consequently, $(\mathrm{val}(K \setminus \{0\}), <, \odot)$ is an ordered abelian group. ☐

In the following, we shall always identify $\mathrm{val}(K)$ with the tropical semiring based on the value group of K, where we have set $\mathbb{0} := \mathrm{val}(0)$. The valuation map satisfies the following crucial properties.

Proposition 2.4. *Let $(K, <, +, \cdot, 0, 1)$ be an ordered group and $(\mathrm{val}(K), <, \oplus, \odot, \mathbb{0}, \mathbb{1})$ the tropical semiring of its value group. The valuation map, restricted to the set K_+ of non-negative elements of K, is a homomorphism of ordered semirings, i.e., for any $\boldsymbol{a}, \boldsymbol{b} \in K_+$, we have:*

$$
\begin{aligned}
\mathrm{val}(\boldsymbol{a} + \boldsymbol{b}) &= \mathrm{val}(\boldsymbol{a}) \oplus \mathrm{val}(\boldsymbol{b}) \\
\mathrm{val}(\boldsymbol{ab}) &= \mathrm{val}(\boldsymbol{a}) \odot \mathrm{val}(\boldsymbol{b}) \ . \\
\boldsymbol{a} \geq \boldsymbol{b} &\Rightarrow \mathrm{val}(\boldsymbol{a}) \geq \mathrm{val}(\boldsymbol{b})
\end{aligned}
$$

Proof. The identity $\mathrm{val}(\boldsymbol{ab}) = \mathrm{val}(\boldsymbol{a}) \odot \mathrm{val}(\boldsymbol{b})$, and the implication $\boldsymbol{a} \geq \boldsymbol{b} \Rightarrow \mathrm{val}(\boldsymbol{a}) \geq \mathrm{val}(\boldsymbol{b})$ are direct consequences of the definition of the order on $\mathrm{val}(K)$ and of the operation

$\odot$. We now prove the equality $\mathrm{val}(\boldsymbol{a}+\boldsymbol{b}) = \mathrm{val}(\boldsymbol{a}) \oplus \mathrm{val}(\boldsymbol{b})$. Consider two positive elements $\boldsymbol{a}, \boldsymbol{b} \in K$. Without loss of generality, we can assume that $\boldsymbol{a} \geq \boldsymbol{b}$. We have the inequalities $2\boldsymbol{a} \geq \boldsymbol{a} + \boldsymbol{b} \geq \boldsymbol{a}/2$, and thus $\mathrm{val}(\boldsymbol{a} + \boldsymbol{b}) = \mathrm{val}(\boldsymbol{a}) = \max(\mathrm{val}(\boldsymbol{a}), \mathrm{val}(\boldsymbol{b}))$. $\square$

2.2.3 Signed tropical numbers, and the signed valuation map

We now enhance the valuation map with a sign information. The *signed tropical numbers* $\mathbb{T}_\pm = \mathbb{T}_+ \cup \mathbb{T}_-$ consist of two copies of $\mathbb{T}$, the set of *positive tropical numbers* $\mathbb{T}_+$ and the set of *negative tropical numbers* $\mathbb{T}_-$. These elements are respectively denoted as a and $\ominus a$ for $a \in \mathbb{T}$. The elements a and $\ominus a$ are distinct unless $a = \mathbb{0}_\mathbb{T}$. In the latter case, these two elements are identified, *i.e.*, we have $\mathbb{0}_\mathbb{T} = \ominus \mathbb{0}_\mathbb{T}$. The *sign* of the elements a and $\ominus a$ are $\mathrm{sign}(a) = 1$ and $\mathrm{sign}(\ominus a) = -1$, respectively, when a is not $\mathbb{0}_\mathbb{T}$, and $\mathrm{sign}(\mathbb{0}_\mathbb{T}) = 0$. The *reflection map* $a \mapsto \ominus x$ sends a positive element a to $\ominus a$, and a negative element $\ominus a$ to a. 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 $\mathrm{sign}(x)\,\mathrm{sign}(y)$. The *positive part* and the *negative part* of an element $a \in \mathbb{T}_\pm$ are the tropical numbers a^+ and a^- defined by:

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

When $\mathbb{T} = \mathrm{val}(K)$ is the tropical semiring of an ordered field, we define the *signed valuation map* $\mathrm{sval} : K \to \mathbb{T}_\pm$ by:

$$\mathrm{sval} : \boldsymbol{a} \mapsto \begin{cases} \mathrm{val}(\boldsymbol{a}) & \text{if } \boldsymbol{a} \geq 0 \ , \\ \ominus \mathrm{val}(\boldsymbol{a}) & \text{otherwise.} \end{cases}$$

Consequently, $\boldsymbol{a} \in K$ and $\mathrm{sval}(\boldsymbol{a})$ have the same sign, $\mathrm{sval}(-\boldsymbol{a}) = \ominus\,\mathrm{sval}(\boldsymbol{a})$, and the modulus of $\mathrm{sval}(\boldsymbol{a})$ is $\mathrm{val}(|\boldsymbol{a}|)$. Moreover, the signed valuation map preserves the multiplication: we have $\mathrm{sval}(\boldsymbol{ab}) = \mathrm{sval}(\boldsymbol{a}) \odot \mathrm{sval}(\boldsymbol{b})$ for any $\boldsymbol{a}, \boldsymbol{b} \in K$. Note that we do not equip the tropical signed numbers with an additive law, as $a \oplus (\ominus a)$ would not be defined. Similarly, we do not define an order on $\mathbb{T}_\pm$. However, $\mathbb{T}_\pm$ can be embedded into a semiring, called the *symmetrized* tropical semiring; see Section 7.2.1. In the following, an element of an ordered field will be written in bold and its signed value with a standard font, *e.g.*, $a = \mathrm{val}(\boldsymbol{a})$. Modulus, the reflection map, positive and negative parts, and signed valuation extend to matrices entry-wise.

2.2.4 Hahn series

Given a (totally) ordered abelian group $(G, <, +, 0)$, the set of *Hahn series* $\mathbb{R}[\![t^G]\!]$ with value group $(G, <, +, 0)$ and with real coefficient consists of formal series

$$\boldsymbol{a} := \sum_{g \in G} a_g t^g \ ,$$

where the coefficients a_g are real numbers, and the *support* $\mathrm{Supp}(\boldsymbol{a}) := \{g \in G \mid a_g \neq 0\}$ is *anti well-ordered* (*i.e.*, every non-empty subset of $\mathrm{Supp}(\boldsymbol{a})$ admits a greatest element). Note that Hahn series are sometimes defined as series with well-ordered support, instead of anti well-ordered. As we use the $(\max, +)$ notation for tropical semirings, we find it more natural to adopt the anti well-ordered definition. We shall equip the set of Hahn series with an $\mathcal{L}_{\mathrm{or}}$-structure $(\mathbb{R}[\![t^G]\!], <, +, \cdot, 0, 1)$. The constant symbol 0 is interpreted as the unique Hahn series with empty support, while 1 is the Hahn series $1_{\mathbb{R}} t^{0_G}$, where $1_{\mathbb{R}}$ is the multiplicative identity of $\mathbb{R}$ and 0_G the additive identity of G. Addition and multiplication are defined, for any $\boldsymbol{a} = \sum_{g\in G} a_g t^g$ and $\boldsymbol{b} = \sum_{g \in G} b_g t^g$, by:

$$\boldsymbol{a} + \boldsymbol{b} := \sum_{g \in G} (a_g + b_g) t^g \ ,$$

$$\boldsymbol{a}\boldsymbol{b} := \sum_{g\in G} \left(\sum_{\substack{\alpha,\beta \in G, \\ \alpha+\beta = g}} a_\alpha b_\beta \right) t^g \ .$$

Theorem 2.5. *Hahn series, equipped with the addition and multiplication of formal series, forms a field.*

Proof. We give a sketch of proof for the sake of completeness, and refer to [Pas77, Chap. 13, Theorem 2.11] for details. Let $A = \mathrm{Supp}(\boldsymbol{a})$ and $B = \mathrm{Supp}(\boldsymbol{b})$. To prove the stability of the addition, it is sufficient to observe that $\mathrm{Supp}(\boldsymbol{a} + \boldsymbol{b}) \subseteq A \cup B$ is anti well-ordered, since A and B are.

For the multiplication, one first show that $\mathrm{Supp}(\boldsymbol{a}\boldsymbol{b})$ is well-ordered. It comes from the fact that $\mathrm{Supp}(\boldsymbol{a}\boldsymbol{b})$ is a subset of the Minkowski sum $A + B$, and that $A + B$ is anti well-ordered. Then, one can show that for any $g \in \mathrm{Supp}(\boldsymbol{a}\boldsymbol{b})$, there is a finite number of $(\alpha, \beta) \in A \times B$ such that $\alpha + \beta = g$. Consequently, the coefficient of t^g in $\boldsymbol{a}\boldsymbol{b}$ is a well-defined real number.

Proving that a $\boldsymbol{a} \neq 0$ has a multiplicative inverse is slightly more involved. In a nutshell, $\boldsymbol{a}$ can always be written as $\boldsymbol{a} = ct^g(1 - \boldsymbol{w})$ with c a non-zero real number, $g \in G$ and $\boldsymbol{w}$ a Hahn series with support included in the negative elements of G. Clearly $c^{-1}t^{-g}$ is a multiplicative inverse of ct^g. One then proves that the geometric series $1 + \sum_{i=1}^{+\infty} \boldsymbol{w}^i$ is a Hahn series, which is the inverse of $1 - \boldsymbol{w}$. □

The *leading coefficient* $\mathrm{lc}(\boldsymbol{x})$ of a Hahn series $\boldsymbol{a} = \sum_{g \in G} a_\alpha t^g$ is the coefficient $a_{\alpha_{\max}}$ where $\alpha_{\max}$ is the greatest element of $\mathrm{Supp}(\boldsymbol{a})$. By convention, we set $\mathrm{lc}(0) := 0$. A non-null Hahn series $\boldsymbol{a}$ is *positive* if $\mathrm{lc}(\boldsymbol{a}) > 0$, and we write $\boldsymbol{a} > 0$ in this case. Similarly, we write $\boldsymbol{a} > \boldsymbol{b}$ if $\boldsymbol{a} - \boldsymbol{b} > 0$. This definition turns $\mathbb{R}[\![t^G]\!]$ into an ordered field. Moreover, the topology induced by this order makes $\mathbb{R}[\![t^G]\!]$ a topological field.

Two Hahn series $\boldsymbol{a}, \boldsymbol{b}$ belongs to the same Archimedean class if and only if $\mathrm{Supp}(\boldsymbol{a})$ and $\mathrm{Supp}(\boldsymbol{b})$ have the same greatest element. Hence, the value group of $\mathbb{R}[\![t^G]\!]$ is G, and we can write the Archimedean valuation map as $\mathrm{val}(\boldsymbol{a}) = \max\{g \in \mathrm{Supp}(\boldsymbol{a})\}$.

Theorem 2.6 ([DW96, Theorem 2.15(iv)])**.** *If G is divisible, then $\mathbb{R}[\![t^G]\!]$ is a real closed field.*

2.2.5 The Hardy field of an o-minimal structure

We now present a particular class of ordered fields of germs real-valued functions, called Hardy fields. Such a field arises from functions that are definable in a o-minimal structure. This framework have been devised by Alessandrini to study the tropicalization of semi-algebraic sets [Ale13].

An *expansion* $\mathcal{L}'$ of a language $\mathcal{L}$ is obtained by adding some new relations, functions and constants to $\mathcal{L}$. We define an *expansion* of an $\mathcal{L}$-structure $\mathfrak{M}$ to be an $\mathcal{L}'$-structure $\mathfrak{M}'$ such that: $\mathcal{L}'$ is an expansion of $\mathcal{L}$, $\mathfrak{M}$ and $\mathfrak{M}'$ have the same domain, and the interpretation of the language $\mathcal{L}$ in $\mathfrak{M}$ coincides with the one in $\mathfrak{M}'$. In the rest of this section, $\mathfrak{R} = (\mathbb{R}, \mathcal{L})$ will denote an expansion of the $\mathcal{L}_{\mathrm{or}}$-structure of the real numbers $\bar{\mathbb{R}}$. In particular, the domain of $\mathfrak{R}$ is the set of real numbers.

A set $A \in \mathbb{R}^k$ is *definable* in $\mathfrak{R}$ if there exists an $\mathcal{L}$-formula $\phi(v_1, \dots v_k, w_1, \dots, w_l)$ and an element $b \in \mathbb{R}^l$ such that $A = \{a \in \mathbb{R}^k \mid \mathfrak{R} \models \phi(a, b)\}$. Given a definable set $A \in \mathbb{R}^k$, a map $F : A \to \mathbb{R}^l$ is *definable* if its graph $\{(a, F(a)) \mid a \in A\} \subseteq \mathbb{R}^{k+l}$ is a definable set.

The structure $\mathfrak{R}$ is *o-minimal* if every subset of $\mathbb{R}$ that is definable in $\mathfrak{R}$ is a finite union of points and intervals with endpoints in $\mathbb{R} \cup \{-\infty, +\infty\}$. Under the o-minimality requirement, definable sets and maps are "well-behaved". For example, the set $\{(x, \sin(1/x)) \mid x > 0\}$ is not definable in an o-minimal structure. We refer the reader to [vdD98] or [Cos00] for more details.

We say that two definable functions $f, g : \mathbb{R} \to \mathbb{R}$ are equivalent, and we write $f \sim g$, if $f(t) = g(t)$ *ultimately*, *i.e.*, for all t large enough. The *germ* $\boldsymbol{f}$ of a definable function f is the equivalence class of f for the relation $\sim$. By abuse of notation, $\boldsymbol{f}$ will also denote a representative of the germ $\boldsymbol{f}$.

Let $H(\mathfrak{R}) := \{\boldsymbol{f} \mid f : \mathbb{R} \to \mathbb{R} \text{ is definable in } \mathfrak{R}\}$ be the set of germs of functions definable in $\mathfrak{R}$. Each function symbol $F \in \mathcal{F}$ has a natural interpretation in $H(\mathfrak{R})$, by defining $F(\boldsymbol{f}_1, \dots, \boldsymbol{f}_{n_F})$ as the germ of the definable function $t \mapsto F(\boldsymbol{f}_1(t), \dots, \boldsymbol{f}_{n_F}(t))$. Besides, the set $\mathbb{R}$ is embedded into $H(\mathfrak{R})$ by identifying each element $a \in \mathbb{R}$ with the constant function with value a. This provides an interpretation of the constant symbols of $\mathcal{L}$ in $H(\mathfrak{R})$. Finally, given a relation R of the language $\mathcal{L}$ and $\boldsymbol{f}_1, \dots, \boldsymbol{f}_{n_R} \in H(\mathfrak{R})$, the set $\{t \mid \mathfrak{R} \models R(\boldsymbol{f}_1(t), \dots, \boldsymbol{f}_{n_R}(t))\}$ is definable, and thus consists in a finite union of points and intervals. Hence, $R(\boldsymbol{f}_1(t), \dots, \boldsymbol{f}_{n_R}(t))$ is either ultimately true or ultimately false. This provides an interpretation of R over $H(\mathfrak{R})$. In particular, $\boldsymbol{f} > 0$ if $\boldsymbol{f}(t) > 0$ ultimately.

Consequently, $H(\mathfrak{R})$ has a natural $\mathcal{L}$-structure, which we denote by $\mathfrak{H}(\mathfrak{R})$. It follows from [Cos00, Prop. 5.9] that $\mathfrak{H}(\mathfrak{R})$ and $\mathfrak{R}$ have the same full theory; see also [Fos10, Lemma 2.2.64]. In other words, the following holds.

Proposition 2.7. *Let $\mathfrak{R}$ be an o-minimal $\mathcal{L}$-structure and $\mathfrak{H}(\mathfrak{R})$ the natural $\mathcal{L}$-structure of the germs of functions that are definable in $\mathfrak{R}$. Then, for any $\mathcal{L}$-sentence ϕ, we have $\mathfrak{R} \models \phi$ if and only if $\mathfrak{H}(\mathfrak{R}) \models \phi$.*

As an expansion of $\bar{\mathbb{R}}$, the structure $\mathfrak{R}$ satisfies the axioms of the theory of real closed fields. By Proposition 2.7, the structure $\mathfrak{H}(\mathfrak{R})$ satisfies the same axioms. As a result,

$H(\mathfrak{R})$ is a real closed field that we will refer to as the *Hardy field* of structure $\mathfrak{R}$. In particular, $H(\mathfrak{R})$ is an ordered field, and it carries a natural topology induced by the ordering. The standard topology on $\mathbb{R}$ coincides with the subspace topology induced from $H(\mathfrak{R})$.

A structure $\mathfrak{R}$ is *polynomially bounded* if for any definable function $f : \mathbb{R} \to \mathbb{R}$, there exists a natural number n such that $|f(t)| \leq t^n$ ultimately. Miller proved [Mil94b] that in an o-minimal and polynomially bounded expansion of $\bar{\mathbb{R}}$, if a definable function f is not ultimately zero, then there exists an exponent $r \in \mathbb{R}$ and a non-zero coefficient $c \in \mathbb{R}$ such that

$$\lim_{t \to +\infty} \frac{f(t)}{t^r} = c \; . \tag{2.1}$$

The set of such exponents r forms a subfield of $\mathbb{R}$, called the *field of exponents* of the structure $\mathfrak{R}$.

It follows that the Archimedean value group of $\mathfrak{H}(\mathfrak{R})$ is the field of exponents of $\mathfrak{R}$, with its additive law. Moreover, the Archimedean valuation can be identified with the map:

$$\boldsymbol{f} \mapsto \lim_{t \to +\infty} \log(|\boldsymbol{f}(t)|)/\log(t) \; .$$

In the following, we will use the structure $\bar{\mathbb{R}}^{\mathbb{R}}$ which expands $\bar{\mathbb{R}}$ by adding to the language $\mathcal{L}_{\mathrm{or}}$ the family of function symbols $(f_r)_{r \in \mathbb{R}}$, and interpreting f_r as the power function that maps a positive number t to t^r, and any non-positive number to 0. The structure $\bar{\mathbb{R}}^{\mathbb{R}}$ is o-minimal, polynomially bounded and its field of exponents is $\mathbb{R}$; see [Mil94a, Mil12]. It follows that the value group of $H(\bar{\mathbb{R}}^{\mathbb{R}})$ is the additive group $(\mathbb{R}, <, +, 0)$. Another structure with the same properties is $\mathbb{R}_{\mathrm{an}*}$, the reals with restricted analytic functions and convergent generalized power series [vdDS98].

2.2.6 Maximal ordered groups and fields

Hahn series forms the more general ordered field we need to consider. Indeed, every ordered field can be embedded in $\mathbb{R}[\![t^G]\!]$ for some group G. Recall that an ordered group G is torsion-free, and thus there exists a unique (up to isomorphism) minimal divisible group that contains G, called the *divisible hull* of G; see [Mar02, Lemma 3.1.8].

Theorem 2.8 ([CD69, Theorem II])**.** *Let K be an ordered field with value group G. Then, there exists a value and order preserving isomorphism π of K into a subfield of $\mathbb{R}[\![t^G]\!]$.*

In addition, assume that $\bar{R} \subseteq K$ is any order-isomorphic copy of the reals and $r \mapsto \bar{r}$ is the unique order-isormorphism of $\mathbb{R}$ onto $\bar{R}$. Let $\Delta \subseteq G$ be a rationally independant basis for the divisible hull of G, and $(\boldsymbol{x}_\delta)_{\delta \in \Delta} \subseteq K_+$ a system of positive representatives of Δ. Then π can be chosen so that $\pi(\bar{r}\boldsymbol{x}_\delta) = rt^\delta$ for each $r \in \mathbb{R}$ and $\delta \in \Delta$.

Kaplansky proved that any valued field of characteric 0 can be embedded into a power series field, and his result extends to fields of positive characteric under some conditions [Kap42, Theorem 6]. So Theorem 2.8 is sometimes called “Kaplansky’s embedding Theorem” [DW96, Theorem 2.17]; see also [Poo93]. Conrad and Dauns in fact extended this result to lattice-ordered fields [CD69]; see also [Ste10].

The additive group of Hahn series is also the more general group we need.

Theorem 2.9 (Hahn’s embedding Theorem [Hah07])**.** *Every ordered group is order-isomorphic to a subgroup of the additive group of Hahn series* $\mathbb{R}[\![t^S]\!]$*, for some ordered set* S*.*

The notion of Archimedean classes extends to ordered groups, and the set S is in fact the set of Archimedean classes of G. Several proofs of this theorem are known, see the books [DW96, Fuc63, Ste10] and the references therein.

Chapter 3

Tropicalizing the simplex method

The simplex method is a family of algorithms that solve classical linear programs on an ordered field K, *i.e.*, problems of the form

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0, \ \boldsymbol{x} \in K^n \ , \end{array} \qquad \mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$$

where $\boldsymbol{A} \in K^{m\times n}$, $\boldsymbol{b} \in K^m$, and $\boldsymbol{c} \in K^n$. In this chapter, we give conditions under which a simplex algorithm can be implemented using only the "tropical" information $\operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$.

This is motivated by the tropical counterpart of linear programming developped in the subsequent chapters, but also by the question of the complexity of linear programming over ordered fields [Meg87], in particular over fields of rational functions [Jer73a, ER89, FAA02]. Intuitively, we will perform arithmetic operations over series expansions of rational functions using only the leading terms.

In a nutshell, to perform the basic operations of the simplex method, pivoting and computing the signs of reduced costs, it is sufficient to know the signs of some minors of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. Hence, to tropicalize, it is sufficient to determine the signs of the minors of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ using only $\operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. However, pivoting rules may be arbitrary procedures. In order to tropicalize, we restrict ourselves to *semi-algebraic* pivoting rules, *i.e.*, pivoting rules that have access to the problem at hand only through the signs of polynomials evaluated on $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. It turns out that, for a polynomial P, the sign of $P\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ can be computed from $\operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$, provided that $\operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ satisfies some genericity conditions. Moreover, under assumptions on P, that sign can be determined in time polynomial in the input size of $\operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$.

This chapter is organized as follows. We first recall basic notions and results on polyhedra and linear programming in Section 3.1. A key idea of this chapter is contained in Section 3.2, where we explain how the sign of a polynomial can be computed by tropical means. Section 3.3 discusses the simplex method in the context of ordered fields. Last, semi-algebraic pivoting rules and the tropical implementation of simplex algorithms are presented in Section 3.4.

In this chapter, we demonstrate the feasibility of the tropicalization of the simplex method. Even if the tropical versions of pivoting and computing reduced costs presented here run in polynomial time, their complexities can be improved. More efficient implementations are presented in Chapter 7.

The tropicalization of simplex operations (pivoting and computing reduced costs) was exposed in [ABGJ13b]. In [ABGJ13a], the tropicalization of *combinatorial* pivoting rules was presented (combinatorial pivoting rules rely only on the signs of the minors of the input matrix). The study of the tropical shadow-vertex rule [ABG14] led to the more general framework of *semi-algebraic* pivoting rules that we adopt here.

3.1 Polyhedra over ordered fields

In this section, we review relevant basic results on linear programming and polyhedra over ordered fields. Throughout this chapter, K denotes an arbitrary ordered field.

A *halfspace* in dimension $n \geq 1$ is a set of the form

$$\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a}, \boldsymbol{b}) := \{\boldsymbol{x} \in K^n \mid \boldsymbol{a}\boldsymbol{x} + \boldsymbol{b} \geq 0\} \tag{3.1}$$

where $\boldsymbol{a} \in K^{1\times n}$ and $\boldsymbol{b} \in K$. When $\boldsymbol{b} = 0$, it is said to be a *linear halfspace*. The boundary

$$\boldsymbol{\mathcal{H}}(\boldsymbol{a}, \boldsymbol{b}) := \{\boldsymbol{x} \in K^n \mid \boldsymbol{a}\boldsymbol{x} + \boldsymbol{b} = 0\} \tag{3.2}$$

of an halfpsace in a *hyperplane*. A *polyhedron* is the intersection of finitely many halfspaces, *i.e.*, a set of the form:

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

When $\boldsymbol{b}$ is the zero vector, $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is a *polyhedral cone*.

3.1.1 Convexity

A subset $\boldsymbol{P}$ of the K-vector space K^n is *convex* if, for any $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{P}$, the set $\boldsymbol{P}$ also contains the *convex hull* $\operatorname{conv}(\boldsymbol{x}, \boldsymbol{y})$ of $\boldsymbol{x}$ and $\boldsymbol{y}$, where:

$$\operatorname{conv}(\boldsymbol{x}, \boldsymbol{y}) := \{\boldsymbol{\lambda}\boldsymbol{x} + \boldsymbol{\mu}\boldsymbol{y} \mid \boldsymbol{\lambda}, \boldsymbol{\mu} \in K_+ \text{ and } \boldsymbol{\lambda} + \boldsymbol{\mu} = 1\} \ .$$

More generally, the convex hull $\operatorname{conv}(\boldsymbol{S})$ of an arbitrary subset $\boldsymbol{S} \subseteq K^n$ is the smallest convex subset containing $\boldsymbol{S}$. When $\boldsymbol{S}$ consists of finitely many points $\boldsymbol{x}^1, \dots, \boldsymbol{x}^k \in K^n$, we have:

$$\operatorname{conv}(\boldsymbol{x}^1, \dots, \boldsymbol{x}^k) := \left\{ \sum_{i=1}^{k} \boldsymbol{\lambda}_k \boldsymbol{x}^k \mid \boldsymbol{\lambda}_i \in K_+ \text{ for all } i \in [k], \text{ and } \sum_{i=1}^{k} \boldsymbol{\lambda}_i = 1 \right\} \ .$$

We say that a point $\boldsymbol{v}$ in a convex set $\boldsymbol{P}$ is an *extreme point*, or *vertex*, of $\boldsymbol{P}$ if it cannot be expressed as a convex combinaison of points in $\boldsymbol{P} \setminus \{\boldsymbol{v}\}$. This means that if $\boldsymbol{v}$ is contained in $\operatorname{conv}(\boldsymbol{x}, \boldsymbol{y})$ for some $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{P}$, then $\boldsymbol{v} = \boldsymbol{x}$ or $\boldsymbol{v} = \boldsymbol{y}$.

A *convex cone* in K^n is a convex set $\boldsymbol{C}$ that is also stable under multiplication by positive scalars, *i.e.*, for any $\boldsymbol{x} \in \boldsymbol{C}$ and any positive $\boldsymbol{\lambda} \in K$, the point $\boldsymbol{\lambda x}$ belongs to $\boldsymbol{C}$. Equivalently, $\boldsymbol{C}$ is a convex cone if and only if for any finite subset of points $\boldsymbol{x}^1, \dots, \boldsymbol{x}^k \in \boldsymbol{C}$, the set $\boldsymbol{C}$ contains their *positive hull* $\operatorname{pos}(\boldsymbol{x}^1, \dots, \boldsymbol{x}^k)$, where:

$$\operatorname{pos}(\boldsymbol{x}^1, \dots, \boldsymbol{x}^k) := \left\{ \sum_{i=1}^{k} \boldsymbol{\lambda}_k \boldsymbol{x}^k \mid \boldsymbol{\lambda}_i \in K_+ \text{ for all } i \in [k] \right\} .$$

As for convex hull, the positive hull $\operatorname{pos}(\boldsymbol{S})$ of an arbitrary subset $\boldsymbol{S} \subseteq K^n$ is the smallest convex cone containing S, or, equivalently, the union of $\operatorname{pos}(\boldsymbol{x}^1, \dots, \boldsymbol{x}^k)$ for all finite number of points $\boldsymbol{x}^1, \dots, \boldsymbol{x}^k \in \boldsymbol{S}$.

A *ray* of a convex cone $\boldsymbol{C}$ is a set of the form $[\boldsymbol{r}] := \{\boldsymbol{\lambda r} \mid \boldsymbol{\lambda} \in K_+\}$ for some non-null vector $\boldsymbol{r} \in \boldsymbol{C}$. We say that a ray $[\boldsymbol{r}]$ is an *extreme ray* of the convex cone $\boldsymbol{C}$ if $\boldsymbol{x} \in [\boldsymbol{r}]$ or $\boldsymbol{y} \in [\boldsymbol{r}]$ whenever $\boldsymbol{r} \in \operatorname{pos}(\boldsymbol{x}, \boldsymbol{y})$ for some $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{C}$.

The "unbounded" part of a convex set $\boldsymbol{P}$ is described by its *recession cone* $\operatorname{rec}(\boldsymbol{P})$, where:

$$\operatorname{rec}(\boldsymbol{P}) := \{\boldsymbol{r} \in K^n \mid \boldsymbol{\lambda r} + \boldsymbol{x} \in \boldsymbol{P} \text{ for all } \boldsymbol{x} \in \boldsymbol{P} \text{ and all } \boldsymbol{\lambda} \in K_+\} .$$

When $\boldsymbol{P}$ is a non-empty polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, its recession cone coincides with the polyhedral cone $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, 0)$. By extension, we say that $[\boldsymbol{r}]$ is an *extreme ray* of the convex set $\boldsymbol{P}$ if $[\boldsymbol{r}]$ is an extreme ray of its recession cone $\operatorname{rec}(\boldsymbol{P})$.

3.1.2 Double description

In the remaining of this section, we shall prove the following fundamental theorem.

Theorem 3.1. *A polyhedron on an ordered field is the convex hull of finitely many point and rays. More precisely, for any $\boldsymbol{A} \in K^{m \times n}$ and $\boldsymbol{b} \in K^m$, there exists two finite sets $\boldsymbol{V}, \boldsymbol{R} \subseteq K^n$ such that:*

$$\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R}) , \tag{3.3}$$

where $\operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R}) := \{\boldsymbol{x} + \boldsymbol{y} \mid \boldsymbol{x} \in \operatorname{conv}(\boldsymbol{V}), \boldsymbol{y} \in \operatorname{pos}(\boldsymbol{R})\}$ is the Minkowksi sum of these two sets.

We shall derive Theorem 3.1 as a corollary of a slightly more precise statement for bounded polyhedra.

Theorem 3.2 ([CC58])**.** *A bounded polyhedron on an ordered field is the convex hull of a finite set of points, the set of its extreme points.*

Corollary 3.3. *If $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is included in the non-negative orthant of K^n, then*

$$\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R}) , \tag{3.4}$$

where $\boldsymbol{V}$ is the set of extreme points of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, and $\boldsymbol{R}$ is the set of its extreme rays.

Proof. Let us write $\boldsymbol{\mathcal{P}} := \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. We homogeneize $\boldsymbol{\mathcal{P}}$ into the polyhedral cone

$$\boldsymbol{\mathcal{C}} = \{(\boldsymbol{x}, \boldsymbol{\lambda}) \in K^n \times K \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}\boldsymbol{\lambda} \geq 0,\ \boldsymbol{\lambda} \geq 0\} \ . \tag{3.5}$$

Clearly, if $\boldsymbol{x}$ belongs to the polyhedron $\boldsymbol{\mathcal{P}}$, then $(\boldsymbol{x}, 1) \in \boldsymbol{\mathcal{C}}$. An conversely, if $(\boldsymbol{x}, \boldsymbol{\lambda})$ is a non-null vector contained in the cone $\boldsymbol{\mathcal{C}}$, then either $\boldsymbol{\lambda} = 0$ and $\boldsymbol{x}$ is in the recession cone of $\boldsymbol{\mathcal{P}}$, or $\boldsymbol{\lambda}^{-1}\boldsymbol{x} \in \boldsymbol{\mathcal{P}}$. Observe that an extreme point $\boldsymbol{x}$ of $\boldsymbol{\mathcal{P}}$ yields an extreme ray $[(\boldsymbol{x}, 1)]$ of $\boldsymbol{\mathcal{C}}$, and an extreme ray $[\boldsymbol{r}]$ of the recession cone of $\boldsymbol{\mathcal{P}}$ provides an extreme ray $[(\boldsymbol{r}, 0)]$ of $\boldsymbol{\mathcal{C}}$. Futhermore, any extreme ray of $\boldsymbol{\mathcal{C}}$ arises in this way.

Since $\boldsymbol{\mathcal{P}}$ is included in the non-negative orthant of K^n, its homogeneization $\boldsymbol{\mathcal{C}}$ is included in the non-negative orthant of K^{n+1}. It follows that the set of rays of $\boldsymbol{\mathcal{C}}$ can be identified with the bounded polyhedron obtained by intersecting $\boldsymbol{\mathcal{C}}$ with the hyperplane $\{(\boldsymbol{x}, \boldsymbol{\lambda}) \in K^{n+1} \mid \sum_{j=1}^{n} \boldsymbol{x}_j + \boldsymbol{\lambda} = 1\}$. By Theorem 3.2, this bounded polyhedron is the convex hull of its set $\boldsymbol{G}$ of extreme points. Consequently, the cone $\boldsymbol{\mathcal{C}}$ is the conic hull of $\boldsymbol{G}$, and it is easy to see that $\boldsymbol{G}$ consists of a representative of each extreme ray of $\boldsymbol{\mathcal{C}}$. It follows that $\boldsymbol{\mathcal{P}} = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R})$ where $\boldsymbol{V} = \{\boldsymbol{x} \in K^n \mid (\boldsymbol{x}, 0) \in \boldsymbol{G}\}$ is the set of extreme points of $\boldsymbol{\mathcal{P}}$, and $\boldsymbol{R} = \{\boldsymbol{r} \mid (\boldsymbol{r}, \boldsymbol{\lambda}) \in \boldsymbol{G} \text{ for some } \boldsymbol{\lambda} \neq 0\}$ is the set of extreme rays of $\boldsymbol{\mathcal{P}}$. □

An arbitrary polyhedron may not have extreme points or extreme rays, but it is still finitely generated.

Proof of Theorem 3.1. Observe that an arbitrary polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \subseteq K^n$ is a projection of the polyhedron

$$\{(\boldsymbol{x}^+, \boldsymbol{x}^-) \in K^n \times K^n \mid \boldsymbol{A}\boldsymbol{x}^+ - \boldsymbol{A}\boldsymbol{x}^- + \boldsymbol{b} \geq 0, \boldsymbol{x}^+ \geq 0, \boldsymbol{x}^- \geq 0\} \ ,$$

which is included in the positive orthant, and that this projection preserves convex and conic hulls. The result then follows from Corollary 3.3. □

3.1.3 Classical linear programming

A *linear program* is an optimization problem of the form:

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0,\ \boldsymbol{x} \in K^n \ , \end{array} \tag*{$\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$}$$

where $\boldsymbol{A} \in K^{m \times n}$, $\boldsymbol{b} \in K^m$, and $\boldsymbol{c} \in K^n$. Thus a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ seeks a minimizer of a linear function $\boldsymbol{x} \mapsto \boldsymbol{c}^\top \boldsymbol{x}$ over a polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. When $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is empty, we say that $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is *infeasible*. A linear program is *unbounded* if, for any $\boldsymbol{\nu} \in K$, there exists a feasible point $\boldsymbol{x}$ such that $\boldsymbol{c}^\top \boldsymbol{x} < \boldsymbol{\nu}$. An *optimal solution* of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is a $\boldsymbol{x}^* \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ such that $\boldsymbol{c}^\top \boldsymbol{x}^* \leq \boldsymbol{c}^\top \boldsymbol{x}$ for all $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. If an optimal solution exists, $\boldsymbol{c}^\top \boldsymbol{x}^*$ is called the *optimal value* of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$.

Proposition 3.4. *A linear program* $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ *over an ordered field* K *is either infeasible, unbounded, or admits an optimal solution.*

Proof. If the linear program is infeasible, then the other two possibilities are excluded. Now suppose that $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ is not empty. Then, by Theorem 3.1, there exists two finite sets $\boldsymbol{V},\boldsymbol{R} \subseteq K^n$ such that

$$\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}) = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R}) \ .$$

If there exists a $\boldsymbol{r} \in \boldsymbol{R}$ such that $\boldsymbol{c}^\top \boldsymbol{r} < 0$, then the linear program is unbounded. Otherwise any feasible point $\boldsymbol{x}$ satisfy $\boldsymbol{c}^\top \boldsymbol{x} \geq \min_{\boldsymbol{v}\in\boldsymbol{V}} \boldsymbol{c}^\top \boldsymbol{v}$, and an element of V is an optimal solution. □

Unboundedness can be certified as follows.

Lemma 3.5. *A feasible linear program* $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *is unbounded if and only if there exists a* $\boldsymbol{r}$ *in the polyhedral cone* $\boldsymbol{\mathcal{P}}(\boldsymbol{A},0)$ *such that* $\boldsymbol{c}^\top \boldsymbol{r} < 0$*.*

Proof. Continuing the previous proof, the feasible linear program is unbounded if and only if there exists $\boldsymbol{r} \in \boldsymbol{R}$ such that $\boldsymbol{c}^\top \boldsymbol{r} < 0$, where $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}) = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R})$. Then, observe that $\operatorname{pos}(\boldsymbol{R})$ is the recession cone of $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$, which is exactly the polyhedral cone $\boldsymbol{\mathcal{P}}(\boldsymbol{A},0)$. □

Duality

The *dual linear program* of $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ is :

$$\begin{aligned} &\text{maximize} && -\boldsymbol{b}^\top \boldsymbol{y} \\ &\text{subject to} && \boldsymbol{A}^\top \boldsymbol{y} = \boldsymbol{c},\ \boldsymbol{y} \geq 0,\ \boldsymbol{y} \in K^m \ . \end{aligned} \tag{$\mathbf{LD}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$}$$

Theorem 3.6. *Let* $\boldsymbol{x}$ *be a feasible solution of the linear program* $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *and* $\boldsymbol{y}$ *a feasible solution of the dual problem* $\mathbf{LD}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$*. Then,* $\boldsymbol{c}^\top \boldsymbol{x} \geq -\boldsymbol{b}^\top \boldsymbol{y}$*.*

Proof. Since $\boldsymbol{y}$ is dual feasible, we have $\boldsymbol{c}^\top = \boldsymbol{y}^\top \boldsymbol{A}$. Hence, $\boldsymbol{c}^\top \boldsymbol{x} = \boldsymbol{y}^\top \boldsymbol{A}\boldsymbol{x}$ and $\boldsymbol{c}^\top \boldsymbol{x} + \boldsymbol{b}^\top \boldsymbol{y} = \boldsymbol{y}^\top(\boldsymbol{A}\boldsymbol{x}+\boldsymbol{b})$. Since $\boldsymbol{y} \geq 0$ and $\boldsymbol{A}\boldsymbol{x}+\boldsymbol{b} \geq 0$, it follows that $\boldsymbol{c}^\top \boldsymbol{x} + \boldsymbol{b}^\top \boldsymbol{y} \geq 0$. □

Theorem 3.7 (Complementary Slackness)**.** *Let* $\boldsymbol{x}^*$ *be a feasible solution of the linear program* $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *and* $\boldsymbol{y}^*$ *a feasible solution of the dual problem* $\mathbf{LD}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *such that:*

$$\boldsymbol{y}_i^*(\boldsymbol{A}_i\boldsymbol{x}^* + \boldsymbol{b}_i) = 0 \ \textit{ for all } i \in [m] \ . \tag{3.6}$$

Then, $\boldsymbol{x}^*$ *and* $\boldsymbol{y}^*$ *are optimal solutions of* $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *and* $\mathbf{LD}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *respectively. Moreover,* $\boldsymbol{c}^\top \boldsymbol{x}^* = -\boldsymbol{b}^\top \boldsymbol{y}^*$*.*

Proof. By Weak Duality (Theorem 3.6), $-\boldsymbol{b}^\top \boldsymbol{y}^*$ is a lower bound for the optimal value of $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$, and $\boldsymbol{c}^\top \boldsymbol{x}^*$ is an upper bound for the optimal value of $\mathbf{LD}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$. Hence, it is sufficient to prove the equality $\boldsymbol{c}^\top \boldsymbol{x}^* = -\boldsymbol{b}^\top \boldsymbol{y}^*$. As in the proof of Theorem 3.6, we have $\boldsymbol{c}^\top \boldsymbol{x}^* + \boldsymbol{b}^\top \boldsymbol{y}^* = (\boldsymbol{y}^*)^\top(\boldsymbol{A}\boldsymbol{x}^*+\boldsymbol{b})$. Then, the conditions (3.6) imply that $(\boldsymbol{y}^*)^\top(\boldsymbol{A}\boldsymbol{x}^*+\boldsymbol{b}) = 0$. □

3.2 Computing the sign of a polynomial by tropical means

A key ingredient to tropicalize algorithms is to determine the sign of a polynomial expression on an ordered field K using only the information provided by the valuation map. Given a polynomial $P \in \mathbb{Q}[X_1, \dots, X_l]$, we show that under genericity conditions on $\boldsymbol{\delta} \in K^l$, the sign of $P(\boldsymbol{\delta})$ can be computed using only $\mathrm{sval}(\boldsymbol{\delta})$. More precisely, to compute the sign of $P(\boldsymbol{\delta})$, we solve a linear program over the Newton polytope of the polynomial P. The objective function of the linear program is given by $\mathrm{sval}(\boldsymbol{\delta})$. Hence, if we have an algorithm that solves linear programs over the Newton polytope of P in polynomial time, the sign of $P(\boldsymbol{\delta})$ can be computed in time polynomial in the input size of $\mathrm{sval}(\boldsymbol{\delta})$.

We shall write a multivariate polynomial $P \in \mathbb{Q}[X_1, \dots, X_l]$ as the formal expression:

$$P = \sum_{\alpha \in \Lambda} q_\alpha X^\alpha \ ,$$

where $\Lambda \subseteq \mathbb{N}^l$ is a finite set, the coefficients $q_\alpha \neq 0$ are rationals numbers, and $X^\alpha = \prod_{i=1}^{l} X_i^{\alpha_i}$.

3.2.1 Tropicalization of polynomials

The *tropicalization* of a polynomial $P \in \mathbb{Q}[X_1, \dots, X_l]$ is the formal tropical expression:

$$\mathrm{trop}(P) := \bigoplus_{\alpha \in \Lambda} \mathrm{tsign}(q_\alpha) \odot \delta_1^{\odot \alpha_1} \odot \cdots \odot \delta_l^{\odot \alpha_l} \ , \tag{3.7}$$

where $\mathrm{tsign}(q_\alpha) = \mathbb{1}$ if $q_\alpha > 0$, and $\mathrm{tsign}(q_\alpha) = \ominus\mathbb{1}$ if $q_\alpha < 0$. A tropical vector $\delta \in \mathbb{T}_\pm^l$ is *generic* for the polynomial P if the maximum in

$$\bigoplus_{\alpha \in \Lambda} |\delta|_1^{\odot \alpha_1} \odot \cdots \odot |\delta|_l^{\odot \alpha_l} = \max_{\alpha \in \Lambda} \ \alpha_1 |\delta|_1 + \alpha_2 |\delta|_2 + \cdots + \alpha_l |\delta|_l \ , \tag{3.8}$$

is equal to $\mathbb{0}$ or attained on a unique $\alpha^* \in \Lambda$.

We also say that $\delta \in \mathbb{T}_\pm^l$ is *sign-generic* for P if, for any two $\alpha, \beta \in \Lambda$ attaining the maximum in (3.8), the terms $\mathrm{tsign}(q_\alpha) \odot \delta_1^{\odot \alpha_1} \odot \cdots \odot \delta_l^{\odot \alpha_l}$ and $\mathrm{tsign}(q_\beta) \odot \delta_1^{\odot \beta_1} \odot \cdots \odot \delta_l^{\odot \beta_l}$ have the same tropical sign. When $\delta \in \mathbb{T}_\pm^l$ is generic, or sign-generic, for P, we write:

$$\mathrm{trop}(P)(\delta) := \mathrm{tsign}(q_{\alpha^*}) \odot \delta_1^{\odot \alpha_1^*} \odot \cdots \odot \delta_l^{\odot \alpha_l^*} \tag{3.9}$$

where α^* is any maximizer in (3.8). Observe that if δ is generic for P, then it is sign-generic. Also notice that the modulus of $\mathrm{trop}(P)(\delta)$ is equal to (3.8). We say that $\alpha \in \Lambda$ is a *maximizer* for $|\mathrm{trop}(P)(\delta)|$ if it attains the maximum in (3.8).

The determinant is a polynomial that plays an important role in this manuscript. The tropicalization of the determinant of a square matrix $M \in \mathbb{T}_\pm^{n \times n}$ will be denoted by $\mathrm{tdet}(M)$. It is defined by:

$$\mathrm{tdet}(M) := \bigoplus_{\sigma \in \mathrm{Sym}(n)} \mathrm{tsign}(\sigma) \odot M_{1\sigma(1)} \odot \cdots \odot M_{n\sigma(n)} \ ,$$

where $\mathrm{Sym}(n)$ is the set of all permutations of $[n]$, and $\mathrm{tsign}(\sigma) = \mathbb{1}$ if σ is an even permutation and $\mathrm{tsign}(\sigma) = \ominus\mathbb{1}$ otherwise. Observe that

$$|\,\mathrm{tdet}(M)| = \max_{\sigma \in \mathrm{Sym}(n)} |M_{1\sigma(1)}| + \cdots + |M_{n\sigma(n)}| \ . \tag{3.10}$$

Computing a maximizer for $|\,\mathrm{tdet}(M)|$ amounts to finding a permutation which attains the maximum in (3.10). Such a permutation is a solution of the assignment problem with costs $(|M_{ij}|)$. It can be found in time $O(n^3)$ using the Hungarian method; see [Sch03, §17.3].

Lemma 3.8. *Consider a polynomial* $P = \sum_{\alpha \in \Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1, \dots, X_l]$ *and* $\boldsymbol{\delta} \in \mathbb{K}^l$. *Suppose that* $\delta = \mathrm{sval}(\boldsymbol{\delta})$ *is sign-generic for the polynomial* P*, then*

$$\mathit{trop}(P)(\delta) = \mathrm{sval}(P(\boldsymbol{\delta})) \ .$$

Proof. First one easily checks that if two elements $\boldsymbol{x}, \boldsymbol{y} \in K$ have the same value and the same sign, then $\mathrm{val}(\boldsymbol{x} + \boldsymbol{y}) = \max(\mathrm{val}(\boldsymbol{x}), \mathrm{val}(\boldsymbol{y}))$ and $\boldsymbol{x} + \boldsymbol{y}$ has the same sign as $\boldsymbol{x}$ and $\boldsymbol{y}$. Similarly, if $\mathrm{val}(\boldsymbol{x}) > \mathrm{val}(\boldsymbol{y})$, then we have $\mathrm{val}(\boldsymbol{x} + \boldsymbol{y}) = \mathrm{val}(\boldsymbol{x})$ and $\boldsymbol{x} + \boldsymbol{y}$ has the same sign as $\boldsymbol{x}$.

Let Λ^* be the set of maximizer for $|\mathrm{trop}(P)(\delta)|$. For any $\alpha^* \in \Lambda^*$, the image under the signed valuation map of the monomial $q_{\alpha^*} \prod_i \boldsymbol{\delta}_i^{\alpha_i^*}$ is $\mathrm{trop}(P)(\delta)$. Consequently, the signed value of $\sum_{\alpha^* \in \Lambda^*} q_{\alpha^*} \prod_i \boldsymbol{\delta}_i^{\alpha_i^*}$ is also $\mathrm{trop}(P)(\delta)$. For every $\alpha \in \Lambda \setminus \Lambda^*$ the monomial $q_\alpha \prod_i \boldsymbol{\delta}_i^{\alpha_i}$ has a value strictly smaller than $|\mathrm{trop}(P)(\delta)|$. Hence, the signed value of $P(\boldsymbol{\delta})$ is $\mathrm{trop}(P)(\delta)$. ☐

When $\delta \in \mathbb{T}_\pm^l$ is (sign-)generic for a polynomial P, computing $\mathrm{trop}(P)(\delta)$ amounts to finding a maximizer for $|\mathrm{trop}(P)(\delta)|$. It turns out that such a maximizer is an optimal vertex of an (abstract) linear program over the polytope $\mathrm{conv}(\Lambda)$, the *Newton polytope* of P.

To see this, let us first suppose that δ does not have $\mathbb{0}$ entries. In that case, $\mathrm{trop}(P)(\delta)$ is not equal to $\mathbb{0}$. Moreover, $|\mathrm{trop}(P)(\delta)|$ is the maximum of the linear function $\alpha \mapsto \sum_{i \in [l]} \alpha_i |\delta|_i$. evaluated on the finite number of points $\alpha \in \Lambda$. By convexity, $|\mathrm{trop}(P)(\delta)|$ is the optimal value of the following optimization problem:

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in [l]} \alpha_i |\delta|_i \\ \text{subject to} \quad & \alpha \in \mathrm{conv}(\Lambda). \end{aligned} \tag{3.11}$$

Hence, the set of maximizers for $|\mathrm{trop}(P)(\delta)|$ is exactly the set of optimal vertices of the linear program (3.11). Observe that the feasible set of (3.12) is included in $\mathbb{R}^l$, while the objective function takes values in the totally ordered abelian group $G = \mathbb{T} \setminus \{\mathbb{0}\}$. Hence, the problem (3.12) is a linear program on $\mathbb{R}^l$ with an abstract linear objective function. Now, if $\delta \in \mathbb{T}_\pm^l$ has some entries equal to $\mathbb{0}$, a small technical difficulty arises.

Lemma 3.9. *Let $\mathbb{T} = \mathbb{T}(G)$. Consider a polynomial $P = \sum_{\alpha\in\Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1,\dots,X_l]$ and suppose that $\delta \in \mathbb{T}_{\pm}^l$. Then, a maximizer for $|trop(P)(\delta)|$ is given by an optimal vertex of the problem:*

$$\begin{aligned} \textit{maximize} \quad & \sum_{i\in[l]} \alpha_i d_i \\ \textit{subject to} \quad & \alpha \in \operatorname{conv}(\Lambda), \end{aligned} \tag{3.12}$$

where d is the vector with entries in the additive group $\mathbb{Q}\times G$, ordered lexicographically, defined by $d_i = (-1, 0)$ if $\delta_i = \mathbb{0}$ and $d_i = (0, |\delta|_i)$ otherwise.

Proof. If δ has no $\mathbb{0}$ entries, then the problems (3.11) and (3.12) have the same optimal solutions. Otherwise, if δ has $\mathbb{0}$ entries, it may happen that $\operatorname{trop}(P)(\delta) = \mathbb{0}$. This is the case if and only if, for all $\alpha \in \Lambda$, there exists an $i \in [l]$ with $\delta_i = \mathbb{0}$ and $\alpha_i > 0$. Consequently, $\operatorname{trop}(P)(\delta) = \mathbb{0}$ if and only if the optimal value of

$$\begin{aligned} \text{maximize} \quad & \sum_{i\in[l]\mid\delta_i=\mathbb{0}} \alpha_i \\ \text{subject to} \quad & \alpha \in \operatorname{conv}(\Lambda), \end{aligned} \tag{3.13}$$

is strictly greater than 0. In this case, every $\alpha \in \Lambda$ is a maximizer for $|\operatorname{trop}(P)(\delta)|$. This holds in particular for an optimal vertex of (3.12)

If $\operatorname{trop}(P)(\delta) \neq \mathbb{0}$, then the optimal value of (3.13) is equal to 0, and $|\operatorname{trop}(P)(\delta)|$ is the optimal value of:

$$\begin{aligned} \text{maximize} \quad & \sum_{i\in[l]\mid\delta_i\neq\mathbb{0}} \alpha_i \\ \text{subject to} \quad & \alpha \in \operatorname{conv}(\Lambda) \\ & \sum_{i\in[l]\mid\delta_i=\mathbb{0}} \alpha_i = 0\ . \end{aligned} \tag{3.14}$$

Furthermore, any optimal vertex of (3.14) is a maximizer for $|\operatorname{trop}(P)(\delta)|$. Observe that (3.14) and (3.12) have the same set of optimal solutions. $\square$

3.2.2 Tropically tractable polynomials

We say that a polynomial $P = \sum_{\alpha\in\Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1,\dots,X_l]$ is *tropically tractable* if there is an algorithm that, given any $\delta \in \mathbb{T}_{\pm}^l$ that is sign-generic for P, returns the the sign of $\operatorname{trop}(P)(\delta)$ in time polynomial in the input size $\langle\delta\rangle$ of δ.

The *(binary) input size* of an integer $z \in \mathbb{Z}$ is the number of bit required to write z in the binary representation. When $z = 0$ only one bit is needed. Otherwise, we need one bit for the sign and $\lceil\log_2(|z|+1)\rceil$ bits for the absolute value $|z|$, hence $\langle\delta\rangle := \lceil\log_2(|z|+1)\rceil + 1$. The input size of a rational number r, which can always be written as $r = p/q$ where p and $q > 0$ are relatively prime integers, is $\langle r\rangle = \langle p\rangle + \langle q\rangle$. The input size of a matrix is the sum of the input sizes of its entries. In particular, the input size of a vector $v \in \mathbb{Q}^l$ is always greater than l.

The notion of input size is a-priori not well defined for elements of an arbitrary group G. Since this is sufficient for our purposes, we shall study the tropical tractability of polynomials over tropical semirings of the form $\mathbb{T}(\mathbb{Q}^r)$, where r is an integer and $\mathbb{Q}^r$ is equipped with component-wise addition and lexicographic order. Note that Hahn's embedding theorem (Theorem 2.9) states that any totally ordered abelian group G is order-isomorphic to an additive subgroup of $\mathbb{R}^{|S|}$ equipped with a lexicographic order, where S is a suitable (possibly infinite) ordered set. Hence, G contains a subgroup which can be identified with a subgroup of $\mathbb{Q}^{|S|}$. The notion of input size is then well-defined for the elements of $\mathbb{Q}^{|S|}$ with a finite number of non-zero components. This of course depends on the embedding into $\mathbb{Q}^{|S|}$, which may not be known a-priori. Here, we assume that such an embedding is known.

We now give sufficient conditions on a polynomial to be tropically tractable.

Proposition 3.10. *Consider a polynomial $P = \sum_{\alpha\in\Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1,\dots,X_l]$ that satisfies the following properties:*

(i) there exists an algorithm which computes $\operatorname{sign}(q_\alpha)$*, for every $\alpha \in \Lambda$, in time polynomial in l;*

(ii) the Newton polytope $\operatorname{conv}(\Lambda)$ *is contained in a L_∞-ball of radius R, where the input size of R is polynomial in l;*

(iii) there exists an algorithm, which given any $\eta \in \mathbb{Q}^l$, returns an optimal vertex of the linear program

$$\begin{array}{ll} \textit{maximize} & \eta^\top \alpha \\ \textit{subject to} & \alpha \in \operatorname{conv}(\Lambda), \end{array} \tag{3.15}$$

in time polynomial in $\langle \eta \rangle$.

Then P is tropically tractable.

Proof. Let $\mathbb{T} = \mathbb{T}(\mathbb{Q}^{r-1})$ for a finite $r > 1$ and $\delta \in \mathbb{T}_\pm^l$ be sign-generic for $\operatorname{trop}(P)$. By Lemmas 3.8 and 3.9, it is sufficient to find an optimal vertex of the problem (3.12), *i.e.*, a maximizer of the linear function $\alpha \mapsto \sum_{i\in[l]} \alpha_i d_i$ which takes values in the lexicographically ordered group $\mathbb{Q}^r$. We shall use instead a real-valued linear objective function, $\alpha \mapsto \eta^\top \alpha$ for some $\eta \in \mathbb{Q}^l$ with an input size bounded by l and $\langle d \rangle$, that provides the same set of optimal solutions.

Note that we are interested in optimal vertices of $\operatorname{conv}(\Lambda)$, hence of elements of Λ. Thus, it is sufficient to find a η such that $\alpha \mapsto \sum_{i\in[l]} \alpha_i d_i$ and $\alpha \mapsto \eta^\top \alpha$ have the same maximizers in Λ. Let us denote $d_i = (d_{ij})_{j\in[r]} \in \mathbb{Q}^r$ for any $i \in [l]$.

Up to multiplying d by the common denominators of the $(d_{ij})_{ij}$, we can assume that the d_{ij} are integers (note that this transformation does not change the sum of the input sizes of the d_{ij}). By assumption (ii), there a exists an integer $R' \geq 1$, whose input size in bounded by a polynomial in l and $\langle \delta \rangle$, such that

$$-R' < \sum_{i\in[l]} \alpha_i \delta_{ij} < R' \tag{3.16}$$

for all points $\alpha \in \Lambda$ and all $j \in [r]$.

The objective vector $\eta \in \mathbb{Q}^l$, defined by $\eta_i = \sum_{j\in[1+r]} d_{ij}(2R')^{-j}$ for all $i \in [l]$, satisfies the required properties. Indeed, $\eta^\top \alpha \geq \eta^\top \beta$ for some $\alpha, \beta \in \Lambda$ if and only if

$$\sum_{j\in[r]} \left(R' + \sum_{i\in[l]} \alpha_i d_{ij} \right) (2R')^{-j} \geq \sum_{j\in[r]} \left(R' + \sum_{i\in[l]} \beta_i d_{ij} \right) (2R')^{-j} . \tag{3.17}$$

By (3.16), the numbers $R' + \sum_{i\in[l]} \alpha_i d_{ij}$ and $R' + \sum_{i\in[l]} \beta_i d_{ij}$ are positive integers strictly smaller than $2R'$. Hence, the left and right-hand side of (3.17) can be thought of as expansions of rationals in base $2R'$. If follows that the inequality (3.17) holds if and only if

$$\left(\sum_{i\in[l]} \alpha_i d_{i1}, \ \ldots, \ \sum_{i\in[l]} \alpha_i d_{ir} \right) \geq_{\text{lex}} \left(\sum_{i\in[l]} \beta_i \delta_{i1}, \ \ldots, \ \sum_{i\in[l]} \beta_i \delta_{ir} \right) . \qquad \square$$

Lemma 3.11. *A determinant is a tropically tractable polynomial. More precisely, given a $M \in \mathbb{T}_\pm^{n\times n}$ which is sign-generic for the $n \times n$ determinant polynomial, the sign of* $\operatorname{tdet}(M)$ *can be computed in $O(n^3)$ operations an in space polynomial in $\langle M \rangle$.*

Proof. This is a consequence of Proposition 3.10. The determinant of a $n \times n$ matrix is the polynomial of $\mathbb{Q}[X_{11}, \ldots, X_{nn}]$ defined by

$$\det = \sum_{\sigma\in S([n])} \operatorname{sign}(\sigma) \prod_{i\in[n]} X_{i\sigma(i)} .$$

A permutation $\sigma \in S([n])$ corresponds to the vector of exponents $\alpha_{ij} \in \mathbb{N}^{n\times n}$ defined for all $i \in [n]$ by $\alpha_{i\sigma(i)} = 1$ and $\alpha_{ij} = 0$ for $j \neq \sigma(i)$. Hence, the Newton polytope of the $n \times n$ determinant is the Birkhoff polytope: its vertices are in bijection with the perfect matchings of the complete bipartite graph between two sets of nodes of cardinality n. This polytope is contained in the L_∞-ball of radius 1 centered at the origin. Hence, Proposition 3.10 (ii) is satisfied.

The sign of a permutation $\sigma \in S([n])$ can be computed in $O(n)$ operations by counting the number of transpositions. Consequently, Proposition 3.10 (i) holds.

Finally, a linear program over the Birkhoff polytope is a maximal assigment problem. It can be solved in strongly polynomial time (in fact in $O(n^3)$ operations) by the Hungarian method; see [Sch03, Theorem 17.3]. Thus, Proposition 3.10 (iii) is satisfied. $\square$

A *separation oracle* for a convex set $C \subseteq \mathbb{R}^l$ is a routine which, given $\alpha \in \mathbb{R}^l$ decides whether $\alpha \in C$, and if not, returns a hyperplane that separates α from C, *i.e.*, finds a vector $d \in \mathbb{R}^l$ such that $d^\top \alpha > \max\{d^\top \beta \mid \beta \in C\}$.

Proposition 3.12. *Consider a polynomial $P = \sum_{\alpha\in\Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1, \ldots, X_l]$ that satisfies Conditions* (ii) *and* (i) *of Proposition 3.10, and such that there exists a polynomial-time separation oracle for the Newton polytope* $\operatorname{conv}(\Lambda)$*. Then P is tropically tractable.*

Proof. By the results of Grötschel, Lovász and Schrijver (Theorem 6.6.5 and Lemma 4.2.7 in [GLS88]), an optimal vertex of the linear program (3.15) can be found in $O(l)$ operations, and in space polynomial in l and $\langle \eta \rangle$, if there exists a polynomial-time separation oracle for the Newton polytope $\mathrm{conv}(\Lambda)$, and the vertices of $\mathrm{conv}(\Lambda)$ have an input size bounded by a polynomial in l. □

3.3 The simplex method

In this section, we recall the basic notions needed to present the simplex method.

Basic points

A *basis* of a polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ a subset $I \subseteq [m]$ of cardinality n such that the submatrix $\boldsymbol{A}_I$, formed from the rows with indices in I, is non-singular. The system

$$\bigcap_{i \in I} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i) = \{\boldsymbol{x} \in K^n \mid \boldsymbol{A}_I \boldsymbol{x} + \boldsymbol{b}_I = 0\} \tag{3.18}$$

contains a unique point, called a *basic point* and denoted 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*, and we say that I is a *feasible basis*. By extension, we say that I is a (feasible) basis of a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ if it is a (feasible) basis for its feasible set $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$.

Remark 3.13. A basis is sometimes defined by a partition of the (explicitly bounded) variables $(\boldsymbol{w}_1, \dots, \boldsymbol{w}_m)$ in "basic" and "non-basic" variables, where $\boldsymbol{w} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}$. Observe that I corresponds to the "non-basic" variables as it indexes the zero coordinates of $\boldsymbol{w}$. The set I can also be interpreted as the set of "basic" variables in the dual program.

Basic points are the "algebraic" counterpart of the geometric notion of extreme points.

Proposition 3.14. *Each feasible basic point of a polyhedron is an extreme point. Conversely, each extreme point is a basic point for some feasible basis.*

Proof. Let $\boldsymbol{x}^I$ be a basic point for some basis I. Suppose by contradiction that $\boldsymbol{x}^I$ is not an extreme point of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Then $\boldsymbol{x}^I = \boldsymbol{\lambda}\boldsymbol{y} + (1 - \boldsymbol{\lambda})\boldsymbol{z}$ for some $\boldsymbol{y}, \boldsymbol{z} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ and $0 < \boldsymbol{\lambda} < 1$. As $\boldsymbol{y} \neq \boldsymbol{x}^I$, we have $\boldsymbol{A}_i \boldsymbol{y} + \boldsymbol{b}_i > 0$ for some $i \in I$, otherwise $\boldsymbol{y}$ would be a solution of the system (3.18). Since $(1 - \boldsymbol{\lambda})(\boldsymbol{A}_i \boldsymbol{z} + \boldsymbol{b}_i) = -\lambda(\boldsymbol{A}_i \boldsymbol{y} + \boldsymbol{b}_i) < 0$ and $(1 - \boldsymbol{\lambda}) > 0$, we deduce that $\boldsymbol{A}_i \boldsymbol{z} + \boldsymbol{b}_i < 0$, and thus that $\boldsymbol{z} \notin \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, a contradiction.

Conversely, consider an extreme point $\boldsymbol{x}$ of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Let $I = \{i \in [m] \mid \boldsymbol{A}_i \boldsymbol{x} + \boldsymbol{b}_i = 0\}$. If $\boldsymbol{A}_I$ has a rank smaller than the dimension n, then there exists a vector $\boldsymbol{d} \neq 0_{K^n}$ in the kernel of $\boldsymbol{A}_I$. Hence, for $\boldsymbol{\lambda} > 0$ small enough, the points $\boldsymbol{x} + \boldsymbol{\lambda}\boldsymbol{d}$ and $\boldsymbol{x} - \boldsymbol{\lambda}\boldsymbol{d}$ belongs to $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Hence, $\boldsymbol{x}$ is in the convex hull of two points of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ that are distincts from $\boldsymbol{x}$. Consequently $\boldsymbol{A}_I$ has rank at least n, hence it contains a $n \times n$ submatrix $\boldsymbol{A}_{I'}$ with $\det(\boldsymbol{A}_{I'}) \neq 0$. □

Note however that two distinct bases I, I' can yield the same basic point. This will not happen under the non-degeneracy assumption explained below.

Given a basis, the corresponding basic point can be computed with Cramer's formulæ.

Proposition 3.15 (Cramer's formulæ)**.** *Let I be a basis of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. The components of the basic point $\boldsymbol{x}^I \in K^n$ are given by:*

$$\boldsymbol{x}^I_j = (-1)^{n+1+j} \det(\boldsymbol{A}_{I,\widehat{j}}\ \boldsymbol{b}_I) / \det(\boldsymbol{A}_I) \quad \text{for all } j \in [n] \ , \tag{3.19}$$

where $\boldsymbol{A}_{I,\widehat{j}}$ is the submatrix of $\boldsymbol{A}_I$ obtained by removing the jth column.

Proof. Consider any $k \in [m]$. Expanding the determinant of $\begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix}$ along the last row yields:

$$\begin{aligned} \det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix} &= \left(\sum_{j=1}^{n} (-1)^{n+1+j} \boldsymbol{A}_{kj} \det(\boldsymbol{A}_{I,\widehat{j}}\ \boldsymbol{b}_I) + (-1)^{2n+2} \boldsymbol{b}_k \det(\boldsymbol{A}_I) \right) \\ &= \left(\sum_{j=1}^{n} \boldsymbol{A}_{kj} \boldsymbol{x}^I_j + \boldsymbol{b}_k \right) \det(\boldsymbol{A}_I) \ . \end{aligned} \tag{3.20}$$

Now suppose that $k \in I$. Since the determinant is an alternative form, we have $\det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix} = 0$. Since $\det(\boldsymbol{A}_I) \neq 0$, we deduce that $\sum_{j=1}^n \boldsymbol{A}_{kj} \boldsymbol{x}^I_j + \boldsymbol{b}_k = 0$. Hence, (3.19) provides the unique solution $\boldsymbol{A}_I^{-1}(-\boldsymbol{b}_I)$ of the system (3.18). □

Cramer's formulæ provide the following characterization of feasible bases.

Lemma 3.16. *Let I be a basis of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. The basis I is feasible if and only if:*

$$\det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix} / \det(\boldsymbol{A}_I) \geq 0 \quad \text{for all } k \in [m] \setminus I \ .$$

Proof. By definition, a basis I is feasible if and only if the basic point $\boldsymbol{x}^I$ satisfy the inequalities $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0$. By definition of a basic point, $\boldsymbol{A}_I \boldsymbol{x}^I + \boldsymbol{b}_I = 0$. Hence, it suffices to check the inequalities $\boldsymbol{A}_k \boldsymbol{x}^I + \boldsymbol{b}_k \geq 0$ for $k \in [m] \setminus I$. Equation (3.20) shows that $\boldsymbol{A}_k \boldsymbol{x}^I + \boldsymbol{b}_k$ is equal to $\det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix} / \det(\boldsymbol{A}_I)$. □

Degeneracy

In general, a feasible basic point $\boldsymbol{x}^I$ may be contained in a hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_k, \boldsymbol{b}_k)$ for some $k \notin I$. When this happens we say that the basis I is *degenerate*. A polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is *non-degenerate* if it does not admit a degenerate basis. Under the non-degeneracy assumption, two distinct bases yield two distinct basic points. Geometrically, this implies that the polyhedron is simple. By extension, we say that a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is (primally) non-degenerate when its feasible set $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is non-degenerate. Non-degeneracy corresponds to the following algebraic conditions.

Lemma 3.17. *A polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is non-degenerate if and only if, for every feasible basis I, the following strict inequalities are satisfied:*

$$\det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k \end{pmatrix} / \det(\boldsymbol{A}_I) > 0 \quad \textit{for all } k \in [m] \setminus I \ .$$

Proof. This follows immediately from the arguments in the proof of Lemma 3.16. □

Edges

A subset $K \subseteq [m+n]$ of cardinality $n-1$ defines a *(feasible) edge*

$$\boldsymbol{\mathcal{E}}_K := \bigcap_{i \in K} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i) \cap \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$$

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 may only consist of a single point. However, this does not happen under the non-degeneracy assumption.

A basic point $\boldsymbol{x}^I$ is contained in the n edges defined by the sets $I \setminus \{i_{\mathsf{out}}\}$ for $i_{\mathsf{out}} \in I$. The edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ is contained in a half-line $\{\boldsymbol{x}^I + \boldsymbol{\mu} \boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}} \mid \boldsymbol{\mu} \geq 0\}$ that we direct with the vector $\boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}} \in K^n$, defined as the unique solution $\boldsymbol{d} \in K^n$ of the system:

$$\boldsymbol{A}_{I \setminus \{i_{\mathsf{out}}\}} \boldsymbol{d} = 0 \text{ and } \boldsymbol{A}_{i_{\mathsf{out}}} \boldsymbol{d} = 1 \ . \tag{3.21}$$

The edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ is unbounded if and only if the set

$$\mathsf{Ent}(I, i_{\mathsf{out}}) := \{i \in [m] \setminus I \mid \boldsymbol{A}_i \boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}} < 0\}$$

is empty. Otherwise, the *length* of the edge is given by:

$$\bar{\boldsymbol{\mu}} = \min \left\{ \frac{\boldsymbol{A}_i \boldsymbol{x}^I + \boldsymbol{b}_i}{-\boldsymbol{A}_i \boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}}} \mid i \in \mathsf{Ent}(I, i_{\mathsf{out}}) \right\} .$$

The other endpoint of the edge is $\boldsymbol{x}' = \boldsymbol{x}^I + \bar{\boldsymbol{\mu}} \boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}}$. Clearly, this point is contained in the hyperplanes $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$ for $i \in I \setminus \{i_{\mathsf{out}}\}$, but also for $i \in \mathsf{Ent}(I, i_{\mathsf{out}})$. Moreover, the intersection $\bigcap_{i \in I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$ is reduced to $\boldsymbol{x}'$ for any $i_{\mathsf{ent}} \in \mathsf{Ent}(I, i_{\mathsf{out}})$. Hence, for any such i_{ent}, the set $I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$ is a feasible basis and $\boldsymbol{x}'$ the corresponding basic point. A basis I' is said to be *adjacent* to a basis I if it is of the form $I' = I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$ for some $i_{\mathsf{out}} \in I$ and $i_{\mathsf{ent}} \in \mathsf{Ent}(I, i_{\mathsf{out}})$. In that case, we also say that the basic point $\boldsymbol{x}^{I'}$ is *adjacent* to the basic point $\boldsymbol{x}^I$. For a non-degenerate polyhedron, the set $\mathsf{Ent}(I, i_{\mathsf{out}})$ is either empty or reduced to a singleton.

Reduced costs and optimal bases

Moving along an edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ from the basic point $\boldsymbol{x}^I$ decreases the objective function $\boldsymbol{x} \mapsto \boldsymbol{c}^\top \boldsymbol{x}$ if and only if the *reduced cost* $\boldsymbol{y}^I_{i_{\mathsf{out}}} = \boldsymbol{c}^\top \boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}}$ is negative.

Lemma 3.18. *The vector of reduced costs $\boldsymbol{y}^I = (\boldsymbol{y}^I_{i_{\mathsf{out}}})_{i_{\mathsf{out}} \in I}$ at a basis I is the unique solution of the following system of equations:*

$$(\boldsymbol{A}_I)^\top \boldsymbol{y} = \boldsymbol{c} \ . \tag{3.22}$$

Proof. By (3.21), the direction vector $\boldsymbol{d}^{I \setminus \{i_{\mathsf{out}}\}}$ is equal to $\boldsymbol{A}_I^{-1} \boldsymbol{e}^{i_{\mathsf{out}}}$, where $\boldsymbol{e}^k$ is the k-th unit vector of $K^{|I|}$. It follows that $\boldsymbol{y}^I_{i_{\mathsf{out}}} = \boldsymbol{c}^\top \boldsymbol{A}_I^{-1} \boldsymbol{e}^{i_{\mathsf{out}}}$. Hence, $\boldsymbol{y}^I = (\boldsymbol{A}_I^{-1})^\top \boldsymbol{c}$, which is the unique solution of (3.22). □

Lemma 3.19. *Let I be a feasible basis. If reduced costs $(\boldsymbol{y}^I_{i_{\mathsf{out}}})_{i_{\mathsf{out}} \in I}$ are non-negative, then the basic point $\boldsymbol{x}^I$ is an optimal solution of the linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$.*

Proof. We can extend $\boldsymbol{y}^I = (\boldsymbol{y}^I_{i_{\mathsf{out}}})_{i_{\mathsf{out}} \in I} \in K^{|I|}$ to a vector K^m by adding components equal to 0. Then, the pair $(\boldsymbol{x}^I, \boldsymbol{y}^I)$ satisfy the complementary slackness conditions (Theorem 3.7). □

We say that a feasible basis I is *optimal* if the reduced costs at I are non-negative. Note that, in case of degeneracy, a basic point $\boldsymbol{x}^I$ may be an optimal solution while I is *not* an optimal basis.

Example 3.20. Consider the linear program:

$$\text{minimize } \boldsymbol{x}_2 \ \text{ s.t } \ \boldsymbol{x}_1 \geq \boldsymbol{x}_2, \ \boldsymbol{x}_1 \geq 0, \ \boldsymbol{x}_2 \geq 0 \ .$$

The point $(0, 0)$ is an optimal solution. It is a basic point for the basis indexing the inequalities $x_1 \geq x_2$ and $x_1 \geq 0$. However, the vector of reduced costs for this basis is $(-1, 1)$, which have a negative component. Hence this basis is not optimal.

The simplex method

We now present the simplex method. For the sake of simplicity, we restrict the exposition to non-degenerate linear programs. The principle of the simplex method is to pivot from feasible basis to feasible basis by following edges. The signs of the reduced costs indicate which pivot improves the objective value and provide a stopping criterion.

Each iteration of the simplex method starts with a feasible basis I. The reduced costs $\boldsymbol{y}^I$ are computed. If $\boldsymbol{y}^I$ is non-negative, then the current basis I is optimal, and the basic point $\boldsymbol{x}^I$ is an optimal solution of the problem.

If the current basis is not optimal, an edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ with a negative reduced cost $\boldsymbol{y}^I_{i_{\mathsf{out}}}$ is selected. The index i_{out} is called a *leaving index*. If the selected edge is unbounded, then the linear program is unbounded. Otherwise, the algorithm *pivots*, *i.e.*, moves to the other end of the selected edge. By the non-degeneracy assumption, the set $\mathsf{Ent}(I, i_{\mathsf{out}})$ is reduced to a singleton $\{i_{\mathsf{ent}}\}$. The index i_{ent} is called the *entering index*. The other endpoint of the edge is a basic point for the basis $I' = I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$. The basis I' is then used to perform the next iteration.

Algorithm 1 describes the simplex method for a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. We have denoted by $\mathsf{Unbounded}(\boldsymbol{A}, \boldsymbol{b})$ a routine which, given a feasible basis I and a leaving

Algorithm 1: The simplex method for non-degenerate linear programs

Data: $\boldsymbol{A} \in K^{m \times n}, \boldsymbol{b} \in K^m$ and $\boldsymbol{c} \in K^n$
Input: A feasible basis I^1 of the linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$.
Output: Either $\mathsf{Unbounded}$, or an optimal basis of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$.

$k \leftarrow 1$
while $\mathsf{SignRedCosts}(\boldsymbol{A}, \boldsymbol{c})(I^k)$ has a negative entry **do**
    $i_{\mathsf{out}} = \phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})(\{I^1, \dots, I^k\})$
    if $\mathsf{Unbounded}(\boldsymbol{A}, \boldsymbol{b})(I^k, i_{\mathsf{out}})$ **then**
        return $\mathsf{Unbounded}$
    $i_{\mathsf{ent}} \leftarrow \mathsf{Pivot}(\boldsymbol{A}, \boldsymbol{b})(I^k, i_{\mathsf{out}})$
    $I^{k+1} \leftarrow I^k \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$
    $k \leftarrow k + 1$
return the optimal basis I^k

index $i_{\mathsf{out}} \in I$, returns true if the edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ is unbounded. Otherwise, the routine $\mathsf{Pivot}(\boldsymbol{A}, \boldsymbol{b})$ returns the entering index i_{ent}. Similarly, $\mathsf{SignRedCosts}(\boldsymbol{A}, \boldsymbol{c})$ is a function that returns the signs of the reduced costs $\boldsymbol{y}^I$.

Given an initial feasible basis I^1, the simplex method builds a sequence of bases $I^1, I^2, \dots, I^N$. At every iteration $k \geq 1$, the leaving index i_{out} is chosen by a function $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ which takes as input $\{I^1, \dots, I^k\}$ the history up to time k. The map ϕ is called a *pivoting rule*.

Proposition 3.21. *Suppose that* $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ *is a non-degenerate linear program, and that the pivoting rule* ϕ *always returns a leaving index* i_{out} *such that reduced cost* $\boldsymbol{y}^I_{i_{\mathsf{out}}}$ *is negative. Then, Algorithm 1 terminates and is correct.*

Proof. Since a feasible basis is given as input, the linear program is always feasible. If an unbounded edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ is encountered, then its direction vector $\boldsymbol{d}$ satisfies $\boldsymbol{c}^\top \boldsymbol{d} = \boldsymbol{y}^I_{i_{\mathsf{out}}} < 0$. For any $\boldsymbol{\lambda} \in K_+$, the point $\boldsymbol{x}^I + \boldsymbol{\lambda} \boldsymbol{d}$ belongs to the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. Consequently, for any $\boldsymbol{\nu} \in K$, we can find a point $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ such that $\boldsymbol{c}^\top \boldsymbol{x} < \boldsymbol{\nu}$ and the linear program is unbounded.

Otherwise the problem admits an optimal solution. By non-degeneracy, each edge has a positive length. Since the pivoting rule always chooses a leaving index with a negative reduced cost, each pivot operation strictly improves the value of the objective function. Consequently, the algorithm terminates, and provides an optimal basis. $\square$

In the following, we shall always assume that a pivoting rule always selects a leaving index with a negative reduced cost.

3.4 Tropical implementation of the simplex method

We now explain how to implement the operations of the simplex method on a linear program $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ (pivoting, computing the signs of the reduced costs, and evaluating the pivoting rule) by tropical means, *i.e.*, using only the signed valuation of $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$. For pivoting and the reduced costs, we shall see that we only need to compute the signs of minors of $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$. As explained in Section 3.2, determinants are tropically tractable polynomials, so the signs of the minors of $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$ can be computed in polynomial time from $\operatorname{sval}\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$. Pivoting rules may be arbitrary procedures. In order to tropicalize, we restrict ourselves to pivoting rules that rely on the signs of polynomials, so that the results of Section 3.2 apply. This does not seem to be a strong restriction, since most known pivoting rules fit in this context.

We begin with the pivoting, and the computation of the signs of reduced costs.

Proposition 3.22. *There exists three maps* $\mathsf{SignRedCosts}^{\mathbb{T}}, \mathsf{Unbounded}^{\mathbb{T}}$ *and* $\mathsf{Pivot}^{\mathbb{T}}$ *satisfying*

$$\mathsf{SignRedCosts}^{\mathbb{T}}(A,c) = \mathsf{SignRedCosts}(\boldsymbol{A},\boldsymbol{c})$$
$$\mathsf{Unbounded}^{\mathbb{T}}(A,b) = \mathsf{Unbounded}(\boldsymbol{A},\boldsymbol{b})$$
$$\mathsf{Pivot}^{\mathbb{T}}(A,b) = \mathsf{Pivot}(\boldsymbol{A},\boldsymbol{b})$$

for any linear programs $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *such that* $\left(\begin{smallmatrix}A & b\\ c & \mathbb{0}\end{smallmatrix}\right) = \operatorname{sval}\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$ *is sign-generic for the polynomials providing the minors of* $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$, *i.e., all polynomials* P *such that* $P\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$ *is a minor of* $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$.

Furthermore, the values of $\mathsf{SignRedCosts}^{\mathbb{T}}(A,c), \mathsf{Unbounded}^{\mathbb{T}}(A,b)$ *and* $\mathsf{Pivot}^{\mathbb{T}}(A,b)$ *can be computed in* $O(n^4), O(m^2n^3)$ *and* $O(m^2n^3)$ *tropical operations respectively, and in space bounded by a polynomial in the input size of* A,b,c.

Proof. The signs of the reduced costs at a basis I are given by the Cramer's formulæ of the system (3.22). This involves the computation of the sign of $\det(\boldsymbol{A}_I)$, and of the determinants $\det((\boldsymbol{A}_{I\setminus\{i\}})^\top\ \boldsymbol{c})$ for $i \in I$, hence $n+1$ minors of $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{b}\\ \boldsymbol{c} & 0\end{smallmatrix}\right)$ of size $n \times n$. By Lemma 3.8, we can compute the signs of these determinants by computing their tropical counterparts on $\left(\begin{smallmatrix}A & b\\ c & \mathbb{0}\end{smallmatrix}\right)$. By Lemma 3.11, computing a $n \times n$ tropical minor of $\left(\begin{smallmatrix}A & b\\ c & \mathbb{0}\end{smallmatrix}\right)$ takes $O(n^3)$ operations and uses a space bounded by the input size of A,b,c.

Pivoting, and determining unboundedness, can be implemented as follows. Given I and $i_{\mathsf{out}} \in I$, we determine which of the $m-n$ sets of the form $I' = I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$, for $i_{\mathsf{ent}} \in [m] \setminus I$, is a feasible basis. By Lemma 3.16, this amounts, for each such I', to computing the sign of $\det(\boldsymbol{A}_{I'})$ and of $\det\begin{pmatrix}\boldsymbol{A}_{I'} & \boldsymbol{b}_{I'}\\ \boldsymbol{A}_k & \boldsymbol{b}_k\end{pmatrix}$ for $k \in [m] \setminus I'$, hence one determinant of size $n \times n$ and $m-n$ determinants of size $(n+1)\times(n+1)$. Thus, to test all I', we have $O((m-n)(m-n+1)) = O(m^2)$ determinants to compute. Each tropical determinant takes $O(n^3)$ operations and uses a space bounded by the input size of A,b,c. $\square$

Remark 3.23. The complexity bounds in Proposition 3.22 can be improved. In Chapter 7, we show that these three complexity bounds can be reduced to $O(n(m+n))$ under additional technical assumptions. In particular, the computation of reduced costs in Chapter 7 is based on the iterative Jacobi algorithm of [Plu90] for tropical Cramer systems. In [RGST05] Richter-Gebert, Sturmfels and Theobald relate the solutions of tropical Cramer systems to solutions of transportation problems. Hence, algorithms for transportation problems may also be used to compute the signs of the reduced costs.

3.4.1 Semi-algebraic pivoting rules

We shall restrict ourselves to *semi-algberaic* pivoting rules,*i.e.*, pivoting rules that have access to information on the problem at hand only through the signs of polynomials evaluated on $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. More precisely, we say that a pivoting rule ϕ is semi-algebraic, if $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is determined from $(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ by the signs of a finite number of polynomials $(P_i^\phi)_{i \in [r]} \subseteq \mathbb{Q}[X_{11}, \dots, X_{(m+1)(n+1)}]$ evaluated on $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$.

Formally, let us denote by Ω^ϕ the oracle which takes as input $i \in [r]$ and returns the sign of $P_i^\phi \left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. If a strategy ϕ is semi-algebraic, then $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ takes as input the history $\{I^1, \dots, I^k\}$ and is allowed to call the oracle Ω^ϕ.

We say that a pivoting rule is *tropically tractable* when:

- the polynomials $(P_i^\phi)_i$ are tropically tractables;
- $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ can be defined in the *arithmetic model of computation with oracle*, which means that $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is allowed to perform arithmetic operations $+, -, \times, /$, and call the oracle Ω^ϕ;
- the number of arithmetic operations, calls to the oracle, and the space complexity of $\phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})(\{I^1, \dots, I^k\})$ is bounded by a polynomial in m, n and k.

Observe that a tropically tractable pivoting rule may involve polynomials that are "untractable" in a classical setting. For example, it may use permanents. A permanent is tropically tractable, as its Newton polytope is, as for the determinant, tge Birkoff polytope. However, computing a classical permanent is a $\#P$-complete problem, see [Val79].

Proposition 3.24. *Let ϕ be a semi-algebraic pivoting rule. There exists a map $\phi^{\mathbb{T}}$ satisfying*

$$\phi^{\mathbb{T}}(A, b, c) = \phi(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$$

for all linear programs $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ such that $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right) = \operatorname{sval} \left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ is sign-generic for the the polynomials $(P_i^\phi)_i$.

Furthermore, if ϕ is tropically tractable, then for any sequence of bases $\{I^1, \dots, I^k\}$, the leaving index provided by $\phi^{\mathbb{T}}(A, b, c)(\{I^1, \dots, I^k\})$ can be computed in time polynomial in k and in the input size of A, b, c.

Proof. This is an immediate consequence of Lemma 3.8, and the definition of a (tropically tractable) semi-algebraic pivoting rule. □

Any $\phi^{\mathbb{T}}$ which arises in this way is called a *tropical pivoting rule*.

Examples of semi-algebraic pivoting rules

Most known pivoting rules are semi-algebraic. Consider for example the rule that selects the smallest index with a negative reduced cost (this rule is known as Bland's rule [Bla77]). Since the signs of the reduced costs are given by determinants, Bland's rule is a semi-algebraic pivoting rule which is also tropically tractable. The tropicalization of Bland's rule will use $O(n^4)$ tropical operations to compute the signs of reduced costs (as in Proposition 3.22) and then $O(m)$ operations to determine the smallest index with a negative reduced cost.

Similarly, every pivoting rule that relies only on the signs of the reduced costs is semi-algebraic. This includes the "least entered" rule, introduced by Zadeh [Zad80]. Indeed, this rule selects the improving pivot with the leaving index that has left the basis the least number of times through the execution of the method. In particular, the "least entered" rule is tropically tractable. The "shadow-vertex" rule is also a tropically tractable semi-algebraic pivoting rule, as we shall see in Chapter 6.

The rule originally proposed by Dantzig [Dan98] picks the leaving index of the smallest negative reduced cost. Since the vector of reduced costs $\boldsymbol{y}^I$ at a basis I is the solution of the system (3.22), its i-th entry, for $i \in I$, is given by the Cramer formula

$$\boldsymbol{y}^I_i = (-1)^{n+\mathsf{idx}(i,I)} \det \begin{pmatrix} \boldsymbol{A}_{I \setminus \{i\}} \\ \boldsymbol{c}^\top \end{pmatrix} / \det(\boldsymbol{A}_I) \ ,$$

where $\mathsf{idx}(i, I)$ is the index of i in the ordered set I. Hence, comparing the two reduced costs $\boldsymbol{y}^I_i$ and $\boldsymbol{y}^I_k$ boils down to computing the sign of the expression

$$\det \begin{pmatrix} \boldsymbol{A}_{I \setminus \{i\}} \\ \boldsymbol{c}^\top \end{pmatrix} - \det \begin{pmatrix} \boldsymbol{A}_{I \setminus \{k\}} \\ \boldsymbol{c}^\top \end{pmatrix} \ , \tag{3.23}$$

which is a polynomial in $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. Hence, Dantzig's rule is semi-algebraic. However, it is unclear whether the polynomial (3.23) is tropically tractable.

The "largest improvement" rule selects the pivot that leads to the largest improvent of the objective value. Hence, we need to compare the objective values of adjacent basic points. At a basis I, the objective value is given by:

$$\boldsymbol{c}^\top \boldsymbol{x}^I = \det \begin{pmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{c} & 0 \end{pmatrix} / \det \boldsymbol{A}_I \ .$$

To see this, one can use Equation (3.20) with the row $(\boldsymbol{A}_k \ \boldsymbol{b}_k)$ replaced by $(\boldsymbol{c} \ 0)$. Consequently, the "largest improvement" rule is semi-algebraic, but it is also unclear whether it is tropically tractable.

3.4.2 The tropical simplex method

Algorithm 2 presents our first tropical implementation of the simplex method. This algorithm can be viewed as a purified version of the method, which is especially useful for theoretical purposes. It is the foundation of the practical algorithm which will be

Algorithm 2: The tropical simplex method for non-degenerate linear programs

Data: A tropical signed matrix $A \in \mathbb{T}_{\pm}^{m \times n}$, two vectors $b \in \mathbb{T}_{\pm}^{m}, c \in \mathbb{T}_{\pm}^{n}$
Input: A subset $I^1 \subseteq [m]$ of cardinality n.
Output: Either $\mathsf{Unbounded}$, or a subset $I \subseteq [m]$ of cardinality n.

$k \leftarrow 1$
while $\mathsf{SignRedCosts}^{\mathbb{T}}(A, c)(I^k)$ has a negative entry **do**
  $i_{\mathsf{out}} = \phi^{\mathbb{T}}(A, b, c)(\{I^1, \dots, I^k\})$
  if $\mathsf{Unbounded}^{\mathbb{T}}(A, b)(I^k, i_{\mathsf{out}})$ **then**
    return $\mathsf{Unbounded}$
  $i_{\mathsf{ent}} \leftarrow \mathsf{Pivot}^{\mathbb{T}}(A, b)(I^k, i_{\mathsf{out}})$
  $I^{k+1} \leftarrow I^k \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$
  $k \leftarrow k + 1$
return I^k

presented in Chapter 7, where more efficient versions of the operations of pivoting and computing reduced costs will be given.

Observe that Algorithm 2, is analogous to Algorithm 1, excepts that the maps $\mathsf{Pivot}, \mathsf{Unbounded}, \mathsf{SignRedCosts}$ and ϕ have been replaced by their tropical counterparts. As an immediate application of Propositions 3.22 and 3.24, we have the following theorem.

Theorem 3.25. *Let $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ be a non-degenerate linear program, and ϕ a semi-algebraic pivoting rule. Suppose that $\left(\begin{smallmatrix} A & b \\ c & 0 \end{smallmatrix}\right) = \operatorname{sval}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ is sign-generic for the polynomials providing a minor of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$, and the polynomials $(P_i^{\phi})_i$ defining ϕ.*

Then, for any feasible basis I^1, the tropical simplex method (Algorithm 2), equipped with the tropical pivoting rule $\phi^{\mathbb{T}}$ and applied on the input A, b, c and I^1, correctly determines if $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is unbounded, or provides an optimal basis.

The sequence of bases $I^1, \dots, I^N$ produced by the tropical simplex method is exactly the sequence of bases obtained by the classical simplex method (Algorithm 1), equipped with the pivoting rule ϕ and applied on the input $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}, I^1$.

If furthermore the pivoting rule ϕ is tropically tractable, the k-th iteration of the tropical simplex method can be performed in time polynomial in k and in the input size of A, b, c.

Chapter 4

Tropical linear programming via the simplex method

In this chapter, we use the tropicalization of the simplex method to solve linear programs over an arbitrary tropical semiring $\mathbb{T} = \mathbb{T}(G)$, *i.e.*, problems of the form

$$\begin{array}{ll} \text{minimize} & c^\top \odot x \\ \text{subject to} & A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^- \ , \end{array} \tag{LP(A,b,c)}$$

where $A^+, A^- \in \mathbb{T}^{m\times n}$, $b^+, b^- \in \mathbb{T}^m$ and $c \in \mathbb{T}^n$. One of the main motivation is to obtain an algorithm for mean payoff games, thanks to the reduction presented in Section 1.1.4.

Our approach is the following. A tropical linear program can be *lifted* to a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ over Hahn series such that the valuation of the entries of $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}$ are given by A^+, A^-, b^+, b^- and c. An optimal solution of the Hahn problem $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ provides an optimal solution of the tropical problem $\mathrm{LP}(A, b, c)$. Hence, the tropicalization of the simplex method presented in Chapter 3 provides an algorithm that solves tropical linear programs, provided that A^+, A^-, b^+, b^- and c satisfy genericity conditions.

However, we cannot solve arbitrary tropical linear programs in this way. To overcome this obstacle, we introduce a perturbation scheme, that transforms an arbitrary tropical linear program into an equivalent, but generic, problem. Our main idea is to use tropical semirings based on additive groups of vectors with a lexicographic order.

This chapter is organized as follows. In Section 4.1, we expose basic results on tropical polyhedra and linear programs. In particular, we explain how tropical polyhedra relate to classical polyhedra over Hahn series. In Section 4.2, we show, that under genericity conditions, the valuation map preseves the face poset of an arrangement of hyperplanes. In particular, this entails a geometric notion of tropical basic points and edges. This geometric interpretation of the tropical simplex method presented in Section 4.3, along with the tropical versions of other related notions such as reduced costs or degeneracy. In Section 4.4, we devise the perturbation scheme that allows to solve arbitrary tropical linear programs with the tropical simplex method.

The contents of this chapter are mostly adapted from [ABGJ13b] and [ABGJ13a].

4.1 Tropical polyhedra

In the following, we work with an arbitrary tropical semiring $\mathbb{T} = \mathbb{T}(G)$. We use indiferrently the notations $(G, +, 0)$ and $(G, \odot, \mathbb{1})$ for the group structure on G. The non-Archimedean field we use is any subfield $\mathbb{K}$ of $\mathbb{R}[\![t^G]\!]$ that contains all the series $\{ct^g \mid c \in \mathbb{R}, g \in G\}$. By Theorem 2.8, any ordered field with value group G that contains $\mathbb{R}$ as a subfield can be identified with such a $\mathbb{K}$.

Tropical halfspaces

An *(affine) tropical halfspace* is the set of points $x \in \mathbb{T}^n$ satisfying a tropical linear inequality:

$$\max(\alpha_1 + x_1, \ldots, \alpha_n + x_n, \beta) \geq \max(\delta_1 + x_1 \ldots, \delta_n + x_n, \gamma) \ , \tag{4.1}$$

where $\alpha, \delta \in \mathbb{T}^n$ and $\beta, \gamma \in \mathbb{T}$. When $\beta = \gamma = \mathbb{0}$, it is said to be a *linear tropical halfspace*. Throughout this paper, we assume that half-spaces are defined by non-trivial inequalities:

Assumption A. There is at least one non $\mathbb{0}$ coefficient in the inequality (4.1), *i.e.*,

$$\max\left(\max_{j\in[n]} \alpha_j, \max_{j\in[n]} \delta_j, \beta, \gamma\right) > \mathbb{0}\,.$$

Tropical halfspaces relate to classical halfspaces, see Figure 4.1 for an illustration.

Lemma 4.1. *The tropical halfspace defined by $\alpha, \delta \in \mathbb{T}^n$ and $\beta, \gamma \in \mathbb{T}$ is the image under the valuation map of the intersection of the halfspace*

$$\left\{ \boldsymbol{x} \in \mathbb{K}^n \mid \eta\left(\sum_{j=1}^{n} t^{\alpha_j}\boldsymbol{x}_j + t^{\beta}\right) \geq \sum_{j=1}^{n} t^{\delta_j}\boldsymbol{x}_i + t^{\gamma} \right\} \tag{4.2}$$

with the positive orthant $\mathbb{K}^n_+$, for any $\eta \in \mathbb{R}$ greater than $n+1$.

Proof. Let $x \in \mathbb{T}^n$ be a point in the tropical halfspace (4.1). Then, the lift $\boldsymbol{x} = (t^{x_1}, \ldots, t^{x_n})$ belongs to the Hahn halfspace (4.2). Indeed, we have:

$$\sum_{j=1}^{n} t^{\delta_j}\boldsymbol{x}_j + t^{\gamma} \leq (n+1)t^{\max(\delta_1+x_1,\ldots,\delta_n+x_n,\gamma)} \leq \eta t^{\delta\odot x\oplus\gamma}$$

and

$$\eta\left(\sum_{j=1}^{n} t^{\alpha_j}\boldsymbol{x}_j + t^{\beta}\right) \geq \eta t^{\max(\alpha_1+x_1,\ldots,\alpha_n+x_n,\beta)} = \eta t^{\alpha\odot x\oplus\beta} \geq \eta t^{\delta\odot x\oplus\gamma}\ .$$

Conversely, suppose $\boldsymbol{x} \in \mathbb{K}^n_+$ belongs to the halfspace (4.2). The Hahn series which appears in the inequality defining (4.2) are non-negative. Since the valuation map is an order-preserving homomorphism from $(\mathbb{K}_+, +, \cdot)$ to $(\mathbb{T}, \max, +)$, it follows that $\mathrm{val}(\boldsymbol{x})$ belongs to the tropical halfspace (4.1). $\square$

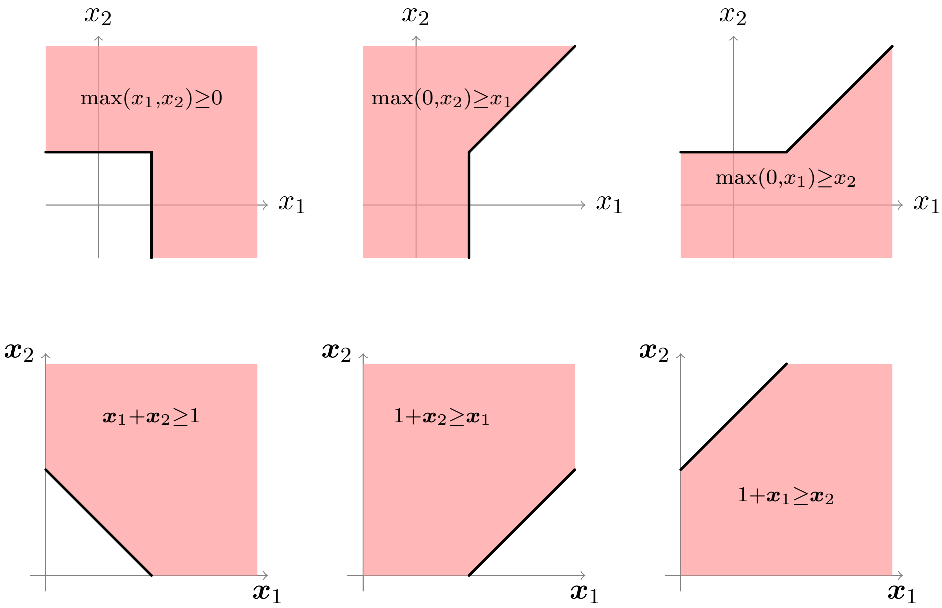

Figure 4.1: Some tropical halfspaces in $\mathbb{T}^2$, and examples of their lifts into halfspaces over the positive orthant $\mathbb{K}^2_+$ of Hahn series.

Lemma 4.2. *Let $\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a},\boldsymbol{b})$ be a halfspace for some $\boldsymbol{a}\in\mathbb{K}^{1\times n}$ and $\boldsymbol{b}\in\mathbb{K}$. Then, the image under the valuation map of $\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a},\boldsymbol{b})\cap\mathbb{K}^n_+$ is a tropical halfspace. More precisely,* $\operatorname{val}(\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a},\boldsymbol{b})\cap\mathbb{K}^n_+)$ *is exactly the set of points $x\in\mathbb{T}^n$ that satisfy*

$$\max(a^+_{11}+x_1,\dots,a^+_{1n}+x_n,b^+)\geq\max(a^-_{11}+x_1,\dots,a^-_{1n}+x_n,b^-)\ , \tag{4.3}$$

where $a^+,a^-\in\mathbb{T}^{1\times n}$ and $b^+,b^-\in\mathbb{T}$ are the values of $\boldsymbol{a}^+=\max(\boldsymbol{a},0),\boldsymbol{a}^-=\min(\boldsymbol{a},0)$ and $\boldsymbol{b}^+=\max(\boldsymbol{b},0),\boldsymbol{b}^-=\min(\boldsymbol{b},0)$ respectively.

Proof. Using the homorphism property of the valuation map, $\operatorname{val}(\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a},\boldsymbol{b})\cap\mathbb{K}^n_+)$ is clearly included in the tropical halfspace (4.3). Conversely, consider any point $x\in\mathbb{T}^n$ satisfying (4.3). We claim that there exists a lift $\boldsymbol{x}\in\mathbb{K}^n_+$ of x, of the form $\boldsymbol{x}=(v_1t^{x_1},\dots,v_nt^{x_n})$ for some vector of positive real numbers $v\in\mathbb{R}^n_+$, which belongs to $\operatorname{val}(\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a},\boldsymbol{b}))$. Let us first treat the case of a linear halfspace, *i.e.*, $\boldsymbol{b}=0$ or equivalently $b^+=b^-=\mathbb{0}$. If the inequality (4.3) is strict at x, then the claim holds with any v with positive entries. Otherwise, $a^+\odot x=a^-\odot x$ and it is sufficient to find of a $v\in\mathbb{R}^n$ which satisfy:

$$\begin{aligned}\sum_{j\in\arg\max\,(a^+\odot x)}\operatorname{lc}(\boldsymbol{a}^+_{1j})v_j &> \sum_{j\in\arg\max\,(a^-\odot x)}\operatorname{lc}(\boldsymbol{a}^-_{1j})v_j \\ v_j &> 0 \qquad\qquad \text{for all } j\in[n]\ .\end{aligned}\tag{4.4}$$

Indeed, given such a v, the Hahn series $\boldsymbol{a}\boldsymbol{x}$ has a positive leading coefficient when $\boldsymbol{x}$ is the lift $(v_jt^{x_j})_j$. The system (4.4) clearly admits a solution, and this proves the claim when $\boldsymbol{b}=0$.

The case $\boldsymbol{b} \neq 0$, easily follows by homogeneization. If $x \in \mathbb{T}^n$ satisfy (4.3), then the point $(x, \mathbb{1}) \in \mathbb{T}^{n+1}$ admits a lift $(c_1 t^{x_1}, \dots, c_n t^{x_n}, c_{n+1}) \in \mathbb{K}_+^{n+1}$, with $c_{n+1} > 0$, which belongs to the linear halfspace $\boldsymbol{\mathcal{H}}^{\geq}((\boldsymbol{a}\ \boldsymbol{b}), 0)$. Consequently, $\boldsymbol{x} = (\frac{c_1}{c_{n+1}} t^{x_1}, \dots, \frac{c_n}{c_{n+1}} t^{x_n})$ is nonnegative and belongs to $\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a}, \boldsymbol{b})$. □

Remark 4.3. The proof above shows that any point x in the tropical halfspace (4.4) has a pre-image by the valuation map in the *interior* of the Hahn halfspace $\boldsymbol{\mathcal{H}}^{\geq}(\boldsymbol{a}, \boldsymbol{b})$.

It follows from the two previous lemmas that we can always assume that each variable (comprising the "affine" variable) appears on at most one side of the inequality defining a tropical halfspace. In other words, any tropical halfspace can be concisely describe by a signed row vector $a = (a_{1j}) \in \mathbb{T}_{\pm}^{1 \times n}$ and a signed scalar $b \in \mathbb{T}_{\pm}$ as:

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

See [GK11, Lemma 1], for an elementary proof.

Tropical s-hyperplanes

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

$$\mathcal{H}(a, b) = \{x \in \mathbb{T}^n \mid a^+ \odot x \oplus b^+ = a^- \odot x \oplus b^-\} \ , \tag{4.5}$$

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

Lemma 4.4. *For any $\boldsymbol{a} \in \mathbb{K}^n$ and any $\boldsymbol{b} \in \mathbb{K}$, let $a = \operatorname{sval}(\boldsymbol{a})$ and $b = \operatorname{sval}(\boldsymbol{b})$. Then:*

$$\operatorname{val}(\boldsymbol{\mathcal{H}}(\boldsymbol{a}, \boldsymbol{b}) \cap \mathbb{K}_+^n) = \mathcal{H}(a, b) \ . \tag{4.6}$$

Proof. Clearly, $\operatorname{val}(\boldsymbol{\mathcal{H}}(\boldsymbol{a}, \boldsymbol{b}) \cap \mathbb{K}_+^n) \subseteq \mathcal{H}(a, b)$. The converse inclusion is a straightfoward consequence of Lemma 4.2. Indeed, if $x \in \mathcal{H}(a, b)$, then x belongs to the two tropical halfspaces $\mathcal{H}^{\geq}(a, b)$ and $\mathcal{H}^{\leq}(a, b)$. Hence, x admits two lifts $\boldsymbol{x}^1, \boldsymbol{x}^2 \in \mathbb{K}_+^n$, one on each side of the hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{a}, \boldsymbol{b})$. Thus the line segment between $\boldsymbol{x}^1$ and $\boldsymbol{x}^2$ intersects the hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{a}, \boldsymbol{b})$. Since $\boldsymbol{x}^1$ and $\boldsymbol{x}^2$ have nonnegative entries, and share the same value x, any point in their convex hull is contained in $\mathbb{K}_+^n$ and has value x. □

Remark 4.5. The set $\mathcal{H}(a, b)$ is said to be *signed* because it corresponds to the tropicalization of the intersection of a Hahn hyperplane with the non-negative orthant. 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 Hahn hyperplane.

Tropical polyhedra

A *tropical polyhedron* is the intersection of finitely many tropical affine halfspaces. 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 :

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

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

Example 4.6. The tropical polyhedron depicted in Figure 1.3 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 1.3 is $\mathcal{H}^{\geq}(A_1, b_1) = \{x \in \mathbb{T}^2 \mid \max(x_1 - 5, x_2 - 3) \geq 0\}$. 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:

$$\begin{aligned} \max(x_2, 0) &\geq x_1 - 7 \,, \\ \max(x_1 - 7, x_2 - 2) &\geq 0 \,, \\ x_1 &\geq \max(x_2 - 6, 0) \,, \end{aligned}$$

which define the half-spaces respectively depicted in purple, green and khaki in Figure 1.3.

Proposition 4.7. *Consider a tropical polyhedra* $\mathcal{P}(A, b)$ *for some* $A \in \mathbb{T}_{\pm}^{m \times n}$ *and* $b \in \mathbb{T}_{\pm}^{m}$*. Then there exist* $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ *and* $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ *such that*

$$\mathcal{P}(A, b) = \operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_{+}^{n}) \ . \tag{4.7}$$

Proof. Lifting the inequalities as in Lemma 4.1, the proposition 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:

$$\begin{aligned} \boldsymbol{A}^+ &= (\eta t^{A_{ij}^+}) \text{ and } \boldsymbol{A}^- = (t^{A_{ij}^-}) \\ \boldsymbol{b}^+ &= (\eta t^{b_i^+}) \text{ and } \boldsymbol{b}^- = (t^{b_i^-}) \end{aligned}$$

where η is a real number strictly greater than $n+1$. Indeed, in this case, if $x \in \mathcal{P}(A, b)$, then $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} > 0$ at $\boldsymbol{x} = (t^{x_1}, \ldots, t^{x_n})$. Hence $\boldsymbol{x}$ belongs to $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_{+}^{n}$. The converse inclusion $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_{+}^{n}) \subseteq \mathcal{P}(A, b)$ follows from the homomorphism property of the valuation map. $\square$

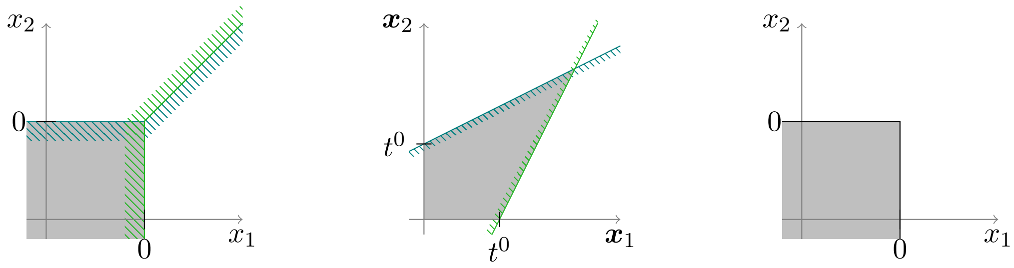

Figure 4.2: Left: the tropical polyhedron $\mathcal{P}$ described in (4.8); middle: the Puiseux polyhedron $\boldsymbol{\mathcal{P}}$ obtained by lifting the inequality representation of $\mathcal{P}$ as in (4.9); right: the set $\operatorname{val}(\boldsymbol{\mathcal{P}})$, which is strictly contained in $\mathcal{P}$.

For arbitrary $\boldsymbol{A} \in \mathbb{K}^{m\times n}$ and $\boldsymbol{b} \in \mathbb{K}^m$, the image by the valuation map of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}^n_+$ is always contained in $\mathcal{P}(A, b)$, where $A = \operatorname{val}(\boldsymbol{A})$ and $b = \operatorname{val}(\boldsymbol{b})$. However, this inclusion may be strict.

Example 4.8. Consider the tropical polyhedron:

$$\mathcal{P} = \{x \in \mathbb{T}^2 \mid \max(0, x_2) \geq x_1,\ \max(0, x_1) \geq x_2,\ x_1 \geq \mathbb{0},\ x_2 \geq \mathbb{0}\} \ . \tag{4.8}$$

A lift of its inequality representation provides the following Puiseux polyhedron:

$$\boldsymbol{\mathcal{P}} = \{\boldsymbol{x} \in \mathbb{K}^2 \mid 1 + 0.5\boldsymbol{x}_2 \geq \boldsymbol{x}_1,\ 1 + 0.5\boldsymbol{x}_1 \geq \boldsymbol{x}_2,\ \boldsymbol{x}_1 \geq 0,\ \boldsymbol{x}_2 \geq 0\} \ . \tag{4.9}$$

See Figure 4.2. By the homomorphism property of the valuation map, we have $\operatorname{val}(\boldsymbol{\mathcal{P}}) \subseteq \mathcal{P}$, but this inclusion is strict for this example. The $\operatorname{val}(\boldsymbol{\mathcal{P}})$ consists of the points $x \in \mathbb{T}^2$ such that $x_1 \leq 0$ and $x_2 \leq 0$. However, $\mathcal{P}$ also contains the half-line $\{(\lambda, \lambda) \mid \lambda > 0\}$. Indeed, suppose that there exist $(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \boldsymbol{\mathcal{P}}$ 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 $1 + 0.5\boldsymbol{x}_1 \geq \boldsymbol{x}_2$ implies that $0.5u_1 \geq u_2$, while $1 + 0.5\boldsymbol{x}_2 \geq \boldsymbol{x}_1$ imposes that $0.5u_2 \geq u_1$, and we obtain a contradiction.

4.1.1 Tropical convexity

We define a *tropical convex set* of $\mathbb{T}^n$ as the image by the valuation map of a convex set of Hahn series contained in the positive orthant $\mathbb{K}^n_+$. Consider a convex combination of two points $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{K}^n_+$:

$$\boldsymbol{z} = \boldsymbol{\lambda x} + \boldsymbol{\mu y} \text{ where } \boldsymbol{\lambda} + \boldsymbol{\mu} = 1 \ , \ \boldsymbol{\lambda} \geq 0 \ , \boldsymbol{\mu} \geq 0 \ .$$

Since $\boldsymbol{x}$ and $\boldsymbol{y}$ have nonnegative entries, and $\boldsymbol{\lambda}, \boldsymbol{\mu} \geq 0$, the value of $\boldsymbol{z}$ is the *tropical convex combination*:

$$\operatorname{val}(\boldsymbol{z}) = \lambda \odot x \oplus \mu \odot y \ , \text{ where } \lambda \oplus \mu = \mathbb{1} \ ,$$

of the tropical vectors $x = \mathrm{val}(\boldsymbol{x})$ and $y = \mathrm{val}(\boldsymbol{y})$ with scaling coefficients $\lambda = \mathrm{val}(\boldsymbol{\lambda})$ and $\mu = \mathrm{val}(\boldsymbol{\mu})$. Hence, a set $\mathcal{P} \subseteq \mathbb{T}^n$ is tropically convex if and only if, for any finite number of points $x^1, \dots, x^k \in \mathcal{P}$, the set $\mathcal{P}$ also contains their *tropical convex hull* $\mathrm{tconv}(x^1, \dots, x^k)$, which is defined by:

$$\mathrm{tconv}(x^1, \dots, x^k) := \left\{ \bigoplus_{i \in [k]} \lambda_i \odot x^i \mid \lambda_i \in \mathbb{T} \text{ for all } i \in [k] \text{ and } \bigoplus_{i \in [k]} \lambda_i = \mathbb{0} \right\} .$$

By analogy with the classical case, we say that a point v in a tropically convex set $\mathcal{P}$ is a *tropical extreme point* of $\mathcal{P}$ if $v \in \mathrm{tconv}(x, y)$ for some $x, y \in \mathcal{P}$ implies that $v = x$ or $v = y$.

It is straightforward to verify that a tropical polyhedron $\mathcal{P}(A, b)$ is stable by tropical convex hull, and thus is tropically convex. Alternatively, this follows from Proposition 4.7.

We define similarly a *tropical convex cone* of $\mathbb{T}^n$ as the image by the valuation map of a convex cone of Hahn series contained in $\mathbb{K}^n_+$. Equivalently, $\mathcal{C} \subseteq \mathbb{T}^n$ is a tropical convex cone if it contains the *tropical conic hull* $\mathrm{tpos}(x^1, \dots, x^k)$ of any finite number of points $x^1, \dots, x^k \in \mathcal{C}$, where:

$$\mathrm{tpos}(x^1, \dots, x^k) := \left\{ \bigoplus_{i \in [k]} \lambda_i \odot x^i \mid \lambda_i \in \mathbb{T} \text{ for all } i \in [k] \right\} .$$

Clearly, a tropical polyhedral cone $\mathcal{P}(A, \mathbb{0})$ is a tropical convex cone.

A point r in a tropical convex cone $\mathcal{C}$ defines a *tropical ray* $[r] := \{\lambda \odot r \mid \lambda \in \mathbb{T} \setminus \{\mathbb{0}\}\}$ of $\mathcal{C}$. We say that $[r]$ is a *tropical extreme ray* of $\mathcal{C}$ if $x \in [r]$ or $y \in [r]$ whenever $r \in \mathrm{tpos}(x, y)$ for some $x, y \in \mathcal{C}$. Equivalently, $r = x \oplus y$ implies $r = x$ or $r = y$.

The *tropical recession cone* of a tropical convex set $\mathcal{P} \subseteq \mathbb{T}^n$ is

$$\mathrm{trec}(\mathcal{P}) := \{r \in \mathbb{T}^n \mid x \oplus (\lambda \odot r) \in \mathcal{P} \text{ for all } x \in \mathcal{P} \text{ and all } \lambda \in \mathbb{T}\} .$$

Proposition 4.9. *If $\mathcal{P}(A, b)$ is a non-empty tropical polyhedron, its tropical recession cone is the tropical polyhedral cone $\mathcal{P}(A, \mathbb{0})$.*

Proof. Consider any $A \in \mathbb{T}_{\pm}^{m \times n}$ and $b \in \mathbb{T}_{\pm}^m$. Let r be an element of the tropical recession cone of $\mathcal{P}(A, b)$. By contradiction, suppose that r does not belong to $\mathcal{P}(A, \mathbb{0})$. Then, $A_i^+ \odot r < A_i^- \odot r$ for some $i \in [m]$. Clearly, this implies $A_i^- \odot r > \mathbb{0}$. Choose any $x \in \mathcal{P}(A, b)$. By definition of the tropical recession cone, for any $\lambda \in \mathbb{T}$, we have

$$(A_i^+ \odot x \oplus b_i^+) \oplus (\lambda \odot A_i^+ \odot r) \geq (A_i^- \odot x \oplus b_i^-) \oplus (\lambda \odot A_i^- \odot r) . \tag{4.10}$$

Since $A_i^+ \odot r < A_i^- \odot r$, we obtain:

$$A_i^+ \odot x \oplus b_i^+ \geq \lambda \odot A_i^- \odot r .$$

As the latter inequality holds for any $\lambda \in \mathbb{T}$, and $A_i^- \odot r > \mathbb{0}$, we obtain a contradiction.

Conversely, let $r \in \mathcal{P}(A, 0)$ and $x \in \mathcal{P}(A, b)$. Then, for any $\lambda \in \mathbb{T}$, the inequality (4.10) is satisfied for all $i \in [m]$, and thus $x \oplus (\lambda \odot r)$ belongs to $\mathcal{P}(A, b)$. Hence, r is an element of the tropical recession cone. $\square$

4.1.2 Homogeneization

It is sometimes convenient to homogeneize a tropical polyhedron $\mathcal{P}(A,b) \subseteq \mathbb{T}^n$ into the tropical into the polyhedral cone $\mathcal{C}(A,b) \subseteq \mathbb{T}^{n+1}$, defined by

$$\mathcal{C}(A,b) := \mathcal{P}((A\ b), \mathbb{0}) = \{(x,\lambda) \in \mathbb{T}^n \times \mathbb{T} \mid A^+ \odot x \oplus b^+ \odot \lambda \geq A^- \odot x \oplus b^+ \odot \lambda\} \ . \quad (4.11)$$

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

$$\begin{aligned} \mathcal{P}(A,b) &\longrightarrow \{y \in \mathcal{C} \mid y_{n+1} = \mathbb{1}\} \\ x &\longmapsto (x, \mathbb{1}) \end{aligned} \quad (4.12)$$

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

As a tropical cone, $\mathcal{C}(A,b)$ is closed under tropical scalar multiplication. For this reason, we identify $\mathcal{C}(A,b)$, with its image in the *tropical projective space* $\mathbb{TP}^n$. The tropical projective space $\mathbb{TP}^n$ consists of the equivalent classes of $\mathbb{T}^{n+1}$ for the relation $x \sim y$ which holds for $x, y \in \mathbb{T}^{n+1}$ if there exists a $\lambda \in \mathbb{T} \setminus \{\mathbb{0}\}$ such that $x = \lambda \odot y$.

Remark 4.10. Consider the tropical semiring $\mathbb{T} = \mathbb{T}(\mathbb{R})$, and let $r^1, \dots, r^k \in \mathbb{T}^n$ be a set of points with entries in $\mathbb{R}$. Then $\mathcal{P} = \operatorname{tpos}(r^1, \dots, r^k)$ 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 halfspaces which are non-empty proper subsets of $\mathbb{T}^n$ correspond to the "tropical halfspaces" studied in [Jos05]. The tropical projective space defined above compactifies the tropical torus (with boundary).

4.1.3 Tropical double description

As their classical counterparts, tropical polyhedra are exactly the tropical convex sets which are *finitely generated*, *i.e.*, the convex hull of a finite number of points and rays. This has been established in [BH84], see also [GP97]. We refer to [GK11] for more references. We include a proof for the sake of completeness.

Theorem 4.11. *Let $\mathcal{P}(A,b)$ be a tropical polyhedron for some $A \in \mathbb{T}_\pm^{m\times n}$ and $b \in \mathbb{T}_\pm^m$. Then:*

$$\mathcal{P}(A,b) = \operatorname{tconv}(V) \oplus \operatorname{tpos}(R) = \{x \oplus y \mid x \in \operatorname{tconv}(V), y \in \operatorname{tpos}(R)\} \ ,$$

where V is the set of tropical extreme points of $\mathcal{P}(A,b)$ and R the set of its tropical extreme rays. Moreover, the sets V, R are finite.

Proof. It is fact sufficient to prove the result for tropical polyhedral cones. Indeed, we can always homogenize a tropical polyhedron $\mathcal{P}(A,b)$ into the tropical polyhedral cone $\mathcal{C}(A,b)$ defined in (4.11). The rays of the homogeneized cone $\mathcal{C}(A,b)$ are in bijection with the points in $\mathcal{P}(A,b)$ and in its recession cone $\mathcal{P}(A,0)$. Moreover, one easily verifies that $[(x, \mathbb{0})]$ is an extreme ray of $\mathcal{C}(A, b)$ if and only if x is an extreme point of $\mathcal{P}(A, b)$. Similarly, $[(x, \mathbb{1})] \in \mathcal{C}(A, b)$ is an extreme ray if and only if x is an extreme ray of the recession cone of $\mathcal{P}(A, b)$.

We now prove that a tropical polyhedral cone $\mathcal{C}$ is the convex hull of its extreme rays. By Proposition 4.7, there exists a polyhedral cone $\boldsymbol{\mathcal{C}} \subseteq \mathbb{K}_+^n$ whose image under the valuation map is $\mathcal{C}$. It follows that $\mathcal{C}$ is the tropical conic hull of a finite number of points $r^1 = \operatorname{val}(\boldsymbol{r}^1), \dots, r^l = \operatorname{val}(\boldsymbol{r}^l)$, where $[\boldsymbol{r}^1], \dots, [\boldsymbol{r}^l]$ are the extreme rays of $\boldsymbol{\mathcal{C}}$.

It turns out that some points of the generating set $R = \{r^1, \dots, r^l\}$ of $\mathcal{C}$ may not yield extreme rays of $\mathcal{C}$; see Example 4.12. Hence, it may happen that for some $i \in [l]$, the point r^i belongs to the tropical conic hull of the other generators $\operatorname{tpos}(R \setminus \{r^i\})$. Clearly, we can remove these points from R and still have a generating set of $\mathcal{C}$. Let us write $I = \{i \in [l] \mid r^i \notin \operatorname{tpos}(R \setminus \{r^i\})$. We claim that $\{[r^i] \mid i \in I\}$ is exactly the set of extreme rays of $\mathcal{C}$.

First, any extreme ray $[x]$ of $\mathcal{C}$ can be decomposed in $x = \bigoplus_{i \in I} \lambda_i \odot r^i$ for some $\lambda \in \mathbb{T}^I$. It follows from the extremality of $[x]$ that $r^i \in [x]$ for some $i \in I$.

Second, consider any $q \in I$. We shall prove that $[r^q]$ is an extreme ray of $\mathcal{C}$. By contradiction, suppose that $r^q = x \oplus y$ for some $x, y \in \mathcal{C}$. There exist $\lambda, \mu \in \mathbb{T}^I$ such that $x = \bigoplus_{i \in I} \lambda_i \odot r^i$ and $y = \bigoplus_{i \in I} \mu_i \odot r^i$. Hence,

$$r^q = (\lambda_q \oplus \mu_q) \odot r^q \oplus \left(\bigoplus_{i \in I \setminus \{q\}} (\lambda_i \oplus \mu_i) \odot r^i \right) . \tag{4.13}$$

This imply the two inequalities:

$$\begin{aligned} r^q &\geq (\lambda_q \oplus \mu_q) \odot r^q \ , \\ r^q &\geq \bigoplus_{i \in I \setminus \{q\}} (\lambda_i \oplus \mu_i) \odot r^i \ . \end{aligned} \tag{4.14}$$

Equality cannot occur in the last inequality, since r^q is not contained in the tropical conic hull of $\{r^i \mid i \in I \setminus \{q\}\}$. Therefore, there must exists a coordinate $j \in [n]$ such that $r^q_j > \bigoplus_{i \in I \setminus \{q\}} (\lambda_i \oplus \mu_i) \odot r^i_j$. Consequently, $r^q_j = (\lambda_q \oplus \mu_q) \odot r^q_j$, by (4.13) and (4.14). It follows that $\lambda_q \oplus \mu_q = \mathbb{1}$. Hence, $\lambda_q = \mathbb{1}$ or $\mu_q = \mathbb{1}$. Without loss of generality, let us assume that $\lambda_q = \mathbb{1}$. Then,

$$x = r^q \oplus \left(\bigoplus_{i \in I \setminus \{q\}} \lambda_i \odot r^i \right) .$$

By (4.14), we have $r^q \geq \bigoplus_{i \in I \setminus \{q\}} \lambda_i \odot r^i$ and thus $x = r^q$. □

Example 4.12. Consider the tropical polyhedron whose feasible set is an usual square:

$$2 \geq x_1 \geq 1 \ , \ 2 \geq x_2 \geq 1 \ .$$

It can be lifted to the square in $\mathbb{K}^2$:

$$t^2 \geq \boldsymbol{x}_1 \geq t^1 \ , \ t^2 \geq \boldsymbol{x}_2 \geq t^1 \ .$$

The point (t^2, t^2) is extreme in $\mathbb{K}^2$. However, its value $(2,2)$ is not an extreme point of the tropical polyhedron. Indeed, $(2,2) = (2,1) \oplus (1,2)$. In fact, the usual square, as a tropical polyhedron, is a triangle: it is the tropical convex hull of $(1,1), (2,1)$ and $(1,2)$.

Remark 4.13. In the proof of Theorem 4.11, we actually showed that if $\mathcal{P}(A,b) = \operatorname{conv}(V) \oplus \operatorname{tpos}(R)$ for some finite sets V, R, then V contains the set of extreme points of $\mathcal{P}(A,b)$, and R its set of extreme rays.

The "converse" of Theorem 4.11 also holds: a finitely generated tropical convex set is a tropical polyhedron.

Theorem 4.14. *Let $V, R \subseteq \mathbb{T}^n$ be two finite sets. Then the tropical convex set*

$$\operatorname{tconv}(V) \oplus \operatorname{tpos}(R) := \{x \oplus y \mid x \in \operatorname{tconv}(V), y \in \operatorname{tpos}(R)\} \tag{4.15}$$

is a tropical polyhedron.

The classical counterpart of Theorem 4.14 can be proved using separation hyperplanes. The same approach also works in the tropical case. The *polar* $C^\diamond$ of a tropical convex cone C parametrizes the set of tropical linear halfspaces containing C, *i.e.*,

$$C^\diamond := \{(\alpha, \beta) \in \mathbb{T}^n \times \mathbb{T}^n \mid \alpha^\top \odot x \geq \beta^\top \odot x \text{ for all } x \in C\} \ .$$

By definition, the tropical cone C is included in the intersection of the tropical halfspaces parametrized by $C^\diamond$. When C is finitely generated, the converse inclusion holds, thanks to the following separation theorem.

Theorem 4.15. *Let $\mathcal{C} = \operatorname{tpos}(R)$ be a tropical convex cone generated by a finite set $R \subseteq \mathbb{T}^n$. If $v \in \mathbb{T}^n$ does not belong to $\mathcal{C}$, then there exists a tropical halfspace that contains $\mathcal{C}$ and does not contain v.*

Proof. This tropical separation theorem holds for general convex cone (see [CGQ04, Zim77, CGQS05]), but in tropical semirings which are complete, or conditionnaly complete, for their natural ordering. However, in the case of *finitely* generated cones, the completeness requirement can be dispensed with. The proof below is and adaptation of [CGQ04] to our setting.

It is sufficient to prove the theorem for a point $v \in \mathbb{T}^n \setminus \mathcal{C}$ with finite entries. Indeed, suppose that $v_j = \mathbb{0}$ for some $j \in [n]$ and let $J = \{j \in [n] \mid v_j > \mathbb{0}\}$. Then, the projection v_J of v in $\mathbb{T}^J$ has finite entries. The projection $\mathcal{C}_J$ of $\mathcal{C}$ is a tropical convex cone which is finitely generated by the projections of the generators $r \in R$. A tropical halfspace in $\mathbb{T}^J$ separating v_J from $\mathcal{C}_J$ extends to a tropical halfspace of $\mathbb{T}^n$ separating v from $\mathcal{C}$.

We now assume that $v \in \mathbb{T}^n \setminus \mathcal{C}$ has finite entries. If $\mathcal{C}$ is included in one of the coordinate hyperplane $\{x \in \mathbb{T}^n \mid x_j = \mathbb{0}\}$ for some $j \in [n]$, then the inequality $x_j \leq \mathbb{0}$ provides a tropical halfspace separating $\mathcal{C}$ from v. Hence we can restrict to the case where, for any $j \in [n]$, there exists a generator $r \in R$ such that $r_j > \mathbb{0}$. Without loss of generality, we may also assume that $r \neq \mathbb{0}_{\mathbb{T}^n}$ for every $r \in R$. Let $\lambda \in \mathbb{T}^R$ be defined for every $r \in R$ by

$$\lambda_r := \min_{j \in [n]} v_j - r_j \ ,$$

with the convention $-\mathbb{0} = +\infty$. Note that $\lambda_r \in \mathbb{T}$ is well-defined since v has finite entries, and there exists a $j \in [n]$ such that $r_j > \mathbb{0}$. Consider the point $\pi := \bigoplus_{r \in R} \lambda_r \odot r$. Equivalently, for any $j \in [n]$:

$$\pi_j = \max_{r \in R} \lambda_r + r_j \ . \tag{4.16}$$

Clearly, $\pi_j > \mathbb{0}$ since λ has finite entries and at least one generator $r \in R$ satisfies $r_j > \mathbb{0}$.

We claim that the tropical halfspace:

$$\mathcal{H}^{\geq} = \{x \in \mathbb{T}^n \mid \max_{j \in [n]} x_j - v_j \geq \max_{j \in [n]} x_j - \pi_j\}$$

separate the tropical cone $\mathcal{C}$ from v. Observe that $-v_j \in \mathbb{T}$ and $-\pi_j \in \mathbb{T}$, since v and π have finite entries.

First we show that v does not belong to $\mathcal{H}^{\geq}$. Consider any $j \in [n]$ and let $r^* \in G$ be a generator that attains the maximum in (4.16), *i.e.*, such that $\pi_j = \lambda_{r^*} + r^*_j$. Since π_j is finite, so is r^*_j. As $\lambda_{r^*} \leq v_j - r^*_j$, it follows that $\pi_j \leq v_j$. However, the equality $\pi = v$ cannot occurs, as $\pi \in \mathcal{C}$ and $v \notin \mathcal{C}$. Hence, there exists at least one $j \in [n]$ such that $\pi_j < v_j$. Consequently,

$$\max_{j \in [n]} v_j - v_j = 0 < \max_{j \in [n]} v_j - \pi_j \ ,$$

and v does not belong to $\mathcal{H}^{\geq}$.

Second, we prove the inclusion $\mathcal{C} \subseteq \mathcal{H}^{\geq}$. By convexity, it is sufficient to show that every $r \in R$ belongs to $\mathcal{H}^{\geq}$. Fix any generator $r \in R$. By definition of λ_r, we have $\max_{j \in [n]} r_j - v_j = -\lambda_r$. Moreover, $-\lambda_r \geq r_j - \pi_j$ for every $j \in [n]$, by definition of π. This concludes the proof. □

Corollary 4.16. *Let $R \in \mathbb{T}^n$ be a finite set. Then, the tropical convex cone* $\operatorname{tpos}(R)$ *is a tropical polyhedral cone.*

Proof. By Theorem 4.15, the tropical cone C is the intersection of the tropical halfspaces in its polar:

$$C = \bigcap_{(\alpha,\beta) \in C^\diamond} \{x \in \mathbb{T}^n \mid \alpha^\top x \geq \beta^\top x\} \ .$$

By convexity, the polar $C^\diamond$ of the tropical cone $C = \operatorname{tpos}(R)$ is the intersection of finitely many tropical halfspaces:

$$C^\diamond = \bigcap_{r \in R} \{(\alpha, \beta) \in \mathbb{T}^n \times \mathbb{T}^n \mid r^\top \odot \alpha \geq r^\top \odot \beta\} \ .$$

By Theorem 4.11, $C^\diamond$ is a finitely generated convex cone, *i.e.*, there exists a finite set $G \in \mathbb{T}^{2n}$ such that $C^\diamond = \operatorname{tpos}(G)$. Hence, if $x \in \mathbb{T}^n$ satisfies the inequality $\alpha^\top \odot x \geq \beta^\top \odot x$

for all $(\alpha, \beta) \in G$, then x also satisfies this inequality for any $(\alpha, \beta) \in C^{\diamond}$ by convexity. Consequently, C is the following tropical polyhedral cone:

$$C = \bigcap_{(\alpha,\beta)\in G} \{x \in \mathbb{T}^n \mid \alpha^\top x \geq \beta^\top x\} . \qquad \square$$

Proof of Theorem 4.14. Let us homogeneize the convex set $P = \operatorname{tconv}(V) \oplus \operatorname{tpos}(R) \subseteq \mathbb{T}^n$ into the convex cone $C = \operatorname{pos}(V' \cup R') \subseteq \mathbb{T}^{n+1}$ where $V' = \{(v, \mathbb{1}) \mid v \in V\}$ and $R' = \{(r, \mathbb{0}) \mid r \in R\}$. By Corollary 4.16, there exists a matrix $(A\ b) \in \mathbb{T}^{m\times(n+1)}$ such that $C = \mathcal{P}((A\ b), 0)$. If $x \in P$, then $(x, \mathbb{1}) \in C$ and thus $x \in \mathcal{P}(A, b)$. Conversely, if $x \in \mathcal{P}(A, b)$, then $(x, \mathbb{1}) \in C$ and thus $x \in P$. $\square$

4.1.4 Tropical linear programming

A *tropical linear program* is an optimization problem of the form

$$\begin{array}{ll} \text{optimize} & c^\top \odot x \\ \text{subject to} & x \in \mathcal{P}(A, b) , \end{array}$$

where $A \in \mathbb{T}_{\pm}^{m\times n}$, $b \in \mathbb{T}_{\pm}^{m}$ are signed matrices, $c \in \mathbb{T}^n$ is an unsigned vector, and "optimize" means either "maximize" or "minimize". We say that the program is *infeasible* if the tropical polyhedron $\mathcal{P}(A, b)$ is empty. Otherwise, it is said to be *feasible*. A maximization problem is *unbounded* if for any $\nu \in \mathbb{T}$, there exists a $x \in \mathcal{P}(A, b)$ such that $c^\top \odot x > \nu$. Since $\mathbb{0}$ is a lower bound on any tropical number, tropical minimization problems are always bounded. An *optimal solution* of a minimization problem is a $x^* \in \mathcal{P}(A, b)$ such that:

$$c^\top \odot x^* \leq c^\top \odot x \text{ for all } x \in \mathcal{P}(A, b) . \tag{4.17}$$

For a maximization problem, the inequality $\leq$ in (4.17) is replaced by $\geq$.

Lemma 4.17. *A tropical linear maximization problem is either infeasible, unbounded, or admits an optimal solution. A tropical linear minimization problem is either infeasible or admits an optimal solution.*

Proof. If the linear program is infeasible, then the other possibilities are excluded. Now suppose that $\mathcal{P}(A, b)$ is not empty. Then, by Theorem 4.11, there exist a finite number of points $v^1, \dots, v^k \in \mathbb{T}^n$ and $r^1, \dots, r^l \in \mathbb{T}^n$ such that

$$\mathcal{P}(A, b) = \operatorname{tconv}(v^1, \dots, v^k) \oplus \operatorname{tpos}(r^1, \dots, r^l) .$$

First consider a maximization problem. If there exists a $i \in [l]$ such that $c^\top \odot r^i > \mathbb{0}$, then the linear program is unbounded. Otherwise any feasible point x satisfy

$$c^\top \odot x \leq \max_{i\in[k]} c^\top \odot v^i = c^\top \odot \left(\bigoplus_{i\in[l]} v^i \right) ,$$

and the element $\bigoplus_{i\in[l]} v^i$ is an optimal solution.

In case of a feasible minimization problem, any feasible point x can be written as

$$x = \left(\bigoplus_{i\in[k]} \lambda_i \odot v^i\right) \oplus \left(\bigoplus_{i\in[l]} \mu_i \odot r^i\right)$$

where $\bigoplus_{i\in[k]} \lambda_k = \mathbb{1}$. Consequently, we have $c^\top \odot x \geq \bigoplus_{i\in[k]} \lambda_i \odot (c^\top \odot v^i)$. Consider any $i^* \in [k]$ such that $\lambda_{i^*} = \mathbb{1}$. Then, $c^\top \odot x \geq c^\top \odot v^{i^*} \geq \min_{i\in[k]} c^\top \odot v^i$. Consequently, the optimal value of the linear program is $\min_{i\in[k]} c^\top \odot v^i$ and it is attained on some v^i. □

Remark 4.18. The proof of Lemma 4.17 shows that a feasible maximization problem is unbounded if and only if there exists a r in the polyhedral cone $\operatorname{tpos}(r^1, \dots, r^l)$ such that $c^\top r > \mathbb{0}$. The set $\operatorname{tpos}(r^1, \dots, r^l)$ is the tropical recession cone of $\mathcal{P}(A, b)$ and it is equal to $\mathcal{P}(A, 0)$ by Proposition 4.9. Hence, Lemma 3.5 admits a tropical counterpart.

Remark 4.19. The proof of Lemma 4.17 also shows that a feasible a bounded maximization problem admits $\bigoplus_{i\in[l]} v^i$ as an optimal solution. Observe that this point is optimal for all objective vector c that yields a bounded problem. In case of a feasible minimization problem, we proved that there always exists an extreme point v^i which is an optimal solution.

In the following, we shall consider only minimization problems, that we denote as follows:

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

Proposition 4.20. *There is a way to associate to every tropical linear program of the form LP(A, b, c) a linear program over* $\mathbb{K}$

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}),\ \boldsymbol{x} \geq 0 \end{array} \tag{4.18}$$

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 (4.18) *is precisely the feasible set of the tropical linear program LP(A, b, c); in particular, the former program is feasible if, and only if, the latter one is feasible;*

(ii) the valuation of any optimal solution of (4.18) *(if any) is an optimal solution of LP(A, b, c).*

Proof. The lifted matrices $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$ and $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ provided by Proposition 4.7 proves the first part of the proposition. For the second part, choose any $\boldsymbol{c} \in \operatorname{sval}^{-1}(c)$. Since c has tropically non-negative entries, $\boldsymbol{c}$ also has non-negative entries. It follows that $\boldsymbol{c}^\top \boldsymbol{x} > 0$ for all $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}^n_+$. If $\mathcal{P}(A, b)$ is not empty, then $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}^n_+$ is also not

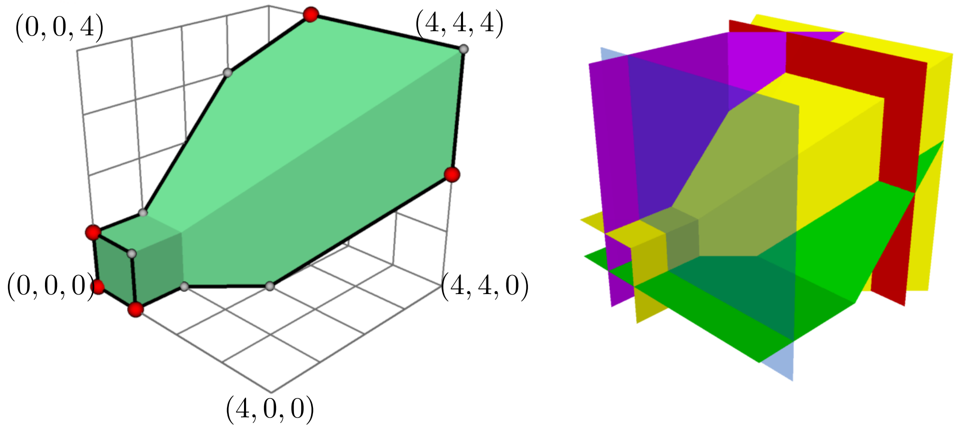

Figure 4.3: The tropical polyhedron defined by the inequalities (4.19) and its external representation.

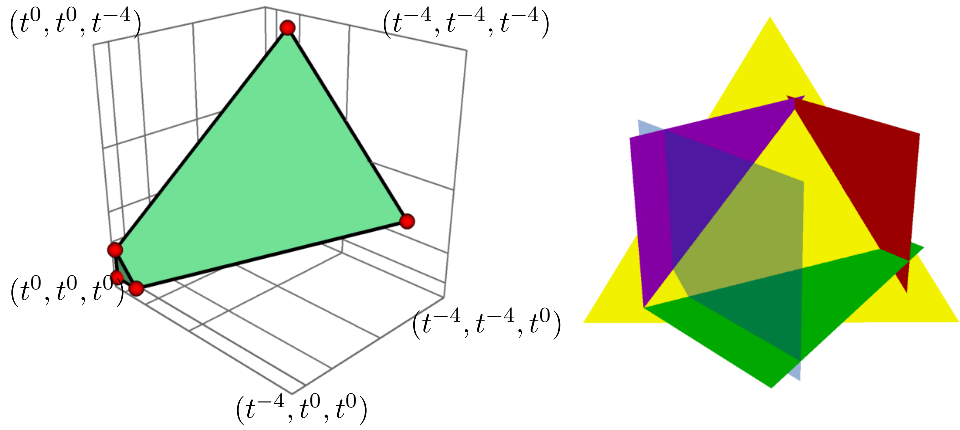

Figure 4.4: A lift of the tropical polyhedron defined by the inequalities (4.19) and its external representation.

empty by the first part of the proposition. Hence, the Hahn linear program (4.18) admits an optimal solution $\boldsymbol{x}^*$ by Proposition 3.4. As $\boldsymbol{c}^\top \boldsymbol{x}^* \leq \boldsymbol{c}^\top \boldsymbol{x}$ for all $\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n$, it follows from the homomorphism property of the valuation map that $c^\top \odot \mathrm{val}(\boldsymbol{x}^*) \leq c^\top \odot x$ for all $x \in \mathrm{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n) = \mathcal{P}(A, b)$. $\square$

Example 4.21. Throughout the rest of this manuscript, we will illustrate some results on the following problem.

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

These constraints define the tropical polyhedron represented in Figure 4.3. A lift of this tropical polyhedron is depicted in Figure 4.4. The optimal value of this tropical linear program is 0 and the set of optimal solutions is the ordinary square:

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

However, over Hahn 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 value $(0, 0, 0)$, which is an optimum for the tropical linear program.

The homogeneization of the polyhedron (4.19) is the cone described by the inequalities:

$$
\begin{aligned}
\max(x_4, x_2 - 1) &\geq \max(x_1 - 1, x_3 - 1) \\
x_3 &\geq \max(x_4, x_2 - 2) \\
x_2 &\geq x_4 \\
x_1 &\geq \max(x_4, x_2 - 3) \\
x_4 &\geq x_2 - 4\ ,
\end{aligned}
\tag{4.20}
$$

where the coordinate x_4 plays the role of the affine component. For the sake of simplicity, the linear half-spaces in (4.20) are still referred to as ($\mathcal{H}_1$)–($\mathcal{H}_5$).

4.2 Generic arrangements of tropical hyperplanes

A set of Hahn hyperplanes $\{\boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)\}_{i \in [m]}$ induces a cell decomposition of the ambient space $\mathbb{K}^n$ into polyhedra. Similarly a set of tropical hyperplanes $\{\mathcal{H}(A_i, b_i)\}_{i \in [m]}$ decomposes the space $\mathbb{T}^n$. In this section, we establish the following relation between these two decompositions (see Figure 4.5 for an illustration).

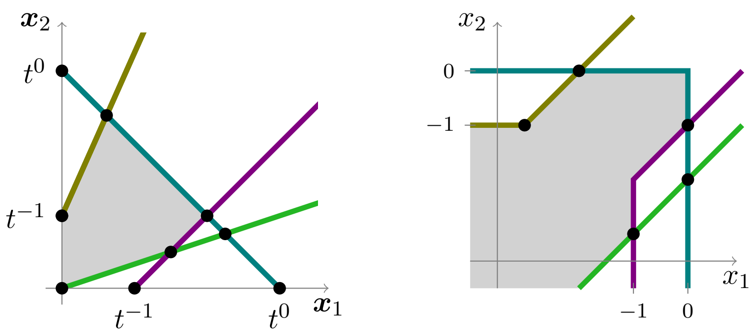

Figure 4.5: Illustration of Theorem 4.22. Left: the cell decomposition of $\mathbb{K}_+^2$ induced by the hyperplane arrangement of Example 1.2. Right: the cell decomposition of $\mathbb{T}^2$ induced by the tropicalization of the hyperplanes of Example 1.2. Note that in the tropical decomposition, the zero-dimensionnal cells that correspond to points with $\mathbb{0}$ entries are not represented.

Theorem 4.22. *Suppose that $(A\ b) \in \mathbb{T}^{m\times(n+1)}$ is sign-generic for all minors polynomials. Then, for all $\boldsymbol{A} \in \operatorname{sval}^{-1}(A)$, $\boldsymbol{b} \in \operatorname{sval}^{-1}(b)$ and $I \subseteq [m]$,*

$$\operatorname{val}\Big(\boldsymbol{\mathcal{P}}_I(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n\Big) \;=\; \mathcal{P}_I(A, b)\,. \tag{4.21}$$

where, for any subset of rows $I \subseteq [m]$, we denote

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

By Theorem 4.22, the set of tropical s-hyperplanes $\{\mathcal{H}(A_i, b_i)\}_{i\in[m]}$ induces a cellular decomposition of $\mathbb{T}^n$ into tropical polyhedra. We call this collection of tropical polyhedra the *signed cells* of the arrangement $\{\mathcal{H}(A_i, b_i)\}_{i\in[m]}$. Notice that the signed cells form an intersection poset thanks to Theorem 4.22.

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 4.3, this is the case for the ordinary line segment $[(1, 0, 1), (1, 1, 1)]$.

Example 4.23. Consider the tropical polyhedral cone $\mathcal{C}$ in $\mathbb{T}^3$ given by the three homogenous constraints

$$x_2 \geq \max(x_1, x_3) \tag{4.22}$$
$$x_1 \geq \max(x_2 - 2, x_3 - 1) \tag{4.23}$$
$$\max(x_1, x_3 + 1) \geq x_2 - 1\ . \tag{4.24}$$

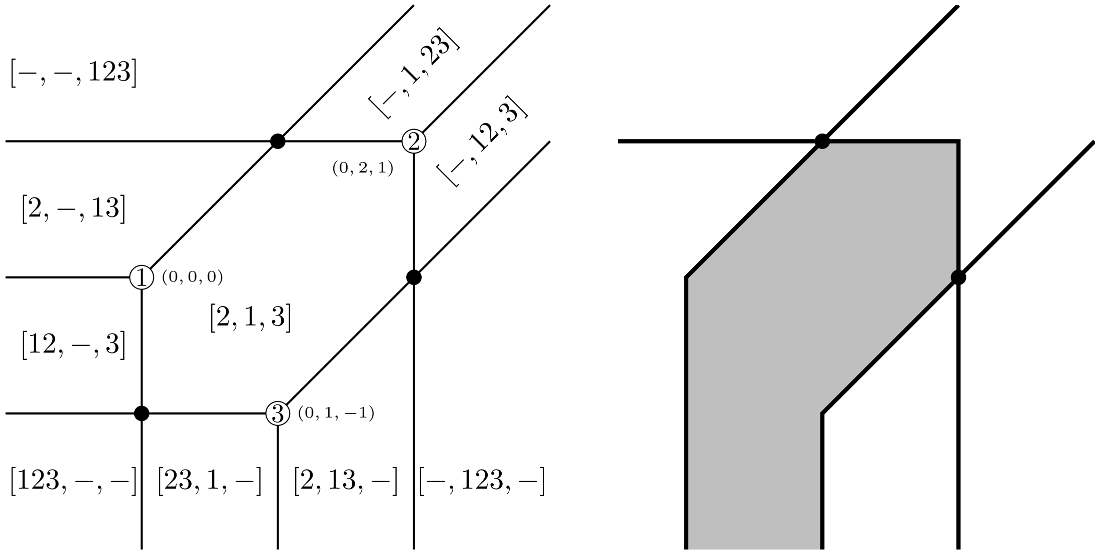

Figure 4.6: Unsigned (left) and signed (right) cell decompositions induced by the three tropical s-hyperplanes in Example 4.23.

This gives rise to an arrangement of three tropical s-hyperplanes in which $\mathcal{C}$ forms one signed cell; see Figure 4.6 (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 ten unsigned maximal cells; in Figure 4.6 (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 "$\geq$" by "$\leq$" in some subset of the constraints above. For instance, exchanging "$\geq$" by "$\leq$" in (4.22) 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 the signed and unsigned cells depend on the arrangement of s-hyperplanes, while several different arrangements may describe the same tropical polyhedron. For example,

$$\{x \in \mathbb{T}^2 \mid x_1 \oplus x_2 \leq 1\} = \{x \in \mathbb{T}^2 \mid x_1 \leq 1 \text{ and } x_2 \leq 1\}\,. \tag{4.25}$$

Even if a canonical external representation exists, see [AK13], it may not satisfy the genericity conditions of Theorem 4.22. Thus this approach does not easily lead to a meaningful notion of faces for tropical polyhedra.

The rest of this section is devoted to prove Theorem 4.22.

4.2.1 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 $j \in \arg(|W_i| \odot x)$.

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 4.7 (there, 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.

Example 4.24. Let W be the matrix formed by the coefficients of the system (4.20), and consider the point $x = (1, 0, 0, 0)$ (corresponding to $(1, 0, 0)$ via the bijection (4.12)). The inequalities ($\mathcal{H}_1$), ($\mathcal{H}_2$) and ($\mathcal{H}_3$) are tight at x. They read

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

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 4.7. 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. A matching can be identified with a bijection $\sigma : I \to J$.

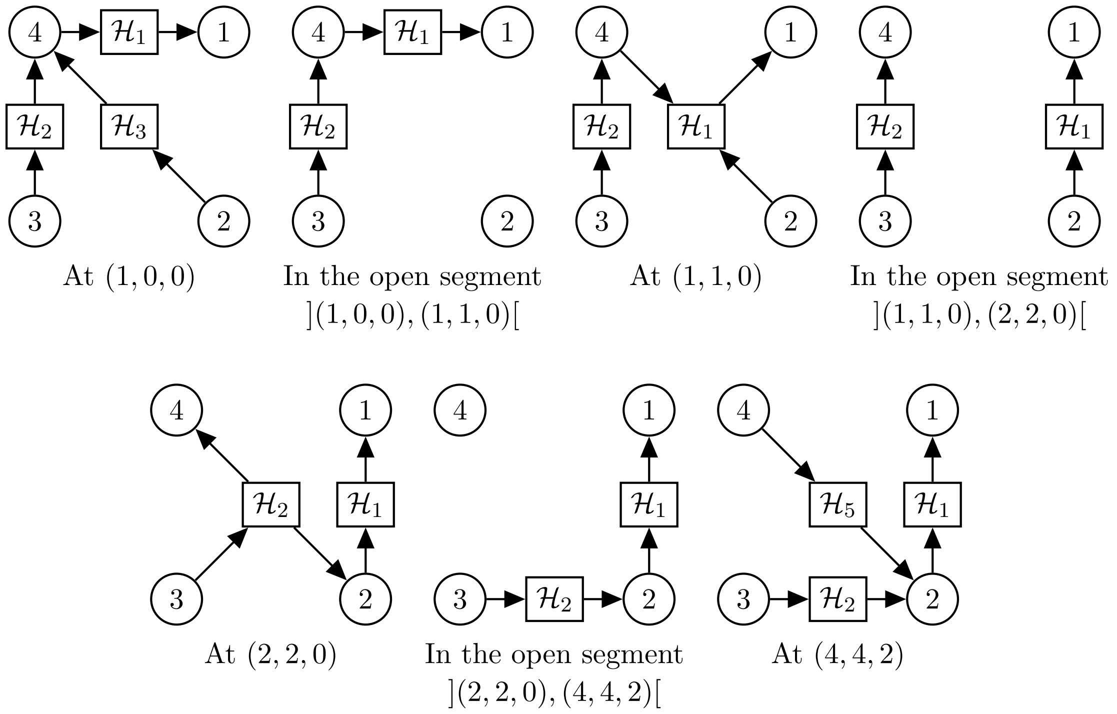

Figure 4.7: Tangent digraphs at various points of the tropical cone obtained by homogenization of the tropical polyhedron defined by the inequalities (4.19). Hyperplane nodes are rectangles and coordinate nodes are circles.

Lemma 4.25. *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 this matching is a solution of the maximal assignment problem with costs $(|W_{ij}|)_{i\in I, j\in J}$.*

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

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

Since x has no $\mathbb{0}$ entries, this implies $\sum_{p=1}^{q} |W_{i_p j_p}| \geq \sum_{p=1}^{q} |W_{i_p \sigma(i_p)}|$ for any bijection $\sigma : I \to J$. □

Lemma 4.26. *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 admits a square submatrix W' which is not generic for the determinant polynomial. Moreover, if the cycle is directed in the tangent digraph $\vec{\mathcal{G}}_x$, then W' is not sign-generic for the determinant polynomial.*

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 $\mathcal{G}_x$. 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 pair-wise 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 made with rows $i_1, \dots, i_q$ and columns $j_1, \dots j_q$. By Lemma 4.25, the bijections σ and τ are both maximizing in $|\operatorname{tdet}(W')|$, hence W' is not generic for the determinant.

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 differs 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, and W' is not sign-generic for the determinant. This completes the proof. □

4.2.2 Cells of an arrangement of signed tropical hyperplanes

Theorem 4.27. *Suppose that $(A\ b) \in \mathbb{T}^{m \times (n+1)}$ is sign-generic for all minors polynomials. 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 $\operatorname{val}(\boldsymbol{\mathcal{C}} \cap \mathbb{K}_+^{n+1}) \subseteq \mathcal{C}$ 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 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:

$$\sum_{j \in \arg(|W_i| \odot x)} \operatorname{lc}(\boldsymbol{w}_{ij}) v_j > 0 \quad \text{for all } i \in I \;, \tag{4.26}$$

$$v_j > 0 \quad \text{for all } j \in [n+1] \;, \tag{4.27}$$

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 (4.26) 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 $\boldsymbol{W}_i \boldsymbol{x} \geq 0$ 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 $j \in \arg(|W_i| \odot x)$ and $f_{ij} = 0$ otherwise. We claim that there exists a $v \in \mathbb{R}^{n+1}$ such that $Fv > 0$ and $v > 0$, or, equivalently, that the following polyhedron is not empty:

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

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

$$F^\top \alpha + \lambda \leq 0 \tag{4.28}$$

$$\sum_{i \in I} \alpha_i + \sum_{j \in [n+1]} \lambda_j > 0 \tag{4.29}$$

Note that if α is the 0 vector, then by (4.29), there exists a $\lambda_j > 0$ for some $j \in [n+1]$, which contradicts (4.28). Thus, the set $K = \{i \in I \mid \alpha_i > 0\}$ is not empty. Let $J \subseteq [n+1]$ be defined by:

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

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 (4.28) yields:

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

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 the nodes $K \cup J$). Then, by Lemma 4.26, the matrix W is not sign-generic for a minor polynomial. This contradicts the sign-genericity of W and proves the claim.

Now we consider the polyhedron $\mathcal{P}(A, b)$. The inclusion $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n) \subseteq \mathcal{P}(A, b)$ 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 4.27 shows that valuation commutes with intersection for halfspaces in general position. This extends to mixed intersection of halfspaces and (signed) hyperplanes.

Proof of Theorem 4.22. We first prove the result when $I = [m]$. In this case, the claim is about the intersection of all (Hahn or signed tropical) hyperplanes in the arrangement. The first inclusion $\operatorname{val}\big(\bigcap_{i=1}^m \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i) \cap \mathbb{K}_+^n\big) \subseteq \bigcap_{i=1}^m \mathcal{H}(A_i, b_i)$ is trivial. Conversely, let $x \in \bigcap_{i=1}^m \mathcal{H}(A_i, b_i)$. The point x belongs to the tropical polyhedron $\mathcal{P}(A, b)$. By Theorem 4.27, x admits a lift in $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \cap \mathbb{K}_+^n$. But observe that the choice of tropical signs for the rows of $(A\ b)$ is arbitrary. Indeed, if $(A'\ b')$ is obtained by multiplying some rows of $(A\ b)$ by $\ominus\mathbb{1}$, then $(A'\ b')$ satisfy the conditions of Theorem 4.27 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 Hahn polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}^s, \boldsymbol{b}^s) \cap \mathbb{K}_+^n$, where $(\boldsymbol{A}^s\ \boldsymbol{b}^s) = \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_m \end{pmatrix} (\boldsymbol{A}\ \boldsymbol{b})$.

Since the Hahn points $\boldsymbol{x}^s$ are non-negative with value x, any point in their convex hull is also non-negative with value 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 $m \geq 2$ hyperplanes. Let S^+ (resp. S^-) be the set of all signs 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 $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_m,\boldsymbol{b}_m)$, 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 $i \notin I$. $\square$

4.3 The simplex method for tropical linear programming

We shall now use the tropical simplex method to solve a tropical linear program. By Proposition 4.20, a solution of $\mathrm{LP}(A,b,c)$ can be found by applying the simplex method to a classical linear program

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b}),\ \boldsymbol{x} \geq 0 \end{array} \tag{4.30}$$

over Hahn series, for some $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right) \in \operatorname{sval}^{-1}\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$. Note that the feasible set of (4.30) is included in the positive orthant. To ease the connection between tropical and classical linear programs, we shall make the following assumption.

Assumption B. The matrix $(A\ b) \in \mathbb{T}_{\pm}^{m\times(n+1)}$ is such that $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ is included in the positive orthant $\mathbb{K}_+^n$ for any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$.

This assumption can be easily satisfied by adding explicitely the (implicit) inequalities $x \geq \mathbb{0}$ to the description of $\mathcal{P}(A,b)$.

Tropical basic points

Proposition-Definition 4.28. *Suppose that $(A\ b) \in \mathbb{T}_{\pm}^{m\times(n+1)}$ is sign-generic for the minor polynomials and satisfies Assumption B. Let I be a subset of $[m]$ of cardinality n such that* $\operatorname{tdet}(A_I) \neq \mathbb{0}$*. If the set*

$$\mathcal{P}_I(A,b) = \{x \in \mathcal{P}(A,b) \mid A_I^+ \odot x \oplus b_I^+ = A_I^- \odot x \oplus b_I^-\} \tag{4.31}$$

is not empty, it contains a unique point x^I. In this case, I is called a (feasible) basis, and x^I a (feasible) basic point, of $\mathcal{P}(A,b)$.

For any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$, the feasible bases of $\mathcal{P}(A,b)$ are exactly the feasible bases of $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$, Moreover, for any feasible basis I, the basic point x^I of $\mathcal{P}(A,b)$ is the value of the basic point $\boldsymbol{x}^I$ of $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$.

Proof. Consider any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$. By Assumption B, the Hahn polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is included in the positive orthant $\mathbb{K}_+^n$. Hence, by Corollary 4.22, the set (4.31) is exactly the image under the valuation map of the set

$$\boldsymbol{\mathcal{P}}_I(\boldsymbol{A}, \boldsymbol{b}) = \{\boldsymbol{x} \in \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) \mid \boldsymbol{A}_I \boldsymbol{x} + \boldsymbol{b}_I = 0\} \tag{4.32}$$

Since $\operatorname{tdet}(A_I) \neq \mathbb{0}$, and A_I is sign-generic for the determinant, it follows from Lemma 3.8 that $\det(\boldsymbol{A}_I) \neq 0$. As a consequence, I is a basis of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, and the intersection $\bigcap_{i \in I} \boldsymbol{\mathcal{H}}(\boldsymbol{A}_i, \boldsymbol{b}_i)$ contains only the Hahn basic point $\boldsymbol{x}^I$. If this point is contained in $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, *i.e.*, if the basis is feasible for $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, then the set (4.32) is reduced to $\{\operatorname{val}(\boldsymbol{x}^I)\}$. Otherwise, the set (4.32) is empty. □

Given a basis, the corresponding basic point can be obtained as follows.

Proposition 4.29. *Suppose that $(A\ b) \in \mathbb{T}_\pm^{m \times (n+1)}$ is sign-generic for the minor polynomials and satisfies Assumption B. Let $I \subseteq [m]$ be a feasible basis of $\mathcal{P}(A, b)$. The jth component of the basic point $x^I \in \mathbb{T}^n$ is given by*

$$x_j^I = (\ominus \mathbb{1})^{\odot n+1+j} \odot \operatorname{tdet}(A_{I,\widehat{j}}\ b_I) \odot (\operatorname{tdet}(A_I))^{\odot -1} = |\operatorname{tdet}(A_{I,\widehat{j}}\ b_I)| - |\operatorname{tdet}(A_I)| \ . \tag{4.33}$$

Proof. By Lemma 3.15, the tropical basic point x^I is the image under the valuation map of the Hahn basic point $\boldsymbol{x}^I$ of the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, for any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$. The rest of the proposition then follows from Cramer's formulæ (Proposition 3.15) and Lemma 3.8. □

Proposition 4.30. *Every extreme point of a tropical polyhedron is a feasible basic point.*

Proof. Let $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$ be the lifted matrix given by Proposition 4.7, so that $\operatorname{val}(\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}))$ coincides with $\mathcal{P}(A, b)$. Let $\boldsymbol{V}$ be the set of basic points of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$. By Proposition 4.28, $V = \operatorname{val}(\boldsymbol{V})$ is the set of tropical basic points of $\mathcal{P}(A, b)$. The set of extreme points of $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is exactly the set of its basic points by Proposition 3.14. Hence, $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b}) = \operatorname{conv}(\boldsymbol{V}) + \operatorname{pos}(\boldsymbol{R})$ for some finite set $\boldsymbol{R} \subseteq \mathbb{K}^n$ by Theorem 3.1. Consequently, $\mathcal{P}(A, b) = \operatorname{tconv}(V) \oplus \operatorname{tpos}(\operatorname{val}(\boldsymbol{R}))$. It then follows from Remark 4.13 that V contains the set of extreme points of $\mathcal{P}(A, b)$. □

However, in contrast with the classical case, a tropical basic point may not be an extreme point. This happens in particular in Example 4.12, where $(2, 2)$ is a basic point but not an extreme point. Observe that the set of basic points actually depends on the external representation chosen for a tropical polyhedron. For example, the tropical polyhedron of Example 4.12 can also be described by:

$$2 \geq \max(x_1, x_2), x_1 \geq 1, x_2 \geq 1 \ .$$

With this representation, $(2, 2)$ is no longer a basic point. In fact, the set of basic points is $\{(2, 1), (1, 2), (1, 1)\}$, and it coincides with the set of extreme points.

Non-degeneracy

By analogy with the classical case, we say that a feasible basis I of a tropical polyhedron $\mathcal{P}(A,b)$ is *degenerate* if the tropical basic point x^I belongs to an s-hyperplane $\mathcal{H}(A_k,b_k)$ for some $k \notin I$. When there is no degenerate basis, we say that the tropical polyhedron $\mathcal{P}(A,b)$, and the tropical linear program $\mathrm{LP}(A,b,c)$, is *non-degenerate*.

Note that even if $(A\ b)$ is generic for the minor polynomials, it may happen $\mathcal{P}(A,b)$ is degenerate. This happens in particular in the tropical counterpart of the degenerate linear program in Example 3.20.

Example 4.31. The tropical polyhedron of $\mathbb{T}^2$ defined by the inequalities:

$$x_1 \leq x_2 \ , \ x_1 \geq \mathbb{0} \ , \ x_2 \geq \mathbb{0} \tag{4.34}$$

has a sign-generic matrix, and $(\mathbb{0},\mathbb{0})$ is a basic point for the three distinct bases $x_1 = x_2, x_1 = \mathbb{0}$, and $x_1 = x_2, x_2 = \mathbb{0}$, and $x_1 = \mathbb{0}, x_2 = \mathbb{0}$.

The following conditions are sufficient to ensure non-degeneracy.

Lemma 4.32. *Suppose that $(A\ b) \in \mathbb{T}_\pm^{m\times(n+1)}$ is sign-generic for the minor polynomials and satisfies Assumption B. If one of the following conditions holds, the tropical polyhedron $\mathcal{P}(A,b)$ is non-degenerate.*

(i) The polyhedron $\mathcal{P}(A,b)$ does not contain a point with $\mathbb{0}$ entries.

(ii) The matrix $(A\ b)$ is of the form $\left(\begin{smallmatrix} A' & b' \\ D & \mathbb{0}\end{smallmatrix}\right)$, where b' has no $\mathbb{0}$ entries, and D is a $n\times n$ diagonal matrix with tropically positive entries on the diagonal.

Proof. Let $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$ and I a feasible basis. By Corollary 4.22, it is sufficient to prove that, for any $k \in [m]\setminus I$, the basic point $\boldsymbol{x}^I$ is not contained in the hyperplane $\boldsymbol{\mathcal{H}}(\boldsymbol{A}_k,\boldsymbol{b}_k)$. By contradiction, suppose that $\boldsymbol{A}_k\boldsymbol{x}^I + \boldsymbol{b}_k = 0$. Then, $\det\left(\begin{smallmatrix} \boldsymbol{A}_I & \boldsymbol{b}_I \\ \boldsymbol{A}_k & \boldsymbol{b}_k\end{smallmatrix}\right) = 0$ by (3.20). Thus $\operatorname{tdet}\left(\begin{smallmatrix} A_I & b_I \\ A_k & b_k\end{smallmatrix}\right) = \mathbb{0}$ by genericity on that minor polynomial. By definition of the tropical determinant, we have

$$\left|\operatorname{tdet}\begin{pmatrix} A_I & b_I \\ A_k & b_k\end{pmatrix}\right| \geq \bigoplus_{j\in[n]} |A_{kj}| \odot |\operatorname{tdet}(A_{I,\widehat{j}}\ b_I)| \oplus |b_k| \odot |\operatorname{tdet}(A_I)| \ . \tag{4.35}$$

(i) If the polyhedron does not contain point with $\mathbb{0}$ entries, the basic point x^I does not have $\mathbb{0}$ entries. By Proposition 4.29, it follows that $|\operatorname{tdet}(A_{I,\widehat{j}}\ b_I)| \neq \mathbb{0}$ for all $j \in [n]$. Moreover, $|\operatorname{tdet}(A_I)| \neq \mathbb{0}$. At least one of the $|A_{k1}|,\dots,|A_{kn}|,|b_k|$ is different from $\mathbb{0}$ by Assumption A. Consequently, we obtain the contradiction $\operatorname{tdet}\left(\begin{smallmatrix} A_I & b_I \\ A_k & b_k\end{smallmatrix}\right) > \mathbb{0}$ by (4.35).

(ii) Now suppose that $(A\ b)$ is of the form $\left(\begin{smallmatrix} A' & b' \\ D & \mathbb{0}\end{smallmatrix}\right)$, where b' has no $\mathbb{0}$ entries, and D is a $n\times n$ diagonal matrix with tropically positive entries on the diagonal. Since $|\operatorname{tdet}(A_I)| \neq \mathbb{0}$, Equation (4.35) imply $b_k = \mathbb{0}$. As the components of b' are not equal to $\mathbb{0}$, we have $A_k = D_l$ for some $l \in [n]$. Hence, $\left|\operatorname{tdet}\left(\begin{smallmatrix} A_I & b_I \\ A_k & b_k\end{smallmatrix}\right)\right| = |D_{ll}| \odot |\operatorname{tdet}(A_{I,\widehat{l}}\, b_I)|$.

Let $\sigma : [n] \mapsto I$ be a maximizing permutation in $\operatorname{tdet}(A_I)$. As $\operatorname{tdet}(A_I) \neq \mathbb{0}$, we have $A_{\sigma(j)j} \neq \mathbb{0}$ for all $j \in [n]$. In particular, $A_{\sigma(l)l} \neq \mathbb{0}$. Consequently, $\sigma(l)$ either indexes the row $(A_k\ b_k) = (D_l\ \mathbb{0})$, or a row of $(A'\ b')$. Since $\sigma(l) \in I$ and $k \notin I$, we deduce that $\sigma(l)$ indexes a row of $(A'\ b')$, and thus that $b_{\sigma(l)} \neq \mathbb{0}$. Finally, we obtain the contradiction:

$$\mathbb{0} = |\operatorname{tdet}(A_{I,\widehat{l}}\ b_I)| \geq |b|_{\sigma(l)} + \sum_{j \in [n] \setminus \{l\}} |A|_{j\sigma(j)} > \mathbb{0}\ . \qquad \square$$

Tropical edges

Proposition-Definition 4.33. *Suppose that $(A\ b) \in \mathbb{T}_{\pm}^{m \times (n+1)}$ is sign-generic for the minor polynomials and satisfies Assumption B. Let K be a subset of $[m]$ of cardinality $n-1$ such that A_K has a maximal square submatrix with a non $\mathbb{0}$ tropical determinant. If the set*

$$\mathcal{P}_K(A, b) = \{x \in \mathcal{P}(A, b) \mid A_K^+ \odot x \oplus b_K^+ = A_I^- \odot x \oplus b_K^-\} \tag{4.36}$$

is not empty, then it is called an edge of $\mathcal{P}(A, b)$.

The edges of $\mathcal{P}(A, b)$ are exactly image under the valuation map 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 4.28. $\square$

Since a bounded edge of a Hahn polyhedron is the convex hull of two of its basic points, a bounded edge of a tropical polyhedron is the tropical convex hull of two of its basic points. We refer to Chapter 7 for a more thorough description of tropical edges.

Tropical reduced costs

We also define a tropical version of reduced costs.

Proposition-Definition 4.34. *Suppose that $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right) \in \mathbb{T}_{\pm}^{(m+1) \times (n+1)}$ satisfies Assumption B and is sign-generic for the minor polynomials. Let I be a feasible basis of LP(A, b, c). The vector of* reduced costs *of LP(A, b, c) at I is the vector $y^I \in \mathbb{T}_{\pm}^{|I|}$ with entries:*

$$y_i^I = (\ominus \mathbb{1})^{n + \mathsf{idx}(i, I)} \odot \operatorname{tdet} \begin{pmatrix} A_{I \setminus \{i\}} \\ c^\top \end{pmatrix} \odot (\operatorname{tdet}(A_I))^{\odot -1} \text{ for all } i \in I \tag{4.37}$$

where $\mathsf{idx}(i, I)$ is the index of i in the ordered set I.

For any $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right) \in \operatorname{sval}^{-1} \left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$, and for any feasible basis I, the reduced costs vector y^I of LP(A, b, c) is the image under the signed valuation map of the reduced costs vector $\boldsymbol{y}^I$ of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$.

Proof. The reduced costs vector $\boldsymbol{y}^I$ at a basis I is the unique solution of the system $\boldsymbol{A}_I^\top \boldsymbol{y} = \boldsymbol{c}$. We then apply Cramer's formulæ to this system, and Lemma 3.8. $\square$

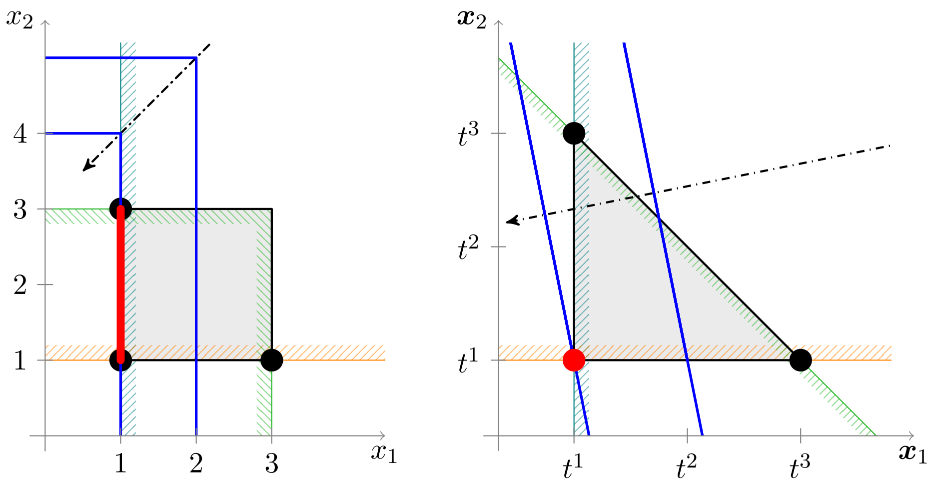

Figure 4.8: Illustration of Example 4.35. The set of optimal solutions of the tropical linear program on the left is the segment between $(1, 1)$ and $(1, 3)$. In particular, the tropical basic point $(1, 3)$ is an optimal solution. However, the corresponding basis is not optimal. Indeed, on a lift of this tropical linear program (on the right), the basic point defined by the blue and green hyperplanes is not optimal.

If, at a feasible basis I all reduced costs $(y_i^I)_{i \in I}$ have a non-negative tropical sign, we say that I is an *optimal basis* of $\mathrm{LP}(A, b, c)$. Observe that at an optimal basis I, the basic point x^I is an optimal solution of $\mathrm{LP}(A, b, c)$. Indeed, $\boldsymbol{x}^I$ is an optimal solution of the Hahn linear program provided by Proposition 4.7.

However, it may happen that a basic point x^I is an optimal solution of $\mathrm{LP}(A, b, c)$, while I is not an optimal basis, *i.e.*, some reduced costs have negative sign. Unlike the classical case, this can happen even on a non-degenerate tropical linear program.

Example 4.35. Consider following the tropical linear program (illustrated in Figure 4.8):

$$\text{minimize } \max(x_1, x_2 - 4) \text{ s.t. } 3 \geq \max(x_1, x_2),\ x_1 \geq 1,\ x_2 \geq 1 \ .$$

It can be described by the matrices

$$A = \begin{pmatrix} \ominus\mathbb{1} & \ominus\mathbb{1} \\ \mathbb{1} & \mathbb{0} \\ \mathbb{0} & \mathbb{1} \end{pmatrix},\ b = \begin{pmatrix} 3 \\ \ominus 1 \\ \ominus 1 \end{pmatrix} \text{ and } c = \begin{pmatrix} \mathbb{1} \\ -4 \end{pmatrix}$$

One easily verify that $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ is sign-generic for the minor polynomials, and that Assumption B is satisfied.

The set of optimal solutions of this tropical linear program is the line segment between the two basic points $(1, 1)$ and $(3, 1)$. The basic point $(3, 1)$ is defined by the system

$$3 = \max(x_1, x_2),\ x_2 = 1 \ .$$

The tropical reduced cost for the constraint $3 \geq \max(x_1, x_2)$ is $\ominus\mathbb{1}$, hence tropically negative.

To see this, it may be easier to look at a lift of this tropical linear program over Hahn series:

$$\text{minimize } \boldsymbol{x}_1 + t^{-4}\boldsymbol{x}_2 \text{ s.t. } t^3 \geq \boldsymbol{x}_1 + \boldsymbol{x}_2,\ \boldsymbol{x}_1 \geq t,\ \boldsymbol{x}_2 \geq t\ .$$

The Hahn basic point corresponding to the tropical basic point $(3, 1)$ is defined by:

$$t^3 = \boldsymbol{x}_1 + \boldsymbol{x}_2,\ \boldsymbol{x}_2 = t\ .$$

The vector of reduced cost for the corresponding basis is the unique solution $\boldsymbol{y} \in \mathbb{K}^2$ of

$$\begin{pmatrix} -1 & 0 \\ -1 & 1 \end{pmatrix} \boldsymbol{y} = \begin{pmatrix} 1 \\ t^{-4} \end{pmatrix}$$

hence is $\begin{pmatrix} 1-t^{-4} \\ -1 \end{pmatrix}$. Its image under the signed valuation map is $\begin{pmatrix} \mathbb{1} \\ \ominus\mathbb{1} \end{pmatrix}$.

The tropical simplex method solves generic tropical linear programs

Proposition 4.36. *Let ϕ be a semi-algebraic pivoting rule. Suppose that LP(A, b, c) a non-degenerate tropical linear program satisfying Assumption B and such that $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ is sign-generic for the minor polynomials, and all polynomials (P_i^ϕ).*

Then, for any feasible basis I of LP(A, b, c), the tropical simplex method, equipped with the tropical pivoting rule $\phi^{\mathbb{T}}$, and applied on A, b, c, I, terminates and returns an optimal basis of LP(A, b, c).

Proof. Consider the matrix $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right) \in \operatorname{sval}^{-1}(\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right))$ given by Proposition 4.20. The conditions of Theorem 3.25 are satisfied and thus the tropical simplex method terminates. The linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ seek a minimum of $\boldsymbol{x} \mapsto \boldsymbol{c}^\top \boldsymbol{x}$ with $\boldsymbol{c} \geq 0$, and the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ is included in the positive orthant. Hence, $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is bounded and the tropical simplex method returns an optimal basis I^* of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. It follows that the basic point $\boldsymbol{x}^{I^*}$ is optimal for $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. By Proposition 4.20, the tropical basic point $x^{I^*} = \operatorname{val}(\boldsymbol{x}^{I^*})$ is optimal for LP(A, b, c). □

We conclude this section by applying the tropical simplex algorithm to the running example 4.21.

Example 4.37. We start from the tropical basic point $(4, 4, 2)$ associated with the basis $I = \{\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_5\}$. For this basis, the 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.

4.4 Perturbation scheme

We fix a totally ordered abelian group G, and a tropical linear program $\mathrm{LP} = \mathrm{LP}(A, b, c)$ on the tropical semiring $\mathbb{T} = \mathbb{T}(G)$. We shall construct a tropical linear program $\widetilde{\mathrm{LP}}$ which is generic and whose solution provides an optimal solution of LP. The problem $\widetilde{\mathrm{LP}}$ is defined on a "bigger" semiring $\mathbb{I} = \mathbb{T}(F \times G \times H)$, where F and H are two groups and $F \times G \times H$ is ordered lexicographically. We shall use $(\mathbb{Z}, +)$ for F, and the additive group $\mathbb{Z}^{\mathbb{N}\times\mathbb{N}}$, with a lexicographic order, for H. For computational purposes, we shall see below that it is sufficient to instantiate H as $\mathbb{Z}^{(m+n+3)\times(n+3)}$ to use the tropical simplex method on tropical linear program defined by m inequalities in dimension n.

Intuitively, a tuple $(f, g, h) \in F \times G \times H$ corresponds to an element of G of the form $fM + g + h\epsilon$, where M is an infinite formal value and ϵ and infinitesimal formal value. An element $g \in G$ is lifted into $(0, \beta, \cdot)$. In contrast, the elements of $\mathbb{I}$ of the form $(f, \cdot, \cdot)$ with $f \neq 0$ correspond to different layers of infinite values, namely $-\infty$ if $f < 0$, and $+\infty$ if $f > 0$. Finally, the semiring $\mathbb{I}$ has its own bottom element, $\mathbb{0}_{\mathbb{I}}$, which also corresponds to $\mathbb{0}_{\mathbb{T}}$.

We define a *canonical embedding* ψ, which maps a tropical signed number $x \in \mathbb{T}_{\pm}$ to $\psi(x) \in \mathbb{I}_{\pm}$ defined by:

$$\psi(x) := \begin{cases} (0, |x|, 0) & \text{if } x \text{ is tropically positive },\\ \ominus(0, |x|, 0) & \text{if } x \text{ is tropically negative },\\ \mathbb{0}_{\mathbb{G}} & \text{if } x = \mathbb{0}_{\mathbb{T}} . \end{cases}$$

The map ψ is extended to matrices component-wise, and we let

$$\mathsf{A} := \psi(A),\ \mathsf{b} = \psi(b),\ \text{and } \mathsf{c} = \psi(c) . \tag{4.38}$$

In order to obtain a non-degenerate linear program, we wish to use Lemma 4.32 (ii). So, we replace the $\mathbb{0}$ entries of b by "infinitely small" but finite entries. We define $\mathsf{d} \in \mathbb{I}_{\pm}^{m}$ to be a vector such that

$$\mathbb{1}_{\mathbb{I}} \gg \mathsf{d}_i > \mathbb{0}_{\mathbb{I}} \text{ for all } i \in [m] . \tag{4.39}$$

For example, we can take $\mathsf{d}_i = (-1, 0, 0)$ for all $i \in [m]$. We want to solve the following linear program over $\mathbb{I}$:

$$\begin{array}{ll} \text{maximize} & \mathsf{c} \odot \mathsf{x} \\ \text{subject to} & \mathsf{A}^{+} \odot \mathsf{x} \oplus (\mathsf{b}^{+} \oplus \mathsf{d}) \geq \mathsf{A}^{-} \odot \mathsf{x} \oplus \mathsf{b}^{-} \\ & \qquad\qquad\qquad \mathsf{x} \geq \mathbb{0} . \end{array} \tag{4.40}$$

However, the matrix of this problem may not be sign-generic. We now use the H-entries of the elements of $\mathbb{I}$ to satisfy the genericity conditions. Let $E = (\varepsilon^{i,j})$ be a basis of the $\mathbb{Z}$-module $H = \mathbb{Z}^{\mathbb{N}\times\mathbb{N}}$. For example, we can use the canonical basis where $\varepsilon^{i,j}$ is the infinite matrix with all entries equal to 0 except the (i, j)-th entry which is equal to 1. We define a *perturbation map* Φ_E, that associates to any $\mathsf{M} \in \mathbb{I}_{\pm}^{p\times q}$, the perturbed matrix

$\widetilde{M} = \varPhi_E(\mathsf{M}) \in \mathbb{I}_{\pm}^{p\times q}$ defined by

$$\widetilde{M}_{ij} = \begin{cases} (f_{ij}, g_{ij}, \varepsilon^{i,j}) & \text{if } \mathsf{M}_{ij} \text{ is tropically positive, and } |\mathsf{M}_{ij}| = (f_{ij}, g_{ij}, \cdot) \\ \ominus(f_{ij}, g_{ij}, -\varepsilon^{i,j}) & \text{if } \mathsf{M}_{ij} \text{ is tropically negative, and } |\mathsf{M}_{ij}| = (f_{ij}, g_{ij}, \cdot) \\ \mathbb{0}_{\mathbb{I}} & \text{if } \mathsf{M}_{ij} = \mathbb{0}_{\mathbb{T}} \ . \end{cases}$$

Lemma 4.38. *Let $\mathsf{M} \in \mathbb{I}_{\pm}^{p\times q}$. The perturbed matrix $\varPhi_E(\mathsf{M})$ is generic for any polynomial $P \in \mathbb{Q}[X_{11}, \dots, X_{pq}]$.*

Proof. Let $\widetilde{M} = \varPhi_E(\mathsf{M})$ and $P = \sum_{\alpha\in\varLambda} q_\alpha X^\alpha$. If $\operatorname{trop}(P)(\widetilde{M}) = \mathbb{0}$ then there is nothing to prove. Otherwise, let $\alpha, \beta \in \varLambda \subseteq \mathbb{N}^{p\times q}$ be two maximizers in $|\operatorname{trop}(P)(\widetilde{M})|$. We have $\sum_{i,j} \pm\alpha_{i,j}\varepsilon^{i,j} = \sum_{i,j} \pm\beta_{i,j}\varepsilon^{i,j}$. Since $E = (\varepsilon^{i,j})$ is a basis, and α, β have non-negative entries, it follows that $\alpha = \beta$. □

We are now considering now the following tropical linear program on $\mathbb{I}$:

$$\begin{aligned} \text{minimize} \quad & \widetilde{c}^\top \odot \mathsf{x} \\ \text{subject to} \quad & \widetilde{A}^+ \odot \mathsf{x} \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{A}^- \odot \mathsf{x} \oplus \widetilde{b}^- \\ & \widetilde{\mathrm{Id}} \odot \mathsf{x} \geq \mathbb{0}_{\mathbb{I}^n} \end{aligned} \tag{$\widetilde{\mathrm{LP}}$}$$

with parameters given by:

$$\begin{pmatrix} \widetilde{c}^\top & \mathbb{0} & \mathbb{0} \\ \widetilde{A} & \widetilde{b} & \widetilde{d} \\ \widetilde{\mathrm{Id}} & \mathbb{0}_{\mathbb{I}^n} & \mathbb{0} \end{pmatrix} = \varPhi_E \begin{pmatrix} \psi(c)^\top & \mathbb{0} & \mathbb{0} \\ \psi(A) & \psi(b) & \mathsf{d} \\ \mathrm{Id} & \mathbb{0}_{\mathbb{I}^n} & \mathbb{0} \end{pmatrix} \tag{4.41}$$

where Id is the $n \times n$ identity matrix of $\mathbb{I}$.

Example 4.39. Let us illustrate our perturbation scheme on a very simple example. Consider the tropical polyhedron $\mathcal{P}(A, b)$ in $\mathbb{T}^2$ defined by:

$$x_1 \geq x_2 \ \text{ and } \ x_2 \geq x_1 \ .$$

This polyhedron consists of the diagonal $x_1 = x_2$ (see Figure 4.9, right). After embedding into $\mathbb{I}^2$, and replacing the $\mathbb{0}$ entries of the right-hand side b by $(-1, 0, 0)$ as in (4.40), we obtain the polyhedron in $\mathbb{I}^2$ depicted in Figure 4.9 (middle), which defined by:

$$\begin{aligned} \mathsf{x}_1 \oplus (-1, 0, 0) &\geq \mathsf{x}_2 \\ \mathsf{x}_2 \oplus (-1, 0, 0) &\geq \mathsf{x}_1 \ . \end{aligned}$$

Finally, applying the perturbation map as in (4.41) provides the polyhedron illustrated in the left of Figure 4.9, which can be described by:

$$\begin{aligned} (0, 0, \varepsilon^{1,1}) \odot \mathsf{x}_1 \oplus (-1, 0, \varepsilon^{1,3}) &\geq (0, 0, -\varepsilon^{1,2}) \odot \mathsf{x}_2 \\ (0, 0, \varepsilon^{2,2}) \odot \mathsf{x}_2 \oplus (-1, 0, \varepsilon^{2,3}) &\geq (0, 0, -\varepsilon^{2,1}) \odot \mathsf{x}_1 \ . \end{aligned}$$

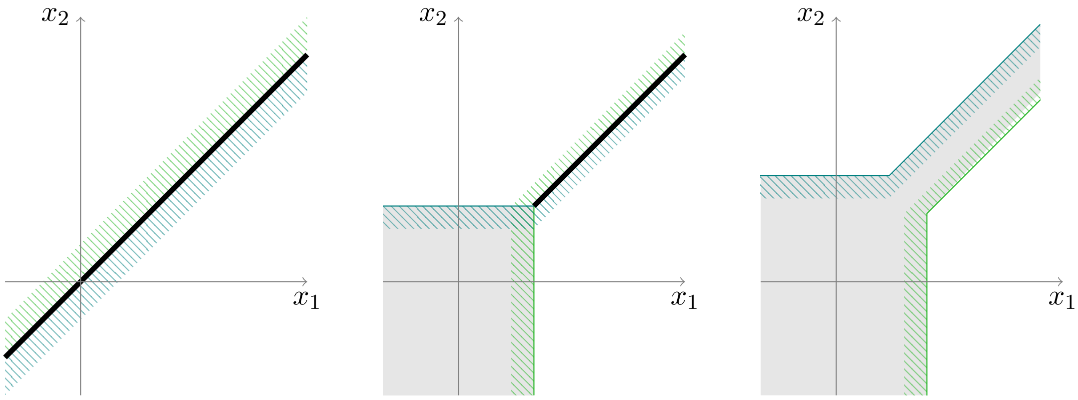

Figure 4.9: Illustration of the two perturbation steps on the polyhedron described in Example 4.39. Left: the original tropical polyhedron $\mathcal{P}(A, b)$, embedded into $\mathbb{I}^2$, which is the diagonal $x_1 = x_2$. Middle: the polyhedron obtained when the $\mathbb{0}$ entries of b have been replaced by the"infinitely small" scalar $(-1, 0, 0)$. Right: the polyhedron obtained after applying the perturbation map Φ_E.

Lemma 4.40. *Suppose that the elements of E are positive. Then, given any feasible point $x \in \mathbb{T}^n$ of LP, its canonical embedding $\mathsf{x} = \psi(x)$ is feasible for $\widetilde{LP}$.*

Proof. Clearly, $\widetilde{\mathrm{Id}} \odot \mathsf{x} \geq \mathbb{0}_{\mathbb{I}^n}$. For the other inequalities, it suffices to show that $\widetilde{A}_i^+ \odot \mathsf{x} \oplus \widetilde{b}_i^+ \geq \widetilde{A}_i^- \odot \mathsf{x} \oplus \widetilde{b}_i^-$ for $i \in [m]$. If $A_i^- \odot x \oplus b_i^- = \mathbb{0}_{\mathbb{T}}$, then we also have $\widetilde{A}_i^- \odot \mathsf{x} \oplus \widetilde{b}_i^- = \mathbb{0}_{\mathbb{I}}$. Since $\widetilde{A}_i^+ \odot \mathsf{x} \oplus \widetilde{b}_i^+ \geq \mathbb{0}_{\mathbb{I}}$, the inequality is satisfied. Otherwise $A_i^+ \odot x \oplus b_i^+ \geq A_i^- \odot x \oplus b_i^- > \mathbb{0}_{\mathbb{T}}$. In this case, we have

$$\begin{aligned}
\widetilde{A}_i^+ \odot \mathsf{x} \oplus \widetilde{b}_i^+ &= (0, A_i^+ \odot x \oplus b_i^+, \varepsilon^+) \\
\widetilde{A}_i^- \odot \mathsf{x} \oplus \widetilde{b}_i^- &= (0, A_i^- \odot x \oplus b_i^-, -\varepsilon^-)
\end{aligned}$$

where ε^+ and ε^- are sum of elements in E. Since the elements of E are positive, it follows that $\varepsilon^+ \geq 0 \geq -\varepsilon^-$. □

Let $\mathbb{I}^{\leq}$ denote the subset of $\mathbb{I}$ consisting of the elements $(f, g, h) \in F \times G \times H$ with $f \leq 0$, together with $\mathbb{0}_{\mathbb{I}}$. We project the elements of $\mathbb{I}^{\leq}$ to $\mathbb{T}$ with the map ρ, defined by $\rho(0, g, \cdot) = g$, and $\rho(f, \cdot, \cdot) = \mathbb{0}_{\mathbb{T}}$ for $f < 0$, along with $\rho(\mathbb{0}_{\mathbb{I}}) = \mathbb{0}_{\mathbb{T}}$. The map ρ is extended to vectors entry-wise.

Lemma 4.41. *Let $\mathsf{x} \in \mathbb{I}^n$ be a point with entries in $\mathbb{I}^{\leq}$. If x is feasible for $\widetilde{LP}$, then $\rho(\mathsf{x})$ is feasible for LP. Besides, if x is optimal for $\widetilde{LP}$, then $\rho(\mathsf{x})$ is optimal for LP.*

Proof. Observe that $\mathbb{I}^{\leq}$ is a subsemiring of $\mathbb{I}$, and that the coefficients defining $\widetilde{\mathrm{LP}}$ belong to $\mathbb{I}^{\leq}$. The lemma then follows from the fact that ρ is a homomorphism of semirings from $\mathbb{I}^{\leq}$ to $\mathbb{T}$ that preserves the order. □

Remark 4.42. The feasible points of $\widetilde{\mathrm{LP}}$ with some entries in $\mathbb{I} \setminus \mathbb{I}^{\leq}$ correspond to rays of the recession cone of $\mathcal{P}(A, b)$. Indeed, consider such a point x. Let $\lambda = -\max_j(\mathsf{x}_j)$ and let $\mathsf{r} = \lambda \odot \mathsf{x}$ be the point obtained by rescaling x by λ. Then, r has entries in $\mathbb{I}^{\leq}$, and satisfies:

$$\widetilde{A}^+ \odot \mathsf{r} \oplus (\lambda \odot \widetilde{b}^+) \geq \widetilde{A}^- \odot \mathsf{r} \oplus (\lambda \odot \widetilde{b}^-) \ .$$

Note that λ is of the form $(-f, \cdot, \cdot)$ for some positive f, whereas $\widetilde{b}^+$ and $\widetilde{b}^-$ have entries of the form $(0, \cdot, \cdot)$ or $\mathbb{0}_{\mathbb{I}}$. Hence, the image of both $\lambda \odot \widetilde{b}^+$ and $\lambda \odot \widetilde{b}^-$ under the projection map ρ is the vector with $\mathbb{0}_{\mathbb{T}}$ entries. It follows that $\rho(\mathsf{r})$ belongs to the polyhedral cone $\mathcal{P}(A, \mathbb{0})$, which is the recession cone of $\mathcal{P}(A, b)$.

Lemma 4.43. *Let I be a feasible basis of $\widetilde{LP}$. Then the basic point x^I have entries in $\mathbb{I}^{\leq}$.*

Proof. Let $(\widetilde{A}' \ \widetilde{d}')$ be the matrix defining the feasible set of $\widetilde{\mathrm{LP}}$, *i.e.*, $\widetilde{A}' = \begin{pmatrix} \widetilde{A} \\ \widetilde{Id} \end{pmatrix}$ and $\widetilde{d}' = \begin{pmatrix} b' \\ \mathbb{0}_{\mathbb{I}^n} \end{pmatrix}$, with $\widetilde{b}' \in \mathbb{I}_{\pm}^m$ being the vector such that $(\widetilde{b}')^+ = \widetilde{b}^+ \oplus \widetilde{d}$ and $(\widetilde{b}')^- = \widetilde{b}^-$. By Proposition 4.29, the components of $\mathsf{x}^I \in \mathbb{I}^n$ are given by

$$\mathsf{x}^I_j = |\operatorname{tdet}(\widetilde{A}'_{I,\widehat{j}} \ \widetilde{b}'_I)| - |\operatorname{tdet}(\widetilde{A}'_I)| \ ,$$

The entries of $|(\widetilde{A' \ b'})|$ belongs to $\mathbb{I}^{\leq}$, hence $|\operatorname{tdet}(\widetilde{A}'_{I,\widehat{j}} \ \widetilde{b}'_I)|$ is also in $\mathbb{I}^{\leq}$. Moreover, the entries of $|\widetilde{A}'|$ are either of the form $(0, \cdot, \cdot)$ or equal to $\mathbb{0}_{\mathbb{I}}$. Since $\operatorname{tdet}(\widetilde{A'}_I) \neq \mathbb{0}_{\mathbb{I}}$, we deduce that $|\operatorname{tdet}(\widetilde{A}'_I)|$ is an element of the form $(0, \cdot, \cdot)$. Hence x^I have entries in $\mathbb{I}^{\leq}$. $\square$

Proposition 4.44. *Suppose that $\widetilde{LP}$ is feasible and let I be a feasible basis of $\widetilde{LP}$. Then, the tropical simplex method, equipped with any tropical pivoting rule, and applied on the input $\begin{pmatrix} \widetilde{A} \\ \widetilde{Id} \end{pmatrix}, \begin{pmatrix} \widetilde{b} \\ \mathbb{0} \end{pmatrix}, \widetilde{c}$ and I, terminates and returns an optimal basis I^* of $\widetilde{LP}$. Let x^{I^*} be the corresponding basic point. Then, $\rho(\mathsf{x}^{I^*})$ is an optimal solution of LP.*

Proof. By Lemma 4.38, the matrix $\begin{pmatrix} \widetilde{A} & \widetilde{b} \\ \widetilde{\mathrm{Id}} & \mathbb{0} \\ \widetilde{c} & \mathbb{0} \end{pmatrix}$ is generic, and thus sign-generic, for the tropicalization of any polynomial. Moreover, $\widetilde{b}$ has non $\mathbb{0}$ entries. Hence, $\widetilde{\mathrm{LP}}$ is non-degenerate by Lemma 4.32(ii) Hence, the tropical simplex method terminates and returns an optimal basis of $\widetilde{\mathrm{LP}}$ by Proposition 4.36. By Lemmas 4.41 and 4.43, $\rho(\widetilde{x}^{I^*})$ is an optimal solution of LP. $\square$

4.4.1 Perturbation into a bounded polyhedron

It is sometimes convenient to obtain a tropical linear program whose feasible set is a bounded polyhedron and that contains no points with $\mathbb{0}$ entries. In particular, this assumption is needed to apply the implementation of tropical simplex method developped in Chapter 7. Hence, we shall add to the tropical linear program (4.40) an "infinitely small" lower bound $\mathsf{x}_j \geq \mathsf{l}_j$ for each variable $j \in [n]$. We require that

$$\mathsf{d}_j \gg \mathsf{l}_i > \mathbb{0}_{\mathbb{I}} \text{ for all } i \in [m] \text{ and } j \in [n] \ .$$

We also add an "infinitely big" upper bound constraint $\mathsf{e}^\top \odot \mathsf{x} \leq \mathsf{u}$, where $\mathsf{e} \in \mathbb{I}_+^n$ is the vector with all entries equal to $\mathbb{1}$, and $\mathsf{u} \gg \mathbb{1}_{\mathbb{I}}$. For example, we may use the parameters:

$$\begin{aligned} \mathsf{d}_1 &= \cdots = \mathsf{d}_m = (-1,0,0) \\ \mathsf{l}_1 &= \cdots = \mathsf{l}_n = (-2,0,0) \\ \mathsf{u} &= (1,0,0)\ . \end{aligned} \tag{4.42}$$

As before, we apply the perturbation map

$$\begin{pmatrix} \widetilde{c}^\top & \mathbb{0} & \mathbb{0} \\ \widetilde{A} & \widetilde{b} & \widetilde{d} \\ \widetilde{\mathrm{Id}} & \ominus\widetilde{l} & \mathbb{0} \\ \ominus\widetilde{e}^\top & \widetilde{u} & \mathbb{0} \end{pmatrix} = \Phi_E \begin{pmatrix} \mathsf{c}^\top & \mathbb{0} & \mathbb{0} \\ \psi(A) & \psi(b) & \mathsf{d} \\ \mathrm{Id} & \ominus\mathsf{l} & \mathbb{0} \\ \ominus\mathsf{e}^\top & \mathsf{u} & \mathbb{0} \end{pmatrix}\ , \tag{4.43}$$

and we denote by $\overline{\mathrm{LP}}$ the following linear program:

$$\begin{aligned} \text{maximize} \quad & \widetilde{c}^\top \odot \bar{\mathsf{x}} \\ \text{subject to} \quad & \widetilde{A}^+ \odot \widetilde{x} \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{A}^- \odot \widetilde{x} \oplus \widetilde{b}^- \\ & \widetilde{\mathrm{Id}} \odot \widetilde{x} \geq \widetilde{l} \\ & \widetilde{u} \geq \widetilde{e}^\top \odot \widetilde{x} \end{aligned} \tag{$\overline{\mathrm{LP}}$}$$

Lemma 4.45. *Suppose that $\mathsf{d}_i \gg \mathsf{l}_j$ for all $i,j \in [m] \times [n]$, that $\mathsf{u} \gg \mathbb{1}$, and that the elements of E are positive. Then, given any feasible point $x \in \mathbb{T}^n$ of LP, the point $\mathsf{x} \in \mathbb{I}^n$, defined by $\mathsf{x}_j = (0, x_j, 0)$ if $x_j \neq \mathbb{T}$ and $\mathsf{x}_j = \widetilde{l}_j$ otherwise, is feasible for $\overline{LP}$.*

Proof. Clearly, x satisfies $\widetilde{u} \geq \widetilde{e}^\top \odot \widetilde{x}$ and $\widetilde{\mathrm{Id}} \odot \widetilde{x} \geq \widetilde{l}$. Now consider an $i \in [m]$. If $A_i^- \odot x \oplus b_i^- = \mathbb{0}_{\mathbb{T}}$, then $\widetilde{A}_i^- \oplus \mathsf{x} \odot \widetilde{b}_i^-$ is of the form $\bigoplus_j \widetilde{A}_{ij}^- \odot \widetilde{l}_j$. Due to our conditions on l, d, it follows that $\widetilde{d}_i \geq \widetilde{A}_i^- \oplus \mathsf{x} \odot \widetilde{b}_i^-$. Otherwise, $A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^- > \mathbb{0}_{\mathbb{T}}$, and the proof of Lemma 4.40 readily applies. □

Lemma 4.46. *Let I be an optimal basis of $\overline{LP}$, and i_u the index of the inequality $\widetilde{u} \geq \widetilde{e} \odot \widetilde{x}$. If $i_u \notin I$, the basic point x^I has entries in $\mathbb{I}^{\leq}$, and $\rho(\mathsf{x}^I)$ is an optimal solution of LP. Otherwise, if $i_u \in I$, pivoting along the edge defined by $I \setminus \{i_u\}$ provides another basis I' with $i_u \notin I'$. Its basic point $\mathsf{x}^{I'}$ is also an optimal solution of $\overline{LP}$.*

Proof. The matrix defining the feasible set of $\overline{\mathrm{LP}}$ is of the form

$$\begin{pmatrix} \widetilde{A}' & \widetilde{d}' \\ \ominus\widetilde{e}^\top & \widetilde{u} \end{pmatrix}$$

where $\widetilde{A}'$ and $\widetilde{d}'$ have entries in $\mathbb{I}^{\leq}$. Suppose that $i_u \notin I$, then the Cramer's formulæ providing x^I involves minors of $(\widetilde{A}'\ \widetilde{d}')$ (as in Lemma 4.41), and thus x^I has entries in $\mathbb{I}^{\leq}$.

Otherwise, let $I = K \cup \{i_u\}$. Let us lift $\overline{\mathrm{LP}}$ to the linear program over Hahn series $\mathbb{R}[\![t^{F\times G\times H}]\!]$ provided by Proposition 4.20:

$$\begin{aligned} \text{maximize} \quad & \widetilde{\boldsymbol{c}}^\top \mathbf{x} \\ \text{subject to} \quad & \widetilde{\boldsymbol{A}}'\mathbf{x} + \widetilde{\boldsymbol{d}}' \geq 0 \\ & \widetilde{\boldsymbol{u}} \geq \widetilde{\boldsymbol{e}}^\top \mathbf{x} \end{aligned} \tag{4.44}$$

Let $\mathbf{x}^I \in \mathbb{K}^n$ be the basic point of (4.44) for the basis I, and $\mathbf{x}^{I'}$ the basic point obtained by pivoting along the edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_u\}}$. We claim that the reduced cost of the edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_u\}}$ is non-positive. This imply that $\widetilde{\boldsymbol{c}}^\top \mathbf{x}^{I'} \leq \widetilde{\boldsymbol{c}}^\top \mathbf{x}^I$. Moreover, I being an optimal basis, $\mathbf{x}^I$ is an optimal solution of (4.44). Consequently, $\widetilde{\boldsymbol{c}}^\top \mathbf{x}^{I'} = \widetilde{\boldsymbol{c}}^\top \mathbf{x}^I$ and thus $\widetilde{c}^\top \odot \mathsf{x}^{I'} = \widetilde{c}^\top \odot \mathsf{x}^I$ by applying the valuation map.

We now prove our claim. Let us denote by $\mathbf{z}$ the reduced cost of the edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_u\}}$, and $\mathbf{y}_k$ the reduced cost of $\boldsymbol{\mathcal{E}}_{I \setminus \{k\}}$ for $k \in I \setminus \{i_u\}$. Note that the Cramer's formulæ defining $\mathbf{y}$ and $\mathbf{z}$ involves only minors of $\begin{pmatrix} \widetilde{\boldsymbol{A}}' \\ \widetilde{\boldsymbol{e}}^\top \\ \widetilde{\boldsymbol{c}}^\top \end{pmatrix}$. It follows that the tropical reduced costs $\mathsf{y} = \operatorname{val}(\mathbf{y})$ and $\mathsf{z} = \operatorname{val}(\mathbf{z})$ have entries in $\mathbb{I}^{\leq}$.

Since I is an optimal basis, $\mathbf{x}^I$ is an optimal solution of (4.44), and $(\mathbf{y}, \mathbf{z})$ and optimal solution of the dual linear program. Consequently,

$$\widetilde{\boldsymbol{c}}^\top \mathbf{x}^I = -(\widetilde{\boldsymbol{d}}')^\top \mathbf{y} - \widetilde{\boldsymbol{u}}\mathbf{z} \tag{4.45}$$

by Theorem 3.7. Since $\widetilde{d}', \mathsf{y}$ and z have entries in $\mathbb{I}^{\leq}$, while $\widetilde{u} \gg \mathbb{1}_{\mathbb{I}}$, it follows that the leading term of the Hahn series (4.45) is given by the leading term of $-\widetilde{\boldsymbol{u}}\mathbf{z}$. Since $\widetilde{\boldsymbol{c}}^\top \mathbf{x}^{I'}$ is non-negative, we deduce that $-\widetilde{\boldsymbol{u}}\mathbf{z}$ is non-negative and thus that $\boldsymbol{z}$ is non-positive. □

4.4.2 Phase I

It remains to detect the feasibility of $\overline{\text{LP}}$. As usual, we use a Phase I method. We add a new variable λ to $\overline{\text{LP}}$ to measure the "infeasibility" of a point. The objective is now to minimize λ. To keep our linear program bounded, we also add upper and lower bound constraints on λ. Let $\delta \in \mathbb{I}^m$ be the unit vector of size m, and $\mathsf{l}_{n+1} \in \mathbb{I}$ a scalar such that $\mathbb{0} < \mathsf{l}_{n+1} \ll \mathsf{l}_j \ll \mathsf{d}_i$ for all $(i,j) \in [m] \times [n]$. If we choose $\mathsf{d}, \mathsf{l}_1, \dots, \mathsf{l}_n$ and u as in (4.42), then we can take

$$\mathsf{l}_{n+1} = (-3, 0, 0) \ . \tag{4.46}$$

Our Phase I linear program is:

$$\begin{aligned} \text{maximize} \quad & \widetilde{m} \odot \lambda \\ \text{subject to} \quad & \widetilde{A}^+ \odot \mathsf{x} \oplus \widetilde{\delta} \odot \lambda \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{A}^- \odot \mathsf{x} \oplus \widetilde{b}^- \\ & \widetilde{\mathrm{Id}} \odot \mathsf{x} \geq \widetilde{l} \\ & \widetilde{\mathbb{1}} \odot \lambda \geq \widetilde{l}_{n+1} \\ & \widetilde{u} \geq \widetilde{e}^\top \odot \mathsf{x} \oplus \widetilde{e}_{n+1} \odot \lambda \end{aligned} \tag{Phase I}$$

where:

$$\begin{pmatrix} \mathbb{0} & \mathbb{0} & \mathbb{0} & \widetilde{m} \\ \widetilde{A} & \widetilde{b} & \widetilde{d} & \widetilde{\delta} \\ \widetilde{\mathrm{Id}} & \ominus\widetilde{l} & \mathbb{0} & \mathbb{0} \\ \mathbb{0} & \ominus\widetilde{l}_{n+1} & \mathbb{0} & \widetilde{\mathbb{1}} \\ \ominus\widetilde{e}^\top & \widetilde{u} & \mathbb{0} & \ominus\widetilde{e}_{n+1} \end{pmatrix} = \varPhi_E \begin{pmatrix} \mathbb{0} & \mathbb{0} & \mathbb{0} & \mathbb{1} \\ \mathsf{A} & \mathsf{b} & \mathsf{d} & \delta \\ \mathrm{Id} & \ominus\mathsf{l} & \mathbb{0} & \mathbb{0} \\ \mathbb{0} & \ominus\mathsf{l}_{n+1} & \mathbb{0} & \mathbb{1} \\ \ominus\mathsf{e}^\top & \mathsf{u} & \mathbb{0} & \mathbb{1} \end{pmatrix} , \tag{4.47}$$

Observe that (4.43) and (4.47) define the same matrices $\widetilde{A}, \widetilde{b}, \widetilde{d}, \widetilde{\mathrm{Id}}$ and $\widetilde{l}$.

We have a feasible basis for Phase I.

Lemma 4.47. *The set I indexing the inequalities $\widetilde{Id} \odot \mathsf{x} \geq \widetilde{l}$ and $\widetilde{u} \geq \widetilde{e}^\top \odot \mathsf{x} \oplus \widetilde{e}_{n+1} \odot \lambda$ is a feasible basis of Phase I.*

Proof. Clearly $\operatorname{tdet}\begin{pmatrix} \widetilde{\mathrm{Id}} & \mathbb{0} \\ \widetilde{e}^\top & \widetilde{e}_{n+1} \end{pmatrix} \neq \mathbb{0}$. Thus it is sufficient to show that the unique solution (x, λ) of the system $\widetilde{\mathrm{Id}} \odot \mathsf{x} = \widetilde{l}$ and $\widetilde{u} = \widetilde{e}^\top \odot \mathsf{x} \oplus \widetilde{e}_{n+1} \odot \lambda$ is feasible for Phase I. Due to our assumption $\mathsf{u} \gg \mathbb{1} \gg \mathsf{l}_j$, it follows that $\lambda \gg \mathbb{1} \gg \mathsf{x}_j$ for all $j \in [n]$. Consequently, $\widetilde{\mathbb{1}} \odot \lambda \geq \widetilde{l}_{n+1}$ and

$$\widetilde{A}^+ \odot \mathsf{x} \oplus \widetilde{\delta} \odot \lambda \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{\delta} \odot \lambda \gg \widetilde{A}^- \mathsf{x} \oplus \widetilde{b}^- \ . \qquad \square$$

Lemma 4.48. *Let I be an optimal basis of Phase I and i_l the index of the inequality $\widetilde{\mathbb{1}} \odot \lambda \geq \widetilde{l}_{n+1}$. Either $i_l \in I$ and $I \setminus \{i_l\}$ is a feasible basis for $\overline{LP}$, or $i_l \notin I$ and $\overline{LP}$ is infeasible.*

Proof. Let $(\mathsf{x}^I, \lambda^I)$ be the basic point of an optimal basis I. First, consider the case $i_l \notin I$. Since Phase I is non-degenerate by Lemma 4.32, we have the strict inequality $\widetilde{\mathbb{1}} \odot \lambda^I > \widetilde{l}_{n+1}$. Hence, the optimal value of Phase I is $\widetilde{m} \odot \lambda^I$, and satisfy:

$$\widetilde{m} \odot \lambda^I > \widetilde{m} \odot \widetilde{l}_{n+1} \odot \widetilde{\mathbb{1}}^{\odot -1} \ . \tag{4.48}$$

By contradiction, suppose that $\overline{\mathrm{LP}}$ admits a feasible point x. Let $\lambda = \widetilde{l}_{n+1} \odot (\widetilde{\mathbb{1}})^{\odot -1}$. We have:

$$\widetilde{A}^+ \odot \mathsf{x} \oplus \widetilde{\delta} \odot \lambda \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{A}^+ \odot \mathsf{x} \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \geq \widetilde{A}^- \mathsf{x} \oplus \widetilde{b}^- \ .$$

Furthermore, $\widetilde{e}_{n+1} \odot \lambda \leq \widetilde{u}$ as $\mathsf{l}_{n+1} \ll \mathbb{1} \ll \mathsf{u}$. Consequently, the point (x, λ) is feasible for Phase I. At this point, the value of the objective function of Phase I is $\widetilde{m} \odot \widetilde{l}_{n+1} \odot (\widetilde{\mathbb{1}})^{\odot -1}$. Using (4.48), this contradicts the optimality of $(\mathsf{x}^I, \lambda^I)$, and thus $\overline{\mathrm{LP}}$ is infeasible.

Second, assume that $i_l \in I$. Then $\lambda^I = \widetilde{l}_{n+1} \odot (\widetilde{\mathbb{1}})^{\odot -1}$. Since $\mathsf{l}_{n+1} \ll \mathsf{d}_i \ll \mathbb{1}$ for all $i \in [m]$, it follows that $\widetilde{\delta} \odot \lambda^I \ll \widetilde{d}$. We obtain:

$$\widetilde{A}^+ \odot \mathsf{x}^I \oplus (\widetilde{b}^+ \oplus \widetilde{d}) = \widetilde{A}^+ \odot \mathsf{x} \oplus \widetilde{\delta} \odot \lambda^I \oplus (\widetilde{b}^+ \oplus \widetilde{d}) \ .$$

As $(\mathsf{x}^I, \lambda^I)$ is feasible for Phase I, it follows that

$$\widetilde{A}_i^+ \odot \mathsf{x}^I \oplus (\widetilde{b}_i^+ \oplus \widetilde{d}_i) \geq \widetilde{A}_i^- \mathsf{x}^I \oplus \widetilde{b}_i^- \text{ for all } i \in [m] \ , \tag{4.49}$$

and the inequality (4.49) holds with equality for $i \in I \cap [m]$. Clearly,

$$\widetilde{u} \geq \widetilde{e}^{\top} \odot \mathsf{x}^I \oplus \widetilde{e}_{n+1} \odot \lambda^I \geq \widetilde{e}^{\top} \odot \mathsf{x}^I \ .$$

Moreover, if I indexes the latter inequality, then we must have $\widetilde{u} = \widetilde{e}^{\top} \odot \mathsf{x}^I$ as $\widetilde{e}_{n+1} \odot \lambda^I \ll \widetilde{u}$. Obviously, the inequalities $\widetilde{\mathrm{Id}} \odot \mathsf{x}^I \geq \widetilde{l}$ are satisfied, and holds with equality when indexed by I.

We have shown that x^I is feasible for $\overline{\mathrm{LP}}$ and that it activates the inequalities indexed by $I \setminus \{i_l\}$. It remains to show that the corresponding submatrix has a non $\mathbb{0}$ tropical determinant. Denote $\widetilde{A}' = \left(\begin{smallmatrix} \widetilde{A} \\ \mathrm{Id} \end{smallmatrix}\right)$ and $\widetilde{\delta}' = \left(\begin{smallmatrix} \widetilde{\delta} \\ \mathbb{0}_{\mathbb{I}^n} \end{smallmatrix}\right)$. Since I is a basis of Phase I, we have $\operatorname{tdet}\begin{pmatrix} \widetilde{A}'_{I\setminus\{i_l\}} & \widetilde{\delta}' \\ \mathbb{0} & \widetilde{\mathbb{1}} \end{pmatrix} \neq \mathbb{0}$. Consequently, $\operatorname{tdet}(\widetilde{A}'_{I\setminus\{i_l\}}) \neq \mathbb{0}$. It follows that, $I \setminus \{i_l\}$ is a feasible basis for $\overline{\mathrm{LP}}$. □

Theorem 4.49 (Tropical simplex method for arbitrary tropical linear programs)**.** *An arbitrary tropical linear program $LP(A, b, c)$ is solved by the following algorithm:*

- *Apply the tropical simplex method to the tropical linear program Phase I, starting with the feasible basis of Lemma 4.47. Let I be the optimal basis of Phase I returned by the algorithm*
- *If $i_l \notin I$, then $LP(A, b, c)$ is infeasible by Lemma 4.48.*
- *Otherwise, apply the tropical simplex method to $\overline{LP}$ with $I \setminus \{i_l\}$ as an initial basis.*
- *Let I^* be the optimal basis of $\overline{LP}$ obtained, possibly after the last pivoting step of Lemma 4.46.*
- *Compute the basic point x^{I^*} of $\overline{LP}$ using Proposition 4.29.*
- *The projection $\rho(\mathsf{x}^{I^*})$ is an optimal solution of $LP(A, b, c)$.*

Remark 4.50. Since the matrix in (4.47) is of size $(m + n + 3) \times (n + 3)$, we use only $(m + n + 3)(n + 3)$ elements of E to obtained the perturbed matrix. Hence, we can use $H = \mathbb{Z}^{(m+n+3)\times(n+3)}$ as a perturbation group, and the canonical basis of H for elements of E. Using the parameters proposed in (4.46) and (4.42), the non $\mathbb{0}$ entries of the matrices (4.47) and (4.43) are of the form (f_i, g_i, h_i), where $|f_i| \leq 3$, the element g_i is either 0 or an entry of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$, and $h_i \in H$ is an element of the basis E of H. Hence the input size of f_i is $O(1)$ and the input size of h_i is $O(mn)$. Consequently, the input size of Phase I and $\overline{\mathrm{LP}}$ are polynomial in the input size of $\mathrm{LP}(A, b, c)$.

Chapter 5

Relations between the complexity of classical and tropical linear programming via the simplex method

In this chapter, we present three results related to the complexity of the simplex method. First, in Section 5.1, we prove that the existence of a pivoting rule which performs a strongly polynomial number of iterations on linear programs over $\mathbb{R}$ would provide a polynomial algorithm for tropical linear programming, and thus mean payoff games. Second, in Section 5.2, we show that if a pivoting rule, used on a tropical linear program, performs a number of iterations which is polynomial in the input size of the tropical entries, the number of iterations is in fact strongly polynomial (*i.e.*, polynomial in the dimensions of the problem). Last, in Section 5.3, we exhibit a class of classical linear programs on which the simplex method, with any pivoting rule, performs a number of iterations which is polynomial in the input size of the problem. Consequently, the corresponding polyhedra have a diameter which is polynomial in the input size.

These three results are based on the following idea. We have seen in Section 3.3 that the simplex method can be implemented using only the signs of polynomials evaluated on the problem to be solved. This also provides the following observation.

Proposition 5.1. *Consider a semi-algebraic pivoting rule ϕ defined by the polynomials $(P_i^{\phi})_i$. The sequence of bases produced by the simplex method applied to a linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ depends only on the signs of the minors of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ and the signs of $P_i^{\phi}\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$.*

We call this collection of signs the *sign pattern* of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, and we denote it by $s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. Any linear program with a sign pattern s is called a *realization* of s.

If the simplex method performs L iterations on an instance $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, then the number of iterations is also equal to L on any realization of the sign pattern $s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, including realizations on other ordered fields. Our first result, in Section 5.1, comes from the fact that the sign pattern of linear program over Hahn series is realizable over the real numbers, by completeness of the theory of real-closed fields.

For a generic tropical linear program, we can also define a sign-pattern, which governs the behavior of the simplex method. In Section 5.2, we show that the tropical realization space of a sign pattern is a semi-linear set. Using simultaneous diophantine approximation, it follows that the sign-pattern of a tropical linear program always have a "short" realization, *i.e.*, with an input size which is polynomial in the dimensions. Consequently, an algorithm which is polynomial in the bit model is in fact strongly polynomial.

Finally, in 5.3, we construct linear programs over $\mathbb{Q}$ which realize the sign-pattern of a tropical instance, and whose input sizes are greater than the *values* of tropical input. Hence, if the simplex method is pseudo-polynomial for the tropical instance, it is polynomial with respect to the input size of the classical instances.

The transfer of complexity from classical to tropical linear programming (Theorem 5.3 below) appeared in [ABGJ13a] in a less general form (restricted to *combinatorial* pivoting rules). The other contents of this chapter are original.

5.1 From classical to tropical linear programming

Let $N_K(n, m, \phi)$ be the maximal length of a run of the simplex method, equipped with a semi-algebraic pivoting rule ϕ, for a non-degenerate classical linear programs of size (n, m), with coefficients in a real closed field K.

Similarly, let $N_{\mathbb{T}}(n, m, \phi^{\mathbb{T}})$ the maximal length of a run of the tropical simplex algorithm, equipped with the tropical rule $\phi^{\mathbb{T}}$, for a tropical linear program satisfying the conditions of Proposition 4.36, with coefficients in a tropical semiring $\mathbb{T} = \mathbb{T}(G)$.

Proposition 5.2. *Let G be a totally ordered abelian group. Then,*

$$N_{\mathbb{T}(G)}(n, m, \phi^{\mathbb{T}}) \leq N_{\mathbb{R}}(n, m, \phi) \ .$$

Proof. Let $\mathrm{LP}(A, b, c)$ be a non-degenerate tropical linear program, with coefficients in the semiring $\mathbb{T}(G)$, satisfying the conditions of Proposition 4.36, and I a feasible basis of this problem. By Theorem 3.25, the number of iterations of the tropical simplex method applied to A, b, c, I is exactly the number of iterations of the classical simplex method over Hahn series applied to $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}, I$ for any $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right) \in \mathrm{sval}^{-1} \left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$.

By Proposition 5.1, the number of iterations of the classical simplex algorithm depends only on I and the sign pattern $s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. We claim that the sign pattern $s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is realizable over the real numbers, *i.e.*, we affirm that there exist $\mathsf{A} \in \mathbb{R}^{m \times n}, \mathsf{b} \in \mathbb{R}^m$ and $\mathsf{c} \in \mathbb{R}^n$ such that $s^{\phi}(\mathsf{A}, \mathsf{b}, \mathsf{c}) = s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. Indeed, let $P_1, \dots, P_r$ be the polynomials defining the sign-pattern. Observe that the realizability of a sign pattern $s \in \{-1, 0, +1\}^r$ by a (n, m) linear program over an ordered field K can be expressed as the following sentence in the language $\mathcal{L}_{\text{or}}$:

$$\exists \left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right) \in K^{(m+1)\times(n+1)} \text{ s.t}$$
$$\Bigg(\bigwedge_{\substack{i\in[r] \\ s_i=+1}} P_i \left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right) > 0 \Bigg) \wedge \Bigg(\bigwedge_{\substack{i\in[r] \\ s_i=0}} P_i \left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right) = 0 \Bigg) \wedge \Bigg(\bigwedge_{\substack{i\in[r] \\ s_i=-1}} P_i \left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right) < 0 \Bigg) . \tag{5.1}$$

Up to embedding, we can always assume that the abelian group G is divisible (Theorem 2.9). In this case, the field of Hahn series $\mathbb{R}[\![t^G]\!]$ is real closed (Theorem 2.6). The sentence (5.1) holds on $\mathbb{R}[\![t^G]\!]$. Since the theory of real-closed fields is complete (Theorem 2.2), and $\mathbb{R}$ is a real-closed field, we conclude that (5.1) also holds on $\mathbb{R}$.

It remains to show that the realization $\left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right)$ of the sign pattern over $\mathbb{R}$ does provide a non-degenerate linear program. Since the tropical linear program $\mathrm{LP}(A, b, c)$ is non-degenerate, so is the Hahn linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ by Corollary 4.22. By Lemma 3.17, the non-degeneracy property is entirely determined by the signs of the minors of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. Since the signs of the minors are part of the sign pattern $s^{\phi}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, and since $\left(\begin{smallmatrix} \mathsf{A} & \mathsf{b} \\ \mathsf{c} & 0 \end{smallmatrix}\right)$ realizes this sign pattern, we conclude that $\mathbf{LP}(\mathsf{A}, \mathsf{b}, \mathsf{c})$ is non-degenerate. □

Theorem 5.3. *Let ϕ be a tropically tractable pivoting rule. Suppose that the simplex method, equipped with ϕ, performs a number of iterations which is polynomial in m and n on all non-degenerate linear programs over $\mathbb{R}$ defined m inequalities in dimension n. Then any tropical linear program can be solved in polynomial time.*

Proof. Let $\mathbb{T} = \mathbb{T}(G)$. Consider a tropical linear program $\mathrm{LP}(A, b, c)$ with $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right) \in \mathbb{T}^{(m+1)\times(n+1)}$. We construct the problems Phase I and $\overline{\mathrm{LP}}$ as in Section 4.4. Applying the tropical simplex method succesively to these two problems solves $\mathrm{LP}(A, b, c)$ by Theorem 4.49. The problem Phase I is described by $m + n + 2$ inequalities in dimension $n + 1$, and $\overline{\mathrm{LP}}$ by $m + n + 1$ inequalities in dimension n. Hence, the two calls to the tropical simplex method performs a total of $N_{\mathbb{T}}(n+1, m+n+3, \phi^{\mathbb{T}}) + N_{\mathbb{T}}(n, m+n+2, \phi^{\mathbb{T}})$ iterations. By Proposition 5.2, this number of iterations is smaller than $N_{\mathbb{R}}(n + 1, m + n+2, \phi^{\mathbb{T}}) + N_{\mathbb{R}}(n, m+n+1, \phi^{\mathbb{T}})$, and the latter is a polynomial in n and m by hypothesis.

The input sizes of Phase I and $\overline{\mathrm{LP}}$ are polynomial in the input size of A, b, c (see Remark 4.50). Then, by Theorem 3.25, each iteration of the tropical simplex method on these problems takes a time polynomial in the input size of A, b, c when ϕ is tropically tractable. □

5.2 A weakly polynomial tropical pivoting rule in fact performs a strongly polynomial number of iterations

Given a semi-algebraic pivoting rule ϕ, we can also define the sign pattern $s^{\phi}(A, b, c)$ for a tropical linear program $\mathrm{LP}(A, b, c)$ if the matrix $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ is sign-generic for the minor polynomials and the polynomials defining ϕ. Indeed, for any of these polynomials, $\mathrm{trop}(P)\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ is well-defined (see Section 3.2), and thus has a tropical sign. Since the execution of the tropical simplex method depends only on this sign pattern, we can also look at other tropical realizations of the sign pattern to obtain bounds on the complexity of the (tropical) simplex method. With this approach, we shall prove the following theorem.

Theorem 5.4. *Let ϕ be a semi-algebraic pivoting rule. Suppose that:*

- *for any tropical linear program $LP(A,b,c)$ with coefficients in the semiring $\mathbb{T}(\mathbb{Q})$ satisfying the conditions of Proposition 4.36, the tropical simplex method, equipped with $\phi^{\mathbb{T}}$, performs a number of iterations which is polynomial in the input size of A,b,c;*
- *for every polynomial $\sum_{\alpha\in\Lambda} q_\alpha X^\alpha \in \mathbb{Q}[X_1,\dots,X_l]$ involved in the definition of ϕ, the Newton polytope $\operatorname{conv}(\Lambda)$ is contained in a L_∞-ball of radius R, where the input size of R is a polynomial in l.*

Then, the tropical simplex method, equipped with $\phi^{\mathbb{T}}$, in fact performs a strongly *polynomial number of iterations on all tropical linear programs, i.e., $N_{\mathbb{T}(G)}(n,m,\phi^{\mathbb{T}})$ is a polynomial in m and n only for any ordered abelian group G.*

Let us first describe the set of tropical realizations of a sign pattern. Let $\mathrm{LP}(A,b,c)$ be a tropical linear program satisfying the conditions of Proposition 4.36 with entries in an arbitrary tropical semiring $\mathbb{T}=\mathbb{T}(G)$, and let s be its sign pattern. Observe that the set of polynomials defining the sign pattern includes all minor polynomials. In particular, this includes the 1×1 minors. As a consequence, for any tropical realization $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)$ of the sign pattern s, the tropical signs of the entries of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)$ are identical. To alleviate the notation, let us consider $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)$ as a vector $\delta\in\mathbb{T}_\pm^l$. By the discussion above, the sign pattern fixes the sign of the entries δ. Consequently, we can identify a realization δ of the sign pattern with a vector consisting of the modulus of the non $\mathbb{0}$ entries of δ, *i.e.*, with a vector $w\in G^k$, where k is the number of non $\mathbb{0}$ entries of δ, and the components of w are exactly the $|\delta_i|$ for the $i\in[l]$ such that $\delta_i\neq\mathbb{0}$. Hence, we identify the set of tropical realizations of a sign pattern with a subset of G^k.

Lemma 5.5. *The set of tropical realizations of a sign pattern s is the union of a finite number of classical convex cones of G^k. Each of this cone is of the form*

$$\{w\in G^k \mid Mw\geq 0,\ M'w>0\} \tag{5.2}$$

for some matrices M,M' with integer entries. Moreveor, each entry of M,M' has an absolute value bounded by $2R$, where R is the radius of a L_∞-ball containing the Newton polytopes of all polynomials defining the sign pattern.

Proof. Let $\delta\in\mathbb{T}_\pm^l$ be a realization of the sign pattern s, and $P=\sum_{\alpha\in\Lambda} q_\alpha X^\alpha\in\mathbb{Q}[X_1,\dots,X_l]$ a polynomial involved in the definition of the sign pattern.

First suppose that $\operatorname{trop}(P)(\delta)=\mathbb{0}$. Then, for every $\alpha\in\Lambda$, there exists a $i\in[l]$ with $\delta_i=\mathbb{0}$ and $\alpha_i>0$. Consequently, for every $\delta'\in\mathbb{T}_\pm^l$ such that $\delta'_i=\mathbb{0}$ when $\delta_i=0$, we have $\operatorname{trop}(P)(\delta')=\mathbb{0}$. In other words, if we restrict the sign pattern to the signs of the entries of δ and the sign of $\operatorname{trop}(P)(\delta)$, the realization space is G^k.

Second, suppose that $\mathrm{trop}(P)(\delta) \neq \mathbb{0}$. Up to replacing P by $-P$ we can assume that $\mathrm{trop}(P)(\delta)$ is tropically positive. Let us define Λ^+, Λ^- as follows:

$$\Lambda^+ := \{\alpha \in \Lambda \mid \mathrm{sign}(q_\alpha) \prod_{i \in [l]} \mathrm{sign}(\delta_i)^{\alpha_i} = +1\}$$
$$\Lambda^- := \{\alpha \in \Lambda \mid \mathrm{sign}(q_\alpha) \prod_{i \in [l]} \mathrm{sign}(\delta_i)^{\alpha_i} = -1\} \ .$$

Since $\mathrm{trop}(P)(\delta)$ is tropically positive, the maximum in $\mathrm{trop}(P)(\delta)$ must be attained only on exponents $\alpha \in \Lambda^+$. Consequently, the modulus of the non $\mathbb{0}$ entries of δ, must satisfy the follwing inequality:

$$\max_{\alpha \in \Lambda^+} \sum_{i | \delta_i \neq \mathbb{0}} \alpha_i |\delta_i| > \max_{\alpha \in \Lambda^-} \sum_{i | \delta_i \neq \mathbb{0}} \alpha_i |\delta_i| \ . \tag{5.3}$$

Conversely, if $\delta' \in \mathbb{T}_\pm$ satisfies (5.3) and have the same $\mathbb{0}$ entries as δ, then δ' is sign-generic for P and $\mathrm{trop}(P)(\delta')$ is tropically positive.

It follows that the realization space of a sign pattern is described by a finite number of inequalities of the form (5.3). Selecting a maximizing term in the left-hand side of each of these inequalities provides a cone of the form (5.2). □

As the set of tropical realizations of a sign pattern is described by linear inequalities, we shall see that we can always find a realization on $\mathbb{T}(\mathbb{Q})$ with a "short" input size, *i.e.*, an input size which is polynomial in m and n. The key tool is simultaneous diophantine approximation. More precisely, we shall use the following result of Frank and Tardos.

Theorem 5.6 ([FT87, Theorem 3.3])**.** *For any rational vector $w \in \mathbb{Q}^l$ and any integer R, there exists an integral vector $\bar{w} \in \mathbb{N}^l$ such that $\|\bar{w}\|_\infty \leq 2^{4l^3} R^{l(l+2)}$ and $\mathrm{sign}(\alpha^\top w) = \mathrm{sign}(\alpha^\top \bar{w})$ for any integral vector $\alpha \in \mathbb{N}^l$ with $\|\alpha\|_1 \leq R-1$.*

We now have all the ingredients to prove our theorem.

Proof of Theorem 5.4. Let s be the sign pattern of a tropical linear program satisfying the conditions of Proposition 4.36 with entries in an arbitrary tropical semiring $\mathbb{T} = \mathbb{T}(G)$. By Lemma 5.5, the realization space of s can be described as a disjunction of conjonctions of linear inequalities with integer coefficients. Consequently, there exists a first-order formulæ $\phi_s(A, b, c)$ in the language of ordered groups $\{<, +, 0\}$ that holds true if and only if $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ is a realization of s. By hypothesis, the sign pattern s is realizable on an ordered abelian group G. Since we can always embed G in a divisible group (Theorem 2.9), it follows that the first-order sentence $\exists A, b, c\ \phi_s(A, b, c)$ holds true in a ordered abelian divisible group. Since $\mathbb{Q}$ is an ordered abelian divisible group, the sign pattern s is realizable on $\mathbb{T}(\mathbb{Q})$ by Theorem 2.1.

Since the sign pattern s is realizable on $\mathbb{T}(\mathbb{Q})$, Theorem 5.6 and Lemma 5.5 tell us that it is realizable by a matrix $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ with entries that have an input size bounded by $O(l^3 + l^2 \log(R))$. By hypothesis, $\log(R)$ is a polynomial in l. For a (n, m) linear program, we consider polynomials on $l = (m+1)(n+1)$ variables. Hence, the sign-pattern can be realized on $\mathbb{T}(\mathbb{Q})$ by a matrix $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ with an input size that is polynomial in m and n. By hypothesis, the tropical simplex method applied to $\mathrm{LP}(A, b, c)$ performs a number of iterations that is polynomial in the input size of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$. Consequently, when applied to the latter instance, the tropical simplex method performs a number of iterations that is polynomial in m and n. Since the number of iterations of the tropical simplex method depends only of the sign pattern s, the number of iterations is still a polynomial in m and n on any realization of the sign pattern s. In particular, this holds for a realization on any tropical semiring $\mathbb{T}(G)$. □

5.3 From tropical to classical linear programming

We now exhibit a class of real linear programs on which the number of iterations of the simplex method is polynomial in the input size, regardless of the pivoting rule. Consequently, the corresponding polyhedra have a diameter which is polynomial in the input size. The idea is to consider the tropical linear programs on which the tropical simplex method is pseudo-polynomial. These instances are *quantized* into real linear programs. The quantized linear programs are combinatorially equivalent to the tropical one, and the values of the tropical input is a lower bound on the input size of the quantized programs.

5.3.1 Edge-improving tropical linear programs

We say that a tropical linear program $\mathrm{LP}(A, b, c)$ is *edge-improving* if it satisfies the conditions of Proposition 4.36, and for any pair of adjacent basic points $x^I, x^{I'}$, the objective values $c^\top \odot x^I$ and $c^\top \odot x^{I'}$ are distinct.

Lemma 5.7. *Let $LP(A, b, c)$ be an edge-improving tropical linear program on n variables with entries in $\mathbb{T}(\mathbb{Z})$. Suppose that the non $\mathbb{0}$ entries of $\left|\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)\right|$ belongs to the interval $[-v, v] \subseteq \mathbb{Z}$. Then, the tropical simplex method, equipped with any pivoting rule, performs at most $O(nv)$ iterations on $LP(A, b, c)$.*

Proof. Let x^I be a basic point. By Proposition 4.29, the components of x^I are of the form

$$x^I_j = |\operatorname{tdet}(A_{I,\widehat{j}}\ b_I)| - |\operatorname{tdet}(A_I)| \ .$$

The matrices $(A_{I,\widehat{j}}\ b_I)$ and A_I are of size $n \times n$. Since the non $\mathbb{0}$ entries of $\left|\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)\right|$ are integers in the interval $[-v, v]$, the non $\mathbb{0}$ components of x^I satisfies $-2nv \leq x^I_j \leq 2nv$. Consequently, $-(2n+1)v \leq c^\top \odot x^I \leq (2n+1)v$.

Suppose that the simplex method starts at the basis I^1. Let I^N be the last basis visited such that $c^\top \odot x^{I^N} > \mathbb{0}$. Since the tropical linear program is edge-improving, I^N is either the last basis visited, or the basis preceding the last. The difference between $c^\top \odot x^{I^1}$ and $c^\top \odot x^{I^N}$ is bounded by $(4n+2)v$. Moreover, $c^\top \odot x^I$ is an integer for any basis I. Since the linear program is edge-improving, $c^\top \odot x^I$ and $c^\top \odot x^{I'}$ differ by at least 1 for any two adjacent bases. □

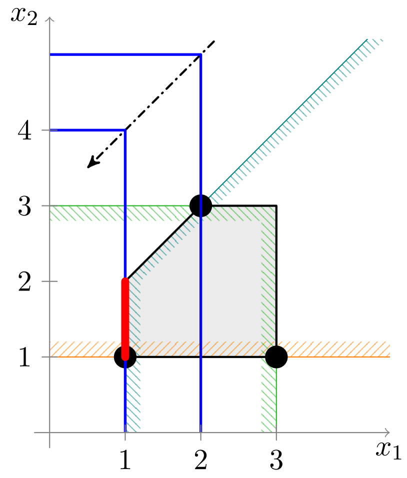

Figure 5.1: Illustration of the edge-improving tropical linear program (5.4). The set of optimal solutions is the segment between $(1, 1)$ and $(1, 2)$. In particular, the tropical basic point $(1, 1)$ is the unique optimal basic point. The basic point/edge graph of this tropical linear program, oriented by the signs of the reduced costs, coincides with the oriented graph of non edge-improving linear program of Example 4.35.

Remark 5.8. Note that genericity for the minor polynomials is not sufficient to ensure an improvement along an edge. In particular, the tropical linear program in Example 4.35 is generic but not edge-improving. Moreover, even under small perturbations of its input, this tropical linear program does not become edge-improving. Hence, the set of edge-improving tropical linear programs is not of measure 0.

However, consider the following tropical linear program, depicted in Figure 5.1.

$$\text{minimize } \max(x_1, x_2 - 4) \;\text{ s.t. }\; 3 \geq \max(x_1, x_2),\; x_1 \geq \max(1, x_2 - 1),\; x_2 \geq 1 \;. \tag{5.4}$$

This problem is edge-improving. Observe that the graph formed by its basic points and edges, and oriented by the signs of the reduced costs, is the same as in Example 4.35.

5.3.2 Quantized linear programs

In the rest of this section, $\mathrm{LP}(A, b, c)$ is a tropical program which satisfies the conditions of Proposition 4.36, and such that the non $\mathbb{0}$ entries of $\left|\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)\right|$ are non-negative integers smaller than v. We consider the sign pattern $\mathsf{signMinors}(A, b, c)$ that consists of the signs of the minors of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)$. We now construct a set of classical linear programs, with entries in $\mathbb{R}$, that realize the sign pattern $\mathsf{signMinors}(A, b, c)$. The idea is to lift $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0}\end{smallmatrix}\right)$ to a matrix $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0\end{smallmatrix}\right)$ whose coefficients are real-valued functions in the variable t (*e.g.*, polynomial functions or rational functions). This provides a family $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ of real linear programs. We say that a real linear program $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ obtained in this way is a *quantization* of $\mathrm{LP}(A, b, c)$ if:

- $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ realizes the sign pattern $\mathsf{signMinors}(A, b, c)$ of the tropical linear program (*i.e.*, the real and the tropical polyhedra are combinatorially equivalent);
- the input size of $\begin{pmatrix} \boldsymbol{A}(t) & \boldsymbol{b}(t) \\ \boldsymbol{c}(t) & 0 \end{pmatrix}$ is greater than v.

Theorem 5.9. *On a quantization of an edge-improving tropical linear program, the classical simplex method, equipped with any pivoting rule, performs a number of iterations which is polynomial in the input size of the problem.*

Proof. Let $\mathrm{LP}(A, b, c)$ be an edge-improving tropical linear program. Let $I^1, \dots, I^N$ be the sequence of bases produced by the simplex method on a quantization of $\mathrm{LP}(A, b, c)$, for a certain pivoting rule (recall that we assume that a pivoting rule always return a leaving index with a negative reduced cost). Since a quantization is combinatorially equivalent to $\mathrm{LP}(A, b, c)$, the sequence $I^1, \dots, I^N$ is a sequence of adjacent bases of $\mathrm{LP}(A, b, c)$ with edges of negative reduced cost between them. By Lemma 5.7, it follows that $N = O(nv)$. Since the input size of the entries of a quantized problem is greater than v, this proves the result. $\square$

We now construct quantizations of a tropical linear program $\mathrm{LP}(A, b, c)$.

Proposition 5.10. *Let $LP(A, b, c)$ be an edge-improving tropical linear program. Suppose that the non $\mathbb{0}$ entries of $|\begin{pmatrix} A & b \\ c & \mathbb{0} \end{pmatrix}|$ are non-negative integers, and let v be the largest entry of $|\begin{pmatrix} A & b \\ c & \mathbb{0} \end{pmatrix}|$. Consider any lift $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{pmatrix} \in \mathrm{sval}^{-1} \begin{pmatrix} A & b \\ c & \mathbb{0} \end{pmatrix}$ such that the entries of $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{pmatrix}$ are polynomial real-valued functions of the form:*

$$t \mapsto \pm \left(\sum_{k=0}^{v} q_k t^k \right) , \tag{5.5}$$

where the q_k are non-negative integers. For any rational number $t \geq 2$, the rational linear program $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ have an input size which is greater than v.

Proof. By assumption, there exists an entry of $|\begin{pmatrix} A & b \\ c & \mathbb{0} \end{pmatrix}|$ which is equal to v. The corresponding entry of $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{pmatrix}$ is the form $\pm(\sum_{k=0}^{v} q_k t^k)$ with $q_v \neq 0$. For any rational $t \geq 2$, the input size of the corresponding entry of the rational matrix $\begin{pmatrix} \boldsymbol{A}(t) & \boldsymbol{b}(t) \\ \boldsymbol{c}(t) & 0 \end{pmatrix}$ is greater than

$$\log_2 \left(\sum_{k=0}^{v} q_k t^k \right) = v \log_2(t) + \log_2(q_v) + \log_2 \left(1 + \sum_{k=0}^{v-1} \frac{q_k}{q_{v'}} t^{k-v} \right) . \tag{5.6}$$

Since the coefficients q_k are non-negative integers, and $t \geq 2$, the expression in (5.6) is greater than v. Since the input size of $\begin{pmatrix} \boldsymbol{A}(t) & \boldsymbol{b}(t) \\ \boldsymbol{c}(t) & 0 \end{pmatrix}$ is greater than the sum of the input sizes of its entries, the result follows. $\square$

Remark 5.11. The coefficients q_k in (5.5) are restricted to be non-negative integers only to easily relate the input size of the quantized problem with v. One can clearly obtain quantizations when the q_k are rational numbers. However in that case, one may need values of t that are larger than 2. Instead of lifting the tropical entries to polynomial functions, one could also consider more general real-valued functions, such as rational functions.

With a lift as in Proposition 5.10, if t is chosen large enough, we will always realize the sign-pattern of the tropical linear program.

Proposition 5.12. *Let $LP(A, b, c)$ be a tropical linear program on n variables with $m \geq n$ constraints with entries in $\mathbb{T}(\mathbb{Z})$. Let $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$ be any lift of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ whose entries are polynomial functions of the form* (5.5)*, and let U be an upper bound on the coefficients q_k of these polynomial functions. If $t \geq 1 + (n+1)!(v+1)^{n+1}U^{n+1}$, then the classical linear program $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ is a quantization of $LP(A, b, c)$.*

Proof. Let $\boldsymbol{M} \in \mathbb{K}^{l \times l}$ be a square submatrix of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{c} & 0 \end{smallmatrix}\right)$. Since the entries of $\boldsymbol{M}$ are of the from (5.5), with coefficient $q_i \leq U$, the determinant of $\boldsymbol{M}$ is of the form

$$\det \boldsymbol{M} = \sum_{k=0}^{lv} r_k t^k$$

where the coefficients r_k are integers with an absolute value smaller that $l!(v+1)^l U^l$. Observe that $\det \boldsymbol{M}$ is a polynomial in t. So, if t is larger than the largest root of $\det \boldsymbol{M}$, then the real number $\det \boldsymbol{M}(t)$ have the same sign as the leading coefficient r_{lv}, which is the sign of the Hahn series $\det \boldsymbol{M}$. The Cauchy bound tells us that the roots of $\det \boldsymbol{M}$ belongs to a disk of radius $1 + \max_{k \in [lv-1]} |r_k|/|r_{lv}|$, see Theorem 8.1.3 and Corollary 8.1.8 in [RS02]. Since $m \geq n$, the biggest square submatrices of $\left(\begin{smallmatrix} A & b \\ c & \mathbb{0} \end{smallmatrix}\right)$ are of size $(n+1) \times (n+1)$. $\square$

Remark 5.13. The bound of Proposition 5.12 is general. For special cases, one can expect to obtain a quantization for smaller values of t.

Chapter 6

Tropical shadow-vertex rule for mean payoff games

In this chapter, we prove that the shadow-vertex pivoting rule is tropically tractable. Following the average-case analysis of Adler, Karp and Shamir in [AKS87], we obtain an algorithm that determines the feasibility of tropical polyhedra, and thus solves mean payoff games, in polynomial time on average. The complexity bound holds when the distribution of the games satisfies a *flip invariance* property. The latter requires that the distribution of the games is left invariant by every transformation consisting, for an arbitrary node of the game, in flipping the orientation of all the arcs incident to this node (see Figure 6.1). Equivalently, the probability distribution on the set of payment matrices A, B is invariant by every transformation consisting in swapping the ith row of A with the ith row of B, or the jth column of A with the jth column of B.

The content of this chapter appeared in [ABG14].

6.1 The shadow-vertex pivoting rule

Given $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{K}^n$, consider the following parametric family of linear programs for increasing values of $\boldsymbol{\lambda} \geq 0$:

$$\begin{array}{rl} \text{minimize} & (\boldsymbol{u} - \boldsymbol{\lambda}\boldsymbol{v})^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0 \end{array} \qquad \mathbf{LP}_{\boldsymbol{\lambda}}$$

The vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ are respectively called *objective* and *co-objective* vectors. For $\boldsymbol{\lambda} = 0$, the problem $(\mathbf{LP}_{\boldsymbol{\lambda}})$ seeks a minimizer of $\boldsymbol{x} \mapsto \boldsymbol{u}^\top \boldsymbol{x}$ over $\boldsymbol{\mathcal{P}} := \boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$, while for $\boldsymbol{\lambda}$ large enough, it corresponds to the maximization of $\boldsymbol{x} \mapsto \boldsymbol{v}^\top \boldsymbol{x}$.

Let us assume that $(\mathbf{LP}_{\boldsymbol{\lambda}})$ admits an optimal basic point $\boldsymbol{x}^{I^1}$ for $\boldsymbol{\lambda}^0 = 0$. Observe that $\boldsymbol{x}^{I^1}$ is also an optimal solution of $(\mathbf{LP}_{\boldsymbol{\lambda}})$ when $\boldsymbol{\lambda}$ lies in a certain closed interval $[\boldsymbol{\lambda}^0, \boldsymbol{\lambda}^1]$. For $\boldsymbol{\lambda} = \boldsymbol{\lambda}^1$, the problem $(\mathbf{LP}_{\boldsymbol{\lambda}})$ admits another optimal basic point $\boldsymbol{x}^{I^2}$ which is adjacent to $\boldsymbol{x}^{I^1}$. When $\boldsymbol{\lambda}$ is continuously increased from 0, we can construct in this way a sequence $\boldsymbol{x}^{I^1}, \dots, \boldsymbol{x}^{I^N}$ of adjacent basic points, and a subdivision $0 = \boldsymbol{\lambda}^0 \leq \boldsymbol{\lambda}^1 \leq \boldsymbol{\lambda}^1 \leq \dots \leq \boldsymbol{\lambda}^N$ of $\mathbb{K}_+$, such that each $\boldsymbol{x}^{I^k}$ is an optimal solution of $(\mathbf{LP}_{\boldsymbol{\lambda}})$ for all $\boldsymbol{\lambda} \in [\boldsymbol{\lambda}^{k-1}, \boldsymbol{\lambda}^k]$. The

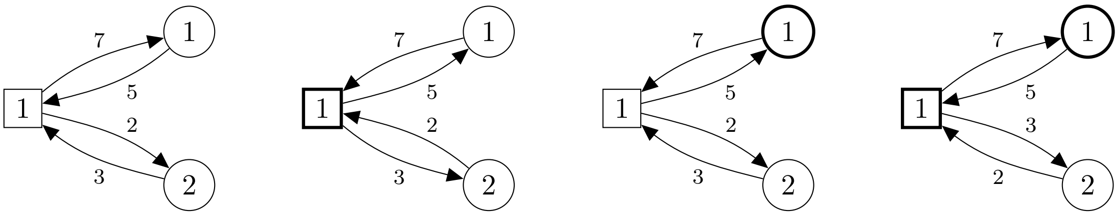

$A = (7\ 2)$ $B = (5\ 3)$ | $A = (5\ 3)$ $B = (7\ 2)$ | $A = (5\ 2)$ $B = (7\ 3)$ | $A = (7\ 3)$ $B = (5\ 2)$

Figure 6.1: A distribution of game satisfying the flip invariance property (with $m = 1$ and $n = 2$), together with the payment matrices. The four configurations are supposed to be equiprobable. The nodes on which the flip operations have been performed are depicted in bold.

last basic point $\boldsymbol{x}^{I^N}$ will be a maximizer of $\boldsymbol{x} \mapsto \boldsymbol{v}^\top \boldsymbol{x}$ over $\boldsymbol{\mathcal{P}}$, unless this problem is unbounded.

The *shadow-vertex rule* ρ is a pivoting rule that provides such a sequence. More precisely, $\rho(\boldsymbol{A}, \boldsymbol{u}, \boldsymbol{v})$ will denote the function which, given a basis I^k that is optimal for $(\mathbf{LP}_{\boldsymbol{\lambda}})$ for all $\boldsymbol{\lambda} \in [\boldsymbol{\lambda}^k, \boldsymbol{\lambda}^{k+1}]$, returns a leaving variable that leads to a basis I^{k+1} such that I^k and I^{k+1} are both optimal for $(\mathbf{LP}_{\boldsymbol{\lambda}})$ at $\boldsymbol{\lambda} = \boldsymbol{\lambda}^{k+1}$.

The shadow-vertex rule was proposed by Gass and Saaty [GS55]. Its name comes from the fact that the sequence of basic points $\boldsymbol{x}^{I^1}, \dots, \boldsymbol{x}^{I^N}$ actually corresponds to a sequence of adjacents basic points in the projection (shadow) of the polyhedron $\boldsymbol{\mathcal{P}}$ in the plane spanned by $(\boldsymbol{u}, \boldsymbol{v})$. We refer to [Bor87] for more details.

The shadow-vertex rule can also be defined algebraically. Given a basis I, we denote by $\boldsymbol{y}^I \in \mathbb{K}^I$ (resp. $\boldsymbol{z}^I \in \mathbb{K}^I$) the reduced costs for the objective vector $\boldsymbol{u}$ (resp. the co-objective vector $\boldsymbol{v}$). Recall that $\boldsymbol{y}^I$ and $\boldsymbol{z}^I$ are defined as the unique solutions $\boldsymbol{y}$ and $\boldsymbol{z}$ of the systems $\boldsymbol{A}_I^\top \boldsymbol{y} = \boldsymbol{u}$ and $\boldsymbol{A}_I^\top \boldsymbol{z} = \boldsymbol{v}$ respectively.

Proposition 6.1. *Let I be an optimal basis of $(\mathbf{LP}_{\boldsymbol{\lambda}})$ for some $\boldsymbol{\lambda} \geq 0$. At basis I, the shadow-vertex rule selects the leaving variable $i_{\mathsf{out}} \in I$ such that:*

$$\boldsymbol{y}^I_{i_{\mathsf{out}}} / \boldsymbol{z}^I_{i_{\mathsf{out}}} = \min\bigl\{ \boldsymbol{y}^I_i / \boldsymbol{z}^I_i \mid i \in I \text{ and } \boldsymbol{z}^I_i > 0 \bigr\} \ . \tag{6.1}$$

If there is no such $i_{\mathsf{out}} \in I$, then $\boldsymbol{x}^I$ maximizes $\boldsymbol{x} \mapsto \boldsymbol{v}^\top \boldsymbol{x}$ over $\boldsymbol{\mathcal{P}}$.

Proof. Observe that the reduced costs for the objective vector $\boldsymbol{u} - \boldsymbol{\lambda} \boldsymbol{v}$ are given by $\boldsymbol{y}^I - \boldsymbol{\lambda} \boldsymbol{z}^I$. Consequently, $\boldsymbol{x}^I$ is an optimal solution of $(\mathbf{LP}_{\boldsymbol{\lambda}})$ for all $\boldsymbol{\lambda} \geq 0$ such that $\boldsymbol{y}^I - \boldsymbol{\lambda} \boldsymbol{z}^I \geq 0$. In particular, this holds for $\boldsymbol{\lambda}^* = \boldsymbol{y}^I_{i_{\mathsf{out}}} / \boldsymbol{z}^I_{i_{\mathsf{out}}}$, where i_{out} is defined in (6.1). Moreover, the reduced cost $\boldsymbol{y}^I_{i_{\mathsf{out}}} - \boldsymbol{\lambda}^* \boldsymbol{z}^I_{i_{\mathsf{out}}}$ equals zero. Hence, any point on the edge $\boldsymbol{\mathcal{E}}_{I \setminus \{i_{\mathsf{out}}\}}$ is an optimal solution of $(\mathbf{LP}_{\boldsymbol{\lambda}})$ for $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$. $\square$

The shadow-vertex rule as a semi-algebraic rule

We claim that the shadow-vertex rule is a semi-algebraic pivoting rule which is tropically tractable. More precisely, we shall see that the leaving variable returned by $\rho(\boldsymbol{A}, \boldsymbol{u}, \boldsymbol{v})(I)$ only depends on the current basis I and on the signs of finitely many minors of the matrix $(\boldsymbol{A}^\top\ \boldsymbol{u}\ \boldsymbol{v})$. The key point is to show that two ratios $\boldsymbol{y}^I_i/\boldsymbol{z}^I_i$ and $\boldsymbol{y}^I_k/\boldsymbol{z}^I_k$ can be compared using the signs of the minors of $(\boldsymbol{A}^\top\ \boldsymbol{u}\ \boldsymbol{v})$.

Lemma 6.2. *Let I be a basis and $i, k \in I$ with $i > k$. Then, we have:*

$$\boldsymbol{y}^I_i/\boldsymbol{z}^I_i - \boldsymbol{y}^I_k/\boldsymbol{z}^I_k = \frac{\det \begin{pmatrix} \boldsymbol{A}_{I\setminus\{i,k\}} \\ \boldsymbol{u}^\top \\ \boldsymbol{v}^\top \end{pmatrix} \det \boldsymbol{A}_I}{\det \begin{pmatrix} \boldsymbol{A}_{I\setminus\{i\}} \\ \boldsymbol{v}^\top \end{pmatrix} \det \begin{pmatrix} \boldsymbol{A}_{I\setminus\{k\}} \\ \boldsymbol{v}^\top \end{pmatrix}} \tag{6.2}$$

Proof. By the Cramer's formulæ, for any $i \in I$ we have:

$$\begin{aligned} \boldsymbol{y}^I_i &= (-1)^{n+\mathsf{idx}(i,I)} \det \begin{pmatrix} \boldsymbol{A}_{I\setminus\{i\}} \\ \boldsymbol{u}^\top \end{pmatrix} / \det(\boldsymbol{A}_I)\ , \\ \boldsymbol{z}^I_i &= (-1)^{n+\mathsf{idx}(i,I)} \det \begin{pmatrix} \boldsymbol{A}_{I\setminus\{i\}} \\ \boldsymbol{v}^\top \end{pmatrix} / \det(\boldsymbol{A}_I)\ , \end{aligned} \tag{6.3}$$

where $\mathsf{idx}(i, I)$ represents the index of i in the ordered set I.

Given $K \subseteq [m+2]$ a subset of cardinality n, let us denote by P_K the polynomial providing the $K \times [n]$ minor of the matrix $X = (X_{ij})$ of $(m+2) \times n$ formal variables, *i.e.*, we have $P_K(\boldsymbol{M}) = \det(\boldsymbol{M}_K)$ for any $\boldsymbol{M} \in \mathbb{K}^{(m+2)\times n}$.

Given a basis $I \subseteq [m]$, let us further define, for $i \in I$, the polynomials Q_i and R_i by:

$$Q_i := P_{I\setminus\{i\}\cup\{m+1\}} \ \text{ and } \ R_i := P_{I\setminus\{i\}\cup\{m+2\}}. \tag{6.4}$$

Then, for any $i \in I$, the reduced costs $\boldsymbol{y}^I_i$ and $\boldsymbol{z}^I_i$ are respectively given by:

$$\begin{aligned} \boldsymbol{y}^I_i &= (-1)^{n+\mathsf{idx}(i,I)} Q_i(\boldsymbol{M})/P_I(\boldsymbol{M})\ , \\ \boldsymbol{z}^I_i &= (-1)^{n+\mathsf{idx}(i,I)} R_i(\boldsymbol{M})/P_I(\boldsymbol{M})\ , \end{aligned} \tag{6.5}$$

where $\boldsymbol{M} = \begin{pmatrix} \boldsymbol{A} \\ \boldsymbol{u}^\top \\ \boldsymbol{v}^\top \end{pmatrix}$. Consequently, the ratio $\boldsymbol{y}^I_i/\boldsymbol{z}^I_i$ is equal to $Q_i(\boldsymbol{M})/R_i(\boldsymbol{M})$. For any two distincts indices $i, k \in I$, we obtain:

$$\boldsymbol{y}^I_i/\boldsymbol{z}^I_i - \boldsymbol{y}^I_k/\boldsymbol{z}^I_k = \frac{(Q_i R_k - Q_k R_i)(\boldsymbol{M})}{(R_i R_k)(\boldsymbol{M})} \tag{6.6}$$

It remains to prove that the polynomial $Q_i R_k - Q_k R_i$ is equal to $P_{I\setminus\{i,k\}\cup\{m+1,m+2\}} P_I$. By Plücker relations (see for instance [GKZ94, Chapter 3, Theorem 1.3]), we know that for any two sequences $1 \leq j_1 < \cdots < j_{n-1} \leq m+2$ and $1 \leq l_1 < \cdots < l_{n+1} \leq m+2$, we have:

$$\sum_{a=1}^{n+1} (-1)^a P_{\{j_1,,\dots,j_{n-1},l_a\}} P_{\{l_1,\dots,\widehat{l_a},\dots,l_{n+1}\}} = 0 \ , \tag{6.7}$$

where $\widehat{l_a}$ means that the index l_a is omitted. Let us apply these relations with

$$\{j_1, \dots, j_{n-1}\} = I \setminus \{i\} \ \text{ and } \ \{l_1, \dots, l_{n+1}\} = I \setminus \{k\} \cup \{m+1, m+2\} \ .$$

If $l_a \in I \setminus \{k\}$, then $P_{\{j_1,,\dots,j_{n-1},l_a\}} = 0$. Hence, the only terms that are non null in (6.7) are obtained for $l_a \in \{i, m+1, m+2\}$. For $l_a = i$, the term reads

$$(-1)^{\mathsf{idx}(i, I\setminus\{k\})} P_{j_1,\dots,j_{n-1},i} P_{I\setminus\{i,k\}\cup\{m+1,m+2\}} \ . \tag{6.8}$$

By exchange of rows on the determinant $P_{j_1,\dots,j_{n-1},i}$, we have:

$$P_{j_1,\dots,j_{n-1},i} = (-1)^{n+\mathsf{idx}(i,I)} P_I. \tag{6.9}$$

Furthermore, we assumed that $i > k$, thus $\mathsf{idx}(i, I) = \mathsf{idx}(i, I \setminus \{k\}) + 1$. It follows that (6.8) is the polynomial $(-1)^{n+1} P_I P_{I\setminus\{i,k\}\cup\{m+1,m+2\}}$.

The indices of $m+1$ and $m+2$ in the ordered set $\{l_1, \dots, l_{n+1}\}$ are respectively n and $n+1$. Thus the terms of (6.7) for $l_a = m+1$ and $l_a = m+2$ are respectively:

$$(-1)^n P_{I\setminus\{i\}\cup\{m+1\}} P_{I\setminus\{k\}\cup\{m+2\}} \quad \text{and} \quad (-1)^{n+1} P_{I\setminus\{i\}\cup\{m+2\}} P_{I\setminus\{k\}\cup\{m+1\}} \ .$$

Finally, we obtain the equality:

$$(-1)^{n+1} P_I P_{I\setminus\{i,k\}\cup\{m+1,m+2\}} + (-1)^n Q_i R_k - (-1)^{n+1} Q_k R_i = 0 \ .$$

This concludes the proof. □

The main result of this section is the following:

Theorem 6.3. *The shadow vertex rule $\rho(\boldsymbol{A}, \boldsymbol{u}, \boldsymbol{v})$ is a semi-algebraic pivoting rule that uses only the signs of the maximal minors of $(\boldsymbol{A}^\top \ \boldsymbol{u} \ \boldsymbol{v})$. The tropical shadow vertex rule $\rho^{\mathbb{T}}(A, u, v)$ returns the leaving variable in $O(n^4)$ operations and in space polynomial in the input size of A, u, v.*

Proof. By Proposition 6.1, the shadow-vertex rule can be implemented using only the signs of $\boldsymbol{z}^I$, and the signs of $\boldsymbol{y}_i^I/\boldsymbol{z}_i^I - \boldsymbol{y}_k^I/\boldsymbol{z}_k^I$ for all $i, k \in I$.

The reduced costs vector $\boldsymbol{z}^I$ is given by the Cramer's formulæ for the system $\boldsymbol{A}_I^\top \boldsymbol{y} = \boldsymbol{v}$ Hence, the signs of $\boldsymbol{z}^I$ can be determined by computing the determinant of $\boldsymbol{A}_I$ and of $\begin{pmatrix} \boldsymbol{A}_{I\setminus\{i\}} \\ \boldsymbol{v}^\top \end{pmatrix}$ for $i \in I$. Hence, $2n+1$ determinants of size $n \times n$.

This allows to determine the set $\Lambda = \{i \in I \mid \boldsymbol{z}_i^I > 0\}$. Then, the leaving variable i_{out} that minimizes the ratio $\boldsymbol{y}_i^I/\boldsymbol{z}_i^I$ for $i \in \Lambda$ can be found by performing $O(n)$ comparisons $\boldsymbol{y}_i^I/\boldsymbol{z}_i^I < \boldsymbol{y}_k^I/\boldsymbol{z}_k^I$. By Lemma 6.2, each of these comparisons can be done by computing four $n \times n$ determinants.

To summarize, we need to compute $O(n)$ determinants of size $n \times n$. By Lemma 3.11, a determinant is tropically tractable. Moreover, a $n \times n$ tropical determinant can be computed tropically in $O(n^3)$ operations. This concludes the proof. □

In the next section, we will use the shadow-vertex rule with the objective vector $\boldsymbol{u} = (\boldsymbol{\epsilon}, \boldsymbol{\epsilon}^2, \dots, \boldsymbol{\epsilon}^n)$ for $\boldsymbol{\epsilon} > 0$ small enough. In that case, there is no need to choose or manipulate $\boldsymbol{\epsilon}$ explicitely. We present a proof which is comparable to the "lexicographic" treatment described in [AKS87, Section 6.1]. When such an objective vector $\boldsymbol{u}$ is used, we will denote the pivoting rule $\rho(\boldsymbol{A}, \boldsymbol{u}, \boldsymbol{v})$ by $\rho_\epsilon(\boldsymbol{A}, \boldsymbol{v})$.

Corollary 6.4. *The pivoting rule $\rho_\epsilon(\boldsymbol{A}, \boldsymbol{v})$ is a semi-algebraic pivoting rule that uses only the signs of the minors of $(\boldsymbol{A}^\top \ \boldsymbol{v})$. The tropical pivoting rule $\rho_\epsilon^{\mathbb{T}}(A, v)$ returns the leaving variable in $O(n^5)$ operations and in space polynomial in the input size of A, v.*

Proof. Let $\boldsymbol{M}$ be a $n \times n$ submatrix of $(\boldsymbol{A}^\top \ \boldsymbol{u} \ \boldsymbol{v}) \in \mathbb{K}^{n \times (m+2)}$ that contains the column $\boldsymbol{u}$. By Theorem 6.3, we only need to show that the sign of $\det(\boldsymbol{M})$ can be computed from the signs of the minors of $(\boldsymbol{A}^\top \ \boldsymbol{v})$. Up to exchange of columns, we can assume that $\boldsymbol{M} = (\boldsymbol{u} \ \boldsymbol{M}')$, where $\boldsymbol{M}'$ is a submatrix of $(\boldsymbol{A}^\top \ \boldsymbol{v})$. Expanding the determinant of $\boldsymbol{M}$ along its first column, we obtain:

$$\det(\boldsymbol{M}) = \sum_{i=1}^{n} \boldsymbol{\epsilon}^i (-1)^{1+i} \det(\boldsymbol{M}'_{[n] \setminus \{i\}}) \ .$$

If $\det(\boldsymbol{M}'_{[n] \setminus \{i\}}) = 0$ for all $i \in [n]$, then clearly the determinant of $\boldsymbol{M}$ vanishes. Otherwise, let i^* be the smallest $i \in [n]$ such that $\det(\boldsymbol{M}'_{[n] \setminus \{i\}}) \neq 0$. Then, if we choose $\boldsymbol{\epsilon} > 0$ small enough, the sign of $\det(\boldsymbol{M})$ will be given by the sign of $(-1)^{1+i^*} \det(\boldsymbol{M}'_{[n] \setminus \{i^*\}})$. Consequently, the sign of $\det(\boldsymbol{M})$ can be obtained by computing the signs of the n determinants $\det(\boldsymbol{M}'_{[n] \setminus \{i\}})$, that are all of size $(n-1) \times (n-1)$.

We have seen in the proof of Theorem 6.3 that we only need to compute the signs of $O(n)$ minors of $(\boldsymbol{A}^\top \ \boldsymbol{u} \ \boldsymbol{v})$, including n minors that involves the column $\boldsymbol{u}$. By the discussion above, the sign of each minor involving $\boldsymbol{u}$ can be computed from n minors of $(\boldsymbol{A}^\top \ \ \boldsymbol{v})$ of size $(n-1) \times (n-1)$. Since a tropical determinant can be computed in $O(n^3)$ operations, we obtain $O(n^5)$ operations for the tropicalization of ρ_ϵ. ☐

6.2 The Parametric Constraint-by-Constraint algorithm

The average-case analysis of [AKS87] applies to the *Parametric Constraint-by-Constraint* algorithm (denoted *PCBC*). We restrict the presentation to polyhedral feasibility problems, following our motivation to their tropical counterparts and mean payoff games. This algorithm applies to polyhedra $\boldsymbol{\mathcal{P}}(\boldsymbol{A}, \boldsymbol{b})$ that satisfy the following assumption.

Assumption C. The matrix $(\boldsymbol{A} \ \boldsymbol{b}) \in \mathbb{K}^{(m+n) \times (n+1)}$ is of the form $\left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{A}' & \boldsymbol{b}' \end{smallmatrix}\right)$, where $\mathbf{Id}$ is the $n \times n$ identity matrix.

Equivalently, we consider polyhedra of the form:

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

We denote by $\boldsymbol{\mathcal{P}}^{(k)} := \boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k]}, \boldsymbol{b}_{[k]})$ the polyhedron defined by the first k inequalities of the system $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0$. Under Assumption C, the polyhedron $\boldsymbol{\mathcal{P}}^{(n)}$ is the positive orthant

Algorithm 3: The parametric constraint by constraint algorithm $PCBC(\boldsymbol{A}, \boldsymbol{b})$

Data: $\boldsymbol{A} \in \mathbb{K}^{(m+n)\times n}$ and $\boldsymbol{b} \in \mathbb{K}^m$ satisfying Assumption C.
Input: None
Output: Either Feasible or Infeasible

$k \leftarrow n$
$I \leftarrow [n]$
while $k < m$ **do**
  if I is a feasible basis for $\boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k+1]}, \boldsymbol{b}_{[k+1]})$ **then**
    $k \leftarrow k+1$
  else
    if $\mathsf{SignRedCosts}(\boldsymbol{A}_{[k]}, \boldsymbol{A}_{k+1})(I)$ are all non-negative **then**
      return Infeasible
    $i_{\mathsf{out}} \leftarrow \rho_\epsilon(\boldsymbol{A}_{[k]}, \boldsymbol{u}, \boldsymbol{A}_{k+1})(I)$
    if $I \setminus \{i_{\mathsf{out}}\} \cup \{k+1\}$ is a feasible basis for $\boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k+1]}, \boldsymbol{b}_{[k+1]})$ **then**
      $I \leftarrow I \setminus \{i_{\mathsf{out}}\} \cup \{k+1\}$
      $k \leftarrow k+1$
    else
      $i_{\mathsf{ent}} \leftarrow \mathsf{Pivot}(\boldsymbol{A}_{[k]}, \boldsymbol{b}_{[k]})(I, i_{\mathsf{out}})$
      $I \leftarrow I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$
return Feasible

$\mathbb{K}^n_+$. The $PCBC$ algorithm consists of m stages, that we index by $k \in \{n, \dots, m+n-1\}$. At stage k, the algorithm determines whether the polyhedron $\boldsymbol{\mathcal{P}}^{(k+1)}$ is empty. At each stage, the simplex algorithm equipped with the pivoting rule ρ_ϵ is used, *i.e.*, throughout the whole execution of $PCBC$, the shadow-vertex rule is used with the objective vector $\boldsymbol{u} = (\boldsymbol{\epsilon}, \dots, \boldsymbol{\epsilon}^n)$, for $\boldsymbol{\epsilon} > 0$ small enough. On the other hand, the co-objective vector $\boldsymbol{v}$ will change at each stage. The vector $(0, \dots, 0)^\top$ is a basic point of $\boldsymbol{\mathcal{P}}^{(n)} = \mathbb{K}^n_+$ minimizing $\boldsymbol{x} \mapsto \boldsymbol{u}^\top \boldsymbol{x}$. This provides an initial basis which is compatible with the shadow-vertex rule at the first stage $k = n$.

At stage k, the co-objective vector is set to $\boldsymbol{A}^\top_{k+1}$. The simplex algorithm thus follows a path in $\boldsymbol{\mathcal{P}}^{(k)}$ consisting of basic points and the edges between them. We stop it as soon as it discovers a point $\boldsymbol{x}' \in \boldsymbol{\mathcal{P}}^{(k)}$ such that $\boldsymbol{A}_{k+1}\boldsymbol{x}' + \boldsymbol{b}_{k+1} \geq 0$ on the path. This point is obviously a basic point of $\boldsymbol{\mathcal{P}}^{(k+1)}$. It follows from the definition of the shadow-vertex rule that $\boldsymbol{x}'$ minimizes the objective function $\boldsymbol{x} \mapsto \boldsymbol{u}^\top \boldsymbol{x}$ over $\boldsymbol{\mathcal{P}}^{(k+1)}$ (see [AKS86, Section 4]). Then, $\boldsymbol{x}'$ can be used as a starting point for the execution of the simplex algorithm during the $(k+1)$-th iteration. If no such point $\boldsymbol{x}'$ is discovered, then the maximum of $\boldsymbol{x} \mapsto \boldsymbol{A}_k \boldsymbol{x} + \boldsymbol{b}_k$ over $\boldsymbol{\mathcal{P}}^{(k)}$ is negative, which shows that the system $\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0$ is infeasible.

We now explain how to tropicalize the $PCBC$ algorithm. As for the simplex method,

Algorithm 4: The tropical parametric constraint by constraint algorithm *TropPCBC*(A, b)

Data: $A \in \mathbb{T}_{\pm}^{m \times n}$ and $b \in \mathbb{T}_{\pm}^{m}$ such that $(A\ b) = \begin{pmatrix} \mathrm{Id} & \mathbb{0} \\ A' & b' \end{pmatrix}$ where Id is the $n \times n$ identity matrix.

Output: Either Feasible or Infeasible.

1 $k \leftarrow n$
2 $I \leftarrow [n]$
3 **while** $k < m$ **do**
4   **if** I is a feasible basis for $\mathcal{P}(A_{[k+1]}, B_{[k+1]})$ **then**
5         $k \leftarrow k+1$
6   **else**
7   **if** $\mathsf{SignRedCosts}^{\mathbb{T}}(A_{[k]}, A_{k+1})(I)$ are all non-negative **then**
8   **return** Infeasible
9         $i_{\mathsf{out}} \leftarrow \rho_{\epsilon}^{\mathbb{T}}(A_{[k]}, u, A_{k+1})(I)$
10   **if** $I \setminus \{i_{\mathsf{out}}\} \cup \{k+1\}$ is a feasible basis for $\mathcal{P}(A_{[k+1]}, b_{[k+1]})$ **then**
11             $I \leftarrow I \setminus \{i_{\mathsf{out}}\} \cup \{k+1\}$
12             $k \leftarrow k+1$
13   **else**
14             $i_{\mathsf{ent}} \leftarrow \mathsf{Pivot}^{\mathbb{T}}(A_{[k]}, b_{[k]})(I, i_{\mathsf{out}})$
15             $I \leftarrow I \setminus \{i_{\mathsf{out}}\} \cup \{i_{\mathsf{ent}}\}$
16 **return** Feasible

it is sufficient to show that $PCBC(\boldsymbol{A}, \boldsymbol{b})$ can be implemented using only the signs of polynomials evaluated on $(\boldsymbol{A}\ \boldsymbol{b})$. Such an implementation is presented in Algorithm 3.

Proposition 6.5. *For any* $\boldsymbol{A} \in \mathbb{K}^{(m+n)\times n}$, $\boldsymbol{b} \in \mathbb{K}^{m}$, *satisyfing Assumption C, Algorithm 3 is an implementation of the Parametric Constraint by Constraint algorithm in the arithmetic model of computation with an oracle that returns the signs of the minors of* $(\boldsymbol{A}\ \boldsymbol{b})$.

Proof. Obviously, $I = [n]$ is a basis of the initial basic point $\boldsymbol{x} = (0, \dots, 0)$ at the first stage $k = n$.

Now suppose that we are the beginning of stage k of the algorithm, with basis I. If $\boldsymbol{A}_{k+1}\boldsymbol{x}^I + \boldsymbol{b}_{k+1} \geq 0$, then the algorithm should go to stage $k+1$. Clearly, this happens if and only if I is a feasible basis for $\boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k+1]}, \boldsymbol{b}_{[k+1]})$ and this is detected at Line 4. Consequently, we can assume that at Line 7, we have $\boldsymbol{A}_{k+1}\boldsymbol{x}^I + \boldsymbol{b}_{k+1} < 0$. If the sign of the reduced costs are non-negative, then $\boldsymbol{x}^I$ maximizes $\boldsymbol{x} \mapsto \boldsymbol{A}_{k+1}\boldsymbol{x} + \boldsymbol{b}_{k+1}$ over $\boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k]}, \boldsymbol{b}_{[k]})$ and thus the linear program is infeasible. Otherwise, the shadow-vertex pivoting rule returns a leaving variable i_{out}. The edge defined by (I, i_{out}) may contain a point such that $\boldsymbol{A}_{k+1}\boldsymbol{x}^I + \boldsymbol{b}_{k+1} = 0$, in that case, the algorithm go to stage $k+1$. This happens if and only if $I \setminus \{i_{\mathsf{out}}\} \cup \{k+1\}$ is a feasible basis for $\boldsymbol{\mathcal{P}}(\boldsymbol{A}_{[k+1]}, \boldsymbol{b}_{[k+1]})$.

If no such point is encountered, the algorithm pivots along the edge defined by (I, i_{out}) and reaches a new basic point. Clearly, this point must satisfies $\boldsymbol{A}_{k+1}\boldsymbol{x} + \boldsymbol{b}_{k+1} < 0$. Consequently, when the body of loop at Line 3 is executed again, the test at Line 4 fails and the algorithm goes directly to Line 7. This shows that Algorithm 4 does implement the Parametric Constraint by Constraint algorithm.

By Lemma 3.16, the feasibility of a basis can be tested at Line 10 using the signs of the minors of $(\boldsymbol{A}\ \boldsymbol{b})$. By Proposition 3.22 and Corollary 6.4, the other operations in Algorithm 4 can also be implemented with the signs of the minors of $(\boldsymbol{A}\ \boldsymbol{b})$. □

As an immediate consequence of Proposition 6.5, the *PCBC* algorithm has a tropical counterpart, *TropPCBC*, which is described in Algorithm 4.

Theorem 6.6. *Let $A \in \mathbb{T}_{\pm}^{(m+n)\times n}$ and $b \in \mathbb{T}_{\pm}^{m}$ be such that $(A\ b)$ is sign-generic for the minor polynomials and $(A\ b) = \left(\begin{smallmatrix} Id & 0 \\ A' & b' \end{smallmatrix}\right)$, where Id is the $n \times n$ identity matrix. Then, the algorithm TropPCBC correctly determines whether $\mathcal{P}(A, b)$ is feasible.*

For all $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$, the total number of bases visited by TropPCBC(A, b) and by PCBC$(\boldsymbol{A}, \boldsymbol{b})$ are equal.

Between two bases, TropPCBC performs $O(n^5 + m^2n^3)$ operations and uses a space bounded by a polynomial in the input size of A, b.

Proof. Observe that Algorithm 4 is exactly Algorithm 3 where we have replaced the oracle giving the signs of the minors by its tropical counterpart. By Lemma 3.8 and Proposition 6.5, it follows that *TropPCBC*(A, b) and *PCBC*$(\boldsymbol{A}, \boldsymbol{b})$ produce the same sequence of bases for any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$. The correctness of *TropPCBC* then follows from the correctness of *PCBC* and Proposition 4.7.

Pivoting from one basis to the next consists of performing once the operations in the loop between Lines 3 and 15. Calling $\mathsf{SignRedCosts}^{\mathbb{T}}$ and $\mathsf{Pivot}^{\mathbb{T}}$ requires $O(n^4)$ and $O(m^2n^3)$ operations respectively by Proposition 3.22. The pivoting rule $\rho_{\epsilon}^{\mathbb{T}}$ returns after $O(n^5)$ operations by Corollary 6.4. Checking the feasibility of a basis requires the computation of $O(m)$ determinants of size $n \times n$ (see Lemma 3.16), and each of these determinants can be computed tropically in $O(n^3)$ operations by Lemma 3.11. Hence, we need $O(mn^3)$ operations to test the feasibility of a basis. In total, we use $O(n^5 + n^4 + m^2n^3 + mn^3) = O(n^5 + m^2n^3)$ operations. Moreover, these operations use a polynomial space. □

6.2.1 Average-case analysis

Given $(\boldsymbol{A}\ \boldsymbol{b}) = \left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{A}' & \boldsymbol{b}' \end{smallmatrix}\right) \in \mathbb{K}^{(m+n)\times(n+1)}$ such that no minor of the matrix $(\boldsymbol{A}'\ \boldsymbol{b}')$ is null, the probabilistic analysis of [AKS87] applies to polyhedra of the form

$$\boldsymbol{\mathcal{P}}_{\boldsymbol{S},\boldsymbol{S}'}(\boldsymbol{A}, \boldsymbol{b}) = \{\boldsymbol{x} \in \mathbb{K}^n \mid \boldsymbol{x} \geq 0,\ (\boldsymbol{S}\boldsymbol{A}'\boldsymbol{S}')\boldsymbol{x} + \boldsymbol{S}\boldsymbol{b}' \geq 0\}\ ,$$

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

Theorem 6.7 ([AKS87]). *For any fixed choice of* $(\boldsymbol{A}\ \boldsymbol{b}) = \left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{A}' & \boldsymbol{b}' \end{smallmatrix}\right) \in \mathbb{K}^{m\times(n+1)}$ *such that no minor of the matrix* $(\boldsymbol{A}'\ \boldsymbol{b}')$ *is null, the total number of basic points visited by the PCBC algorithm on* $\boldsymbol{\mathcal{P}}_{\boldsymbol{S},\boldsymbol{S}'}(\boldsymbol{A},\boldsymbol{b})$ *is bounded by* $O(\min(m^2,n^2))$ *on average.*

Proof. This result is proved in [AKS87] for matrices $(\boldsymbol{A}\ \boldsymbol{b})$ with entries in $\mathbb{R}$. We now show that it holds for matrices $(\boldsymbol{A}\ \boldsymbol{b})$ with entries in an arbitrary real closed field K.

Let $(\boldsymbol{A}\ \boldsymbol{b})$ be a matrix with entries in K that satisfies the conditions of the theorem. By Proposition 6.5, the number of basic points visited by *PCBC* on the polyhedron $\boldsymbol{\mathcal{P}}(\boldsymbol{A},\boldsymbol{b})$ depends only on the sign pattern $\mathsf{signMinors}(\boldsymbol{A},\boldsymbol{b})$ of the minors of $(\boldsymbol{A}\ \boldsymbol{b})$. By completeness of the theory of real closed field (Theorem 2.2, there exists a matrix $(\mathsf{A}\ \mathsf{b})$ with entries in $\mathbb{R}$ that realizes the sign pattern $\mathsf{signMinors}(\boldsymbol{A},\boldsymbol{b})$ (see the proof of Proposition 5.2 for details). Clearly, $(\mathsf{A}\ \mathsf{b})$ satisfies the conditions of the theorem.

Observe that the signs of the minors of $\left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{SAS}' & \boldsymbol{Sb} \end{smallmatrix}\right)$ are entirely determined by $\boldsymbol{S},\boldsymbol{S}'$ and the signs of the minors of $(\boldsymbol{A}\ \boldsymbol{b})$. Consequently, the signs of the minors of $\left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{SAS}' & \boldsymbol{Sb} \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} \mathrm{Id} & 0 \\ \boldsymbol{S}\mathsf{A}\boldsymbol{S}' & \boldsymbol{S}\mathsf{b} \end{smallmatrix}\right)$ coincides. It follows that the *PCBC* algorithm visits the same the number of basic points on $\boldsymbol{\mathcal{P}}_{\boldsymbol{S},\boldsymbol{S}'}(\boldsymbol{A},\boldsymbol{b})$ and $\boldsymbol{\mathcal{P}}_{\boldsymbol{S},\boldsymbol{S}'}(\mathsf{A},\mathsf{b})$. Since the theorem holds on $\mathbb{R}$, it also holds on K. ☐

As a consequence of Theorems 6.6 and 6.7, the algorithm *TropPCBC* also visits a quadratic number of tropical basic points on average. The tropical counterpart of the probabilistic model of [AKS87] can be described as follows. Given $(A\ b) = \left(\begin{smallmatrix} \mathrm{Id} & 0 \\ A' & b' \end{smallmatrix}\right) \in \mathbb{T}_{\pm}^{(m+n)\times(n+1)}$, and $s \in \{\mathbb{1},\ominus\mathbb{1}\}^m$, $s' \in \{\mathbb{1},\ominus\mathbb{1}\}^n$, we define

$$\mathcal{P}_{S,S'}(A,b) = \{x \in \mathbb{T}^n \mid x \geq \mathbb{0},\ (S\odot A'\odot S')^+\odot x\oplus(S\odot b')^+ \geq (S\odot A'\odot S')^-\odot x\oplus(S\odot b')^-\},$$

where $S = \operatorname{diag}(s_1,\dots,s_m)$, $S' = \operatorname{diag}(s'_1,\dots,s'_n)$. As above, we assume that the s_i, s'_j are i.i.d random variables with value equal to $\mathbb{1}$ (resp. $\ominus\mathbb{1}$) with probability 1/2.

Corollary 6.8. *Suppose that* $(A\ b) = \left(\begin{smallmatrix} Id & 0 \\ A' & b' \end{smallmatrix}\right) \in \mathbb{T}_{\pm}^{(m+n)\times(n+1)}$ *is generic for the minor polynomials and that every square submatrix of* $(A'\ b')$ *has a non* $\mathbb{0}$ *tropical determinant.*

The total number of basic points visited by the TropPCBC algorithm on $\mathcal{P}_{S,S'}(A,b)$ *is bounded by* $O(\min(m^2,n^2))$ *on average.*

Proof. Let us pick any $(\boldsymbol{A}\ \boldsymbol{b}) \in \operatorname{sval}^{-1}(A\ b)$. Since $(A\ b)$ is generic for the minor polynomials, it is also sign-generic, and thus *TropPCBC*(A,b) and *PCBC*$(\boldsymbol{A},\boldsymbol{b})$ visits the same number of basic points by Theorem 6.6. It also follows from the genericity of $\left(\begin{smallmatrix} \mathrm{Id} & 0 \\ A' & b' \end{smallmatrix}\right)$ that, for any S, S', the matrix

$$\begin{pmatrix} \mathrm{Id} & \mathbb{0} \\ S\odot A'\odot S' & S\odot b' \end{pmatrix} \tag{6.10}$$

is also generic for the minor polynomials. Let $\boldsymbol{S} \in \operatorname{sval}^{-1}(S)$ and $\boldsymbol{S}' \in \operatorname{sval}^{-1}(S')$. Clearly, $\left(\begin{smallmatrix} \mathbf{Id} & 0 \\ \boldsymbol{SA}'\boldsymbol{S}' & \boldsymbol{Sb}' \end{smallmatrix}\right)$ is a lift of (6.10). Consequently, *TropPCBC* applied to $\mathcal{P}_{S,S'}(A,b)$ visits as many basic points as *PCBC* applied to $\boldsymbol{\mathcal{P}}_{\boldsymbol{S},\boldsymbol{S}'}(\boldsymbol{A},\boldsymbol{b})$ by Theorem 6.6. We conclude with Theorem 6.7. ☐

6.3 Application to mean payoff games

Via the tropical parametric constraint by constraint algorithm, we translate the result of Adler *et al.* to mean payoff games. The probability distribution of games is expressed over their payments matrices A, B, and must satisfy the following requirements:

Assumption D. (i) for all $i \in [m]$ (resp. $j \in [n]$), the distribution of the matrices A, B is invariant by the exchange of the i-th row (resp. j-th column) of A and B.

(ii) almost surely, A_{ij} and B_{ij} are distinct and not equal to $\mathbb{0}$ for all $i \in [m]$, $j \in [n]$. In this case, we introduce the signed matrix $W = (W_{ij}) \in \mathbb{T}_{\pm}^{m \times n}$, defined by $W_{ij} := A_{ij}$ if $A_{ij} > B_{ij}$, and $\ominus B_{ij}$ if $A_{ij} < B_{ij}$.

(iii) almost surely, the matrix W is generic for all minor polynomials.

Let us briefly discuss the requirements of Assumption D. Condition (i) corresponds to the flip invariance property. It handles discrete distributions (see Figure 6.1) as well as continuous ones. In particular, if the distribution of the payment matrices admits a density function f, Condition (i) can be expressed as the invariance of f by exchange operations on its arguments. For instance, if $m = 1$ and $n = 2$, the flip invariance holds if, and only if, for almost all a_{ij}, b_{ij}, $f(a_{1,1}, a_{1,2}, b_{1,1}, b_{1,2}) = f(b_{1,1}, b_{1,2}, a_{1,1}, a_{1,2}) = f(b_{1,1}, a_{1,2}, a_{1,1}, b_{1,2}) = f(a_{1,1}, b_{1,2}, b_{1,1}, a_{1,2})$.

The requirements $A_{ij}, B_{ij} \neq \mathbb{0}$ for all i, j in Condition (ii) ensure that the flip operations always provide games in which the two players have at least one action to play from every position. The matrix W can be thought of as a tropical subtraction "$A \ominus B$", and the conditions $A_{ij} \neq B_{ij}$ ensure that W is well defined. Then, the following result holds:

Lemma 6.9. *If $A_{ij} \neq B_{ij}$ for all i, j, and W is defined as in Condition* (ii) *of Assumption D, then the initial state $j \in [n]$ is winning in the game with matrices A, B if, and only if the tropical polyhedron $\mathcal{P}(W_{\widehat{j}}, W_{[m]\times j})$ is not empty, where $W_{\widehat{j}}$ is the matrix obtained from W by removing the column j, and $W_{[m]\times j}$ is the jth column of W.*

Proof. By Theorem 1.3, the initial state j is winning if, and only if, the system

$$x_j = 0, \ A \odot x \geq B \odot x \ , \tag{6.11}$$

admits a solution. Given $a, b, c, d \in \mathbb{T}$ such that $a \neq c$, it can be easily proved that the inequality $\max(a + x_1, b) \geq \max(c + x_1, d)$ over x_1 is equivalent to $b \geq \max(c + x_1, d)$ if $a < c$, and $\max(a + x_1, b) \geq d$ if $a > c$. Using this principle, we deduce that the system (6.11) is equivalent to

$$x_j = 0, \ W^+ \odot x \geq W^- \odot x \ .$$

Clearly, the latter system admits a solution if and only if the tropical polyhedron $\mathcal{P}(W_{\widehat{j}}, W_{[m]\times\{j\}})$ is not empty. ∎

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

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

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

Theorem 6.10. *Under a distribution satisfying Assumption D, the algorithm TropPCBC determines in polynomial time on average whether an initial state is winning for Player Max in the mean payoff game with payment matrices A, B.*

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

Let us fix two payment matrices A, B satisfying Conditions (ii) and (iii) of Assumption D, and let W be defined as in Condition (ii). Starting from the pair (A, B) of matrices, the successive applications of row/column exchange operations precisely yield 2^{m+n-1} different pairs of matrices. In particular, without loss of generality, we can assume that the n-th columns of A and B have not been switched. Then, the pair of matrices that we obtained are of the form $(A^{s,s'}, B^{s,s'})$, where $s \in \{\mathbb{1}, \ominus\mathbb{1}\}^m$, $s' \in \{\mathbb{1}, \ominus\mathbb{1}\}^{n-1}$, and $A^{s,s'}$ and $B^{s,s'}$ are the matrices obtained from A and B respectively, by exchanging the rows i and the columns j such that $s_i = \ominus\mathbb{1}$ and $s'_j = \ominus\mathbb{1}$. The (i, j)-entries of $A^{s,s'}$ and $B^{s,s'}$ are distinct, and so we can define a matrix $W^{s,s'}$ in the same way we have built W from A and B. Observe that $W^{s,s'}_{\widehat{n}} = S \odot W_{\widehat{n}} \odot S'$ and $W^{s,s'}_{[m]\times n} = S \odot W_{[m]\times n}$, where $S = \operatorname{tdiag}(s_1, \dots, s_m)$ and $S' = \operatorname{tdiag}(s'_1, \dots, s'_{n-1})$. Thus, by Lemma 6.9, the node n is winning in the game with payment matrices $A^{s,s'}, B^{s,s'}$ if and only if the tropical polyhedron $\mathcal{P}_{S,S'}(W_{\widehat{n}}, W_{[m]\times n})$ is not empty. By Theorem 6.6 and Corollary 6.8, the *TropPCBC* algorithm solves the 2^{m+n-1} games obtained by the successive flipping operations in $O(2^{m+n-1} \min(m^2, n^2)(n^5 + m^2 n^3))$ operations and in polynomial space.

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

$$\mathbb{E}[T] = \frac{1}{2^{m+n-1}} \mathbb{E}\Big[\sum_{s,s'} T^{s,s'}\Big] \leq \frac{1}{2^{m+n-1}} (K 2^{m+n-1} \min(m^2, n^2)(n^5 + m^2 n^3))$$

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

Chapter 7

Algorithmics of the tropical simplex method

In this chapter, we present efficient implementations of the tropical pivoting procudure, and of the tropical computation of the signs of reduced costs. We show that these two prodecures can be done using $O(n(m+n))$ tropical operations for a linear program described by m inequalities in dimension n. The algorithms presented in this chapter have been implemented in the library Simplet [Ben14].

The content of this chapter appeared in [ABGJ13b].

7.1 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 basic points of a tropical polyhedron $\mathcal{P}(A, b)$, where $A \in \mathbb{T}_{\pm}^{m\times n}$ and $b \in \mathbb{T}_{\pm}^{m}$, and $\mathbb{T} = \mathbb{T}(G)$ is an arbitrary tropical semiring. 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 $\mathcal{C} = \mathcal{P}(W, \mathbb{0})$. This cone is defined as the intersection of the half-spaces $\mathcal{H}_i^{\geq} := \{x \in \mathbb{T}^{n+1} \mid W_i^+ \odot x \geq W_i^- \odot x\}$ for $i \in [m]$. Similarly, we denote by $\mathcal{H}_i$ the s-hyperplane $\{x \in \mathbb{T}^{n+1} \mid W_i^+ \odot x = W_i^- \odot x\}$. We also let $\mathcal{C}_I := \mathcal{P}_I(W, \mathbb{0})$ for any subset $I \subseteq [m]$.

Throughout this section, we make the following assumptions.

Assumption E. The matrix W is generic for the minor polynomials.

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

Assumption E is is strictly stronger than the sign-genericity of $W = (A\ b)$ for the minor polynomials, and hence, in particular, we can make use of Theorem 4.22. Under Assumption F, the tropical polyhedron $\mathcal{P}(A, b)$ is a bounded subset of G^n. Indeed, as $\mathcal{C}$ is a closed set, Assumption F implies that there exists a vector $l \in G^{n+1}$ such that $x \geq l$

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 4.11. If R contains a point r, then it is easy to verify that $(r, \mathbb{0}) \in \mathcal{C}$, which contradicts Assumption F. Since every $p \in P$ belongs to $\mathcal{P}(A, b)$, the point $(p, \mathbb{1})$ belongs to $\mathcal{C}$, and thus $p_j \geq l_j$ for all $j \in [n]$. It follows that $\mathcal{P}(A, b) = \operatorname{tconv}(P)$ is a bounded subset of G^n.

In rest of this section, we identify the cones $\mathcal{C}$, $\mathcal{H}_i$ and $\mathcal{C}_I$ of $\mathbb{T}^{n+1}$ with their image in the tropical projective space $\mathbb{TP}^n$ (see Section 4.1.2). Through the bijection given in (4.12), the tropical basic point associated with a suitable subset $I \subseteq [m]$ is identified with the unique projective point $x^I \in \mathbb{TP}^n$ in the intersection $\mathcal{C}_I$. Besides, 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 $i_{\mathsf{out}} \in I$.

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 *non-basic* 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 7.1 ([DS04, Proposition 3]). *Let $\mathcal{E}_K = \operatorname{tconv}(x^I, x^{I'})$ be a tropical edge. Then there exist an integer $q \in [n]$ and $q+1$ points $\xi^1, \dots, \xi^{q+1} \in \mathcal{E}_K$ such that*

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

Every ordinary segment is of the form:

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

where the length of the segment μ_p is a positive real number, $J_p \subseteq [n+1]$, and the j-th 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:

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

The vector e^{J_p} is called the *direction* of the segment $[\xi^p, \xi^{p+1}]$. The intermediate points $\xi^2, \dots, \xi^q$ are called *breakpoints*. In the tropical polyhedron depicted in Figure 4.3, 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.

7.1.1 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 $[\xi^1, \xi^2], \dots, [\xi^q, \xi^{q+1}]$ 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 E. This imposes strong combinatorial conditions on the tangent digraphs, which, in turn, allows to easily determine the feasible directions.

We introduce some additional basic notions and notations on directed graphs. Two nodes of a digraph are said to be *connected* if they are connected in the underlying undirected graph. A *connected component* is a set of nodes that are pair-wise connected. 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.

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.

Let us start at the basic point $\xi = x^{K \cup \{i_{\mathsf{out}}\}}$. We shall prove below that, at every basic point, the tangent digraph is spanning tree where every hyperplane node is of degree $(1, 1)$. In other words, for every $i \in K \cup \{i_{\mathsf{out}}\}$, the sets $\arg(W_i^+ \odot \xi)$ and $\arg(W_i^- \odot \xi)$ 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, $\vec{C}_+$, contains $j_{i_{\mathsf{out}}}^+$, and the second one, $\vec{C}_-$ contains $j_{i_{\mathsf{out}}}^-$. From the discussion above, it follows that the set J consists of the coordinate nodes in $\vec{C}^+$.

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, 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 happens, depending 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 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 belong to J, while k must belong to the component $\vec{C}_-$. Hence, the arc a_{new} "reconnects" the two components $\vec{C}_+$ and $\vec{C}_-$. 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 $a_{\mathsf{new}} = (j_{\mathsf{new}}, k)$; a 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 incidents arcs $a_{\mathsf{new}}, a_{\mathsf{old}}$ and a'. By Proposition 7.1, we know that $J \subseteq J'$, hence we must increase the variable j_{new}. Since we want to stay inside the hyperplane $\mathcal{H}_k$, 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 $i \in K \setminus \{k\}$ are of degree $(1, 1)$. Removing the arc a_{old} from the graph provides two connected components, the first one $\vec{C}'_+$ contains the coordinate nodes j_{new}, l as well as the hyperplane node k, while the second one $\vec{C}'_-$ contains j_{old}. The new direction set J' is given by the coordinate nodes in $\vec{C}'_+$.

The tangent digraph in the open segment $]\xi', \xi''[$ is again constant, and 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 $\vec{C}'_+$ and $\vec{C}'_-$, 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 $\vec{C}'_+$ and $\vec{C}'_-$ are reconnected by adding i_{ent} and its two incident arcs.

7.1.2 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 $\{x + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ is included in $\mathcal{D}$. The following lemma will be helpful to prove the feasibility of a direction.

Lemma 7.2. *Let $x \in \mathbb{T}^{n+1}$ with no $\mathbb{0}$ entries. Then, the following properties hold:*

(i) if x belongs to $\mathcal{H}_i^{\geq} \setminus \mathcal{H}_i$, every direction is feasible from x in $\mathcal{H}_i^{\geq}$.

(ii) if x belongs to $\mathcal{H}_i$, the direction e^J is feasible from x in the half-space $\mathcal{H}_i^{\geq}$ if, and only if, $\arg(W_i^+ \odot x) \cap J \neq \emptyset$ or $\arg(W_i^- \odot x) \cap J = \emptyset$.

(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 $\arg(W_i^+ \odot x) \cap J$ and $\arg(W_i^- \odot x) \cap J$ are both empty or both non-empty.

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 Assumption A and the fact that $\boldsymbol{x}$ has no $\mathbb{0}$ entries. Then, for $\lambda > 0$ sufficiently small, we have:

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

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 E, these graphs do not contain any cycle by Lemma 4.26. In other words, they are forests: each connected component is a tree. For such graphs, the following is known:

$$\text{number of nodes} = \text{number of edges} + \text{number of connected components}\,. \tag{7.1}$$

Proposition 7.3. *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 $\arg(W_{i_{\mathsf{out}}}^+ \odot x)$ 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 $n' \geq n-1$. As there is at least one connected component, there is at most $n+n'$ edges by (7.1). Besides, 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 $n' \leq n$. As a result, by using (7.1), 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. Besides, 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 described in Proposition 7.3.

Case (i): Since $n' = n$, the set of hyperplanes 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 hyperplanes 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 $\mathcal{G}_x$ towards 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 4.25, the submatrix W' of W made with columns in $[n]$ and rows in $K \cup \{i_{\mathsf{out}}\}$ satisfies $\operatorname{tper}(|W'|) > \mathbb{0}$. 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 $\arg(W^+_{i_{\mathsf{out}}} \odot x)$ and in $\arg(W^-_{i_{\mathsf{out}}} \odot x)$, 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 $\arg(W^+_i \odot x) \subseteq J$ and $\arg(W^-_i \odot x) \subseteq J$. In contrast, if the node $i \in K$ belongs to C^-, we have $\arg(W^+_i \odot x) \cap J = \arg(W^-_i \odot x) \cap J = \emptyset$. By Lemma 7.2, 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 $\mathcal{H}^{\geq}_{i_{\mathsf{out}}}$, since $x \in \mathcal{H}_{i_{\mathsf{out}}}$ and $\arg(W^+_{i_{\mathsf{out}}} \odot x) \subseteq J$. Finally, for all $i \notin K \cup \{i_{\mathsf{out}}\}$, the point x belongs $\mathcal{H}^{\geq}_i \setminus \mathcal{H}_i$. Indeed, if $x \in \mathcal{H}_i$, then i would be a hyperplane node. Thus, by Lemma 7.2, the direction e^J is feasible in $\mathcal{H}^{\geq}_i$. As $\mathcal{E}_K = (\cap_{i \in K} \mathcal{H}_i) \cap (\cap_{i \notin K} \mathcal{H}^{\geq}_i)$, 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 7.2 shows that the direction e^J is feasible from x in $\mathcal{E}_K$. Indeed, the point x belongs to $\mathcal{H}^{\geq}_i \setminus \mathcal{H}_i$ for $i \notin K$. Besides, for all $i \in K$, the sets $\arg(W^+_i \odot x) \cap J$ and $\arg(W^-_i \odot x) \cap J$ are both non-empty if i belongs to C_1, and both 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 7.2. Indeed, $x \in \mathcal{H}^{\geq}_i \setminus \mathcal{H}_i$ for $i \notin K$. Besides, if $i \in K$, the sets $\arg(W^+_i \odot x) \cap J$ and $\arg(W^-_i \odot x) \cap J$ are both non-empty 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 non-opposite directions which are feasible from x in $\mathcal{E}_K$, which implies than x is a breakpoint. □

Example 7.4. Figure 4.7 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 7.3. 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 $e^{\{1,3,4\}}$.

7.1.3 Moving along an ordinary segment

We now characterize the length μ of an ordinary segment $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ of a tropical edge $\mathcal{E}_K$. We shall see that the tangent digraph is constant in $]\xi, \xi'[$ and that it "acquires" a new arc or a new hyperplane node when the endpoint ξ' is reached. Modifications to the tangent digraph are determined by the following scalars. For all $i \in [m]$, we define:

$$\lambda_i^+(\xi, J) := (|W_i| \odot \xi) - \max_{j\in J}(W_{ij}^+ + \xi_j) ,$$
$$\lambda_i^-(\xi, J) := (|W_i| \odot \xi) - \max_{j\in J}(W_{ij}^- + \xi_j) ,$$

where $W = (W_{ij})$.

By Assumptions A and F, 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 $\mathbb{0}$, in which case we use the convention $\lambda_i^+ = +\infty$ and $\lambda_i^- = +\infty$, respectively. When $\max_{j\in J}(W_{ij}^+ + \xi_j)$ and $\max_{j\in J}(W_{ij}^- + \xi_j)$ are different from $\mathbb{0}$, the scalars $\lambda_i^+(\xi, J)$ and $\lambda_i^-(\xi, J)$ are non-negative elements of the group G, where $\mathbb{T} = \mathbb{T}(G)$. 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^-.

The scalars λ_i^+ and λ_i^+ tells us when the tangent digraph changes, *i.e.*, when the set $\arg(|W_i| \odot x^\lambda)$ is modified. Indeed, let us denote $x^\lambda = \xi + \lambda e^J$, and observe how

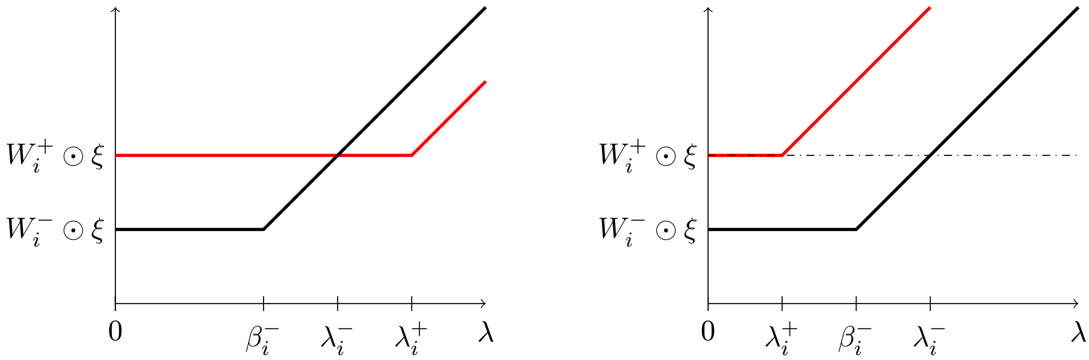

Figure 7.1: Evolution of $W_i^+ \odot (\xi + \lambda e^J)$ (in red) and $W_i^- \odot (\xi + \lambda e^J)$ (in black) with $\lambda \geq 0$ when $\lambda_i^- < \lambda_i^+$ (left) or $\lambda_i^- > \lambda_i^+$ (right).

$\arg(|W_i| \odot x^\lambda)$ varies with $\lambda \geq 0$. For any $i \in [m]$, we have:

$$\arg(|W_i| \odot x^\lambda) = \begin{cases} \arg(|W_i| \odot \xi) & \text{for } \lambda < \min(\lambda_i^+, \lambda_i^-) \\ \arg(|W_i| \odot \xi) \cup \arg\max_{j \in J}(|W_{ij}| + \xi_j) & \text{for } \lambda = \min(\lambda_i^+, \lambda_i^-) \\ \arg\max_{j \in J}(|W_{ij}| + \xi_j) & \text{for } \lambda > \min(\lambda_i^+, \lambda_i^-) \end{cases} \tag{7.2}$$

When $\min(\lambda_i^+, \lambda_i^-) > 0$, then $\arg\max_{j \in J}(|W_{ij}| + \xi_j) \cap \arg(|W_i| \odot \xi) = \emptyset$. Hence, $\arg(|W_i| \odot x^\lambda)$ is constant for $\lambda < \min(\lambda_i^+, \lambda_i^-)$, and gains at least one new element at $\lambda = \min(\lambda_i^+, \lambda_i^-)$. Otherwise, when $\min(\lambda_i^+, \lambda_i^-) = 0$, the set $\arg\max_{j \in J}(|W_{ij}| + \xi_j)$ is included in $\arg(|W_i| \odot \xi)$. In fact, $\arg\max_{j \in J}(|W_{ij}| + \xi_j) = \arg(|W_i| \odot \xi) \cap J$. In this case, $\arg(|W_i| \odot x^\lambda)$ is constant for all $\lambda > 0$.

The distinction between λ_i^+ and λ_i^+ will tell us whether the elements j that will enter $\arg(|W_i| \odot x^\lambda)$ corresponds to tropically positive entries $W_{ij} \in \mathbb{T}^+$ or to tropically negative entries $W_{ij} \in \mathbb{T}^-$. This distinction is crucial in order to detect when x^λ saturate a new inequality.

Indeed, the interpretation of λ_i^+ and λ_i^- differs when one looks at the evolution of $W_i^+ \odot x^\lambda$ and $W_i^- \odot x^\lambda$ with $\lambda \geq 0$ (see Figure 7.1). We have:

$$\begin{aligned} W_i^+ \odot x^\lambda &= \begin{cases} W_i^+ \odot \xi & \text{if } 0 \leq \lambda \leq \lambda_i^+ \\ (W_i^+ \odot \xi) + \lambda - \lambda_i^+ & \text{if } \lambda \geq \lambda_i^+ \end{cases} \\ W_i^- \odot x^\lambda &= \begin{cases} W_i^- \odot \xi & \text{if } 0 \leq \lambda \leq \beta_i^- \\ (W_i^+ \odot \xi) + \lambda - \lambda_i^- & \text{if } \lambda \geq \beta_i^- \end{cases} \end{aligned} \tag{7.3}$$

where $\beta_i^- = \lambda_i^- + (W_i^- \odot \xi) - (W_i^+ \odot \xi)$. In particular $\beta_i^- \leq \lambda_i^-$ and equality holds when $i \in K$.

The endpoint ξ' of the segment $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ 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:

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

Note that $\mathsf{Ent}(\xi, J)$ can also be defined as the set of $i \in [m] \setminus K$ such that $\lambda_i^+ > 0$.

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

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

Alternatively, $i \in K$ belongs to $\mathsf{Br}(\xi, J)$ if and only if $\min(\lambda_i^+, \lambda_i^-) > 0$.

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 analog. 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 notations Br and Ent.

Proposition 7.5. *Let $\{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ 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 $\lambda \geq 0$ satisfying the following conditions:

$$\begin{aligned} &\lambda \leq \min(\lambda_i^+, \lambda_i^-) && \textit{for all } i \in \mathsf{Br} \ , \\ &\lambda \leq \lambda_i^- && \textit{for all } i \in \mathsf{Ent} \textit{ such that } \lambda_i^- \leq \lambda_i^+ . \end{aligned} \tag{7.4}$$

(ii) if $\mu = \lambda_{i_{\mathsf{ent}}}^-$ for some $i_{\mathsf{ent}} \in \mathsf{Ent}$, 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 some $k \in \mathsf{Br}$, then $\xi + \mu e^J$ is a breakpoint.

Proof. Let $x^\lambda := \xi + \lambda e^J$ for all $\lambda \geq 0$. First, We claim that x^λ belongs to $\mathcal{E}_K$ if λ satisfies (7.4).To that end, we shall use repeatedly the evolution of $W_i^+ \odot x^\lambda$ and $W_i^+ \odot x^\lambda$ with λ described in (7.3). We need to show that $x^\lambda \in \mathcal{H}_i$ for $i \in K$ and that $x^\lambda \in \mathcal{H}_i^{\geq}$ for $i \in [m] \setminus K$. Consider an $i \in \mathsf{Br}$. Then $\beta_i^- = \lambda_i^-$. Therefore, for all $0 \leq \lambda \leq \min(\lambda_i^+, \lambda_i^-)$ 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 7.2, $\arg(W_i^+ \odot \xi) \cap J$ and $\arg(W_i^- \odot \xi) \cap J$ are both non-empty. Thus $\lambda_i^+ = \lambda_i^- = \beta_i^- = 0$. Therefore, $x^\lambda \in \mathcal{H}_i$ for all $\lambda \geq 0$ 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 $\mathcal{H}_i^{\geq}$ where $i \in [m] \setminus K$. If $i \notin \mathsf{Ent}$ then $\arg(W_i^+ \odot \xi) \cap J \neq \emptyset$. Consequently, $\lambda_i^+ = 0$. Thus $x^\lambda \in \mathcal{H}_i^{\geq}$ for all $\lambda \geq 0$ as we have:

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

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

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

Now if furthermore $\lambda_i^+ < \lambda_i^-$, then, for $\lambda \geq \lambda_i^+$, we have

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

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

Second, we claim that the solution set of the inequalities (7.4) admits a greatest element $\lambda^* \in \mathbb{R}$. By contradiction, suppose that $x^\lambda \in \mathcal{E}_K$ for all $\lambda \geq 0$. Recall that e^J and $-e^{[n+1]\setminus J}$ coincide as elements of $\mathbb{TP}^n$. Consequently the half-ray $\{\xi - \lambda e^{[n+1]\setminus J} \mid \lambda \geq 0\}$ 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 F.

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

(a) $\lambda^* = \lambda^-_{i_{\mathsf{ent}}} \leq \lambda^+_{i_{\mathsf{ent}}}$ for some $i_{\mathsf{ent}} \in \mathsf{Ent}$. Then $W^+_{i_{\mathsf{ent}}} \odot x^{\lambda^*} = W^-_{i_{\mathsf{ent}}} \odot x^{\lambda^*}$ by (7.3). Moreover, $W^+_{i_{\mathsf{ent}}} \odot \xi > \mathbb{0}$ by Assumptions A and F. As a consequence, $i_{\mathsf{ent}} \notin K$ is a hyperplane node in the tangent graph $\mathcal{G}_{x^{\lambda^*}}$. By Proposition 7.3, we conclude that x^{λ^*} is a basic point for the set $K \cup \{i_{\mathsf{ent}}\}$.

(b) $\lambda^* = \min(\lambda_k^+, \lambda_k^-)$ for some $k \in \mathsf{Br}$. In that case, by (7.2), we have:

$$\arg(|W_i| \odot x^{\lambda^*}) = \arg(|W_i| \odot \xi) \cup \arg\max_{j \in J} \, (|W_{ij}| + \xi_j) \; .$$

The hyperplane node $k \in K$ has at least two incident arcs in $\vec{\mathcal{G}}_\xi$ by Proposition 7.3. Consequently, the set $\arg(|W_i| \odot \xi)$ contains at least two elements. Moreover, $\arg\max_{\, j \in J}(|W_{ij}| + \xi_j)$ contains at least one element. Hence, the set $\arg(|W_i| \odot x^{\lambda^*})$ contains at least three elements,*i.e.*, in the tangent digraph $\vec{\mathcal{G}}_{x^{\lambda^*}}$, the hyperplane node $k \in K$ has at least three incident arcs. By Proposition 7.3, the point x^{λ^*} must be a breakpoint.

Note that the cases (a) and (b) above also prove (ii) and (iii). □

Example 7.6. We now have all the ingredients required to perform a tropical pivot. Feasible directions are given by Proposition 7.3, while Proposition 7.5 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 4.7. By Proposition 7.3 (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 homogenous setting, the inequalities $\mathcal{H}_3$ and $\mathcal{H}_4$ read

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

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 $\mu = \lambda^-_{\mathcal{H}_2} \leq \lambda^+_{\mathcal{H}_2}$. 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.

7.1.4 Incremental update of the tangent digraph

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 $O(n^2 m)$.

Proposition 7.7. *Let $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ be an ordinary segment of $\mathcal{E}_K$.*

(i) every point in $]\xi, \xi'[$ has the same tangent digraph $\vec{\mathcal{G}}_{]\xi,\xi'[}$, which is a subgraph of both $\vec{\mathcal{G}}_\xi$ and $\vec{\mathcal{G}}_{\xi'}$.

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

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

(iii) if ξ' is a breakpoint, then there exists a unique $k^* \in \mathsf{Br}$ such that $\mu = \min(\lambda^+_{k^*}, \lambda^-_{k^*})$, and the set $\arg\max_{j \in J}(|W_{k^* j}| + \xi_j)$ is reduced to a singleton $\{l^*\}$. Moreover,

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

where a_{new} is an arc between k^* and l^*, oriented from l^* to k^* if $\lambda^+_{k^*} < \lambda^-_{k^*}$, and from k^* to l^* otherwise.

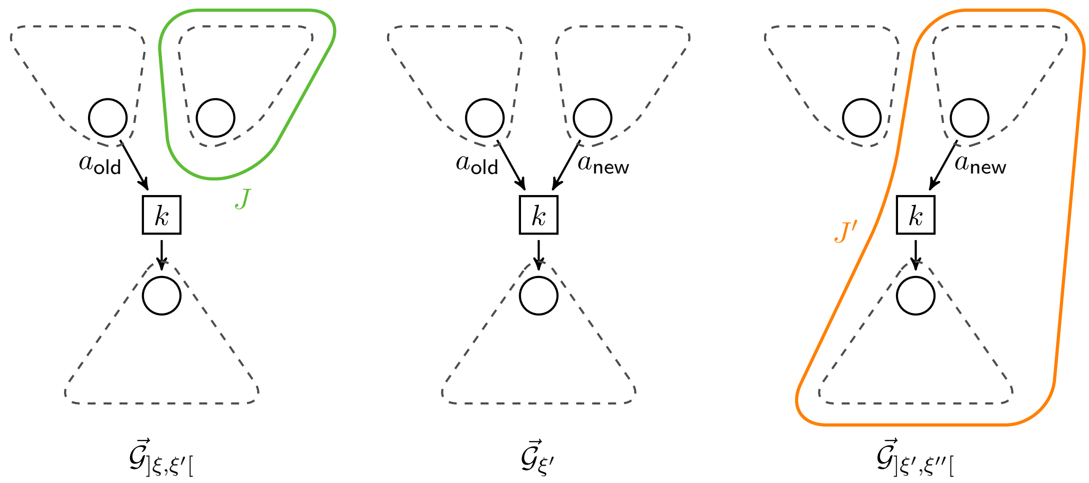

Figure 7.2: Illustration of Proposition 7.7 (iii) and (iv), with a sequence of tangent digraphs around a breakpoint ξ' between two consecutive segments $[\xi, \xi'] \cup [\xi', \xi'']$. The direction of $[\xi, \xi']$, from ξ to ξ', is given by the set of coordinate nodes J, indicated in green. The direction of the second segment, from ξ' to ξ'', is governed by J' depicted in orange.

(iv) 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 (iii) is given in Figure 7.2.

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

(i) Any point in $]\xi, \xi'[$ is of the form x^λ for some $0 < \lambda < \mu$. Consider such a λ. By Proposition 7.3, the tangent digraph $\vec{\mathcal{G}}_{x^\lambda}$ admits $[n+1]$ as its set of coordinate nodes, and the set of hyperplane nodes always contains K.

We now prove that the set of arcs is constant, *i.e.*, we show that for any $i \in K$, the set $\arg(|W_i| \odot x^\lambda)$ does not depend on $\lambda \in]0, \mu[$. Consider a $i \in \mathsf{Br}$. We have $\lambda < \mu$, then in particular $\lambda < \min(\lambda_i^+, \lambda_i^-)$ by Proposition 7.5. Hence, we have $\arg(|W_i| \odot x^\lambda) = \arg(|W_i| \odot \xi)$ by (7.2). Otherwise, let $i \in K \setminus \mathsf{Br}$. Then, $\arg(W_i^+ \odot \xi) \cap J$ and $\arg(W_i^- \odot \xi) \cap J$ are both non-empty, by Lemma 7.2. Consequently, $\min(\lambda_i^+, \lambda_i^-) = 0$ by definition of λ_i^+, λ_i^-. It follows that $\lambda > \min(\lambda_i^+, \lambda_i^-)$. Hence, $\arg(|W_i| \odot x^\lambda) = \arg\max_{j \in J}(|W_{ij}| + \xi_j)$ by (7.2).

(ii) By Proposition 7.3 (C2), $\vec{\mathcal{G}}_{]\xi,\xi'[}$ does not contain the hyperplane node i_{out}. As $\vec{\mathcal{G}}_{]\xi,\xi'[}$ is a subdigraph of $\vec{\mathcal{G}}_\xi$ by (i), we deduce that it is also a subdigraph of $\vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\}$. By

Proposition 7.3 again, the only subdigraph of $\vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\}$ that can be a tangent digraph at a point in $\mathcal{E}_K$ is $\vec{\mathcal{G}}_\xi \setminus \{i_{\mathsf{out}}\}$.

(iii) Since ξ' is a breakpoint, we have $\mu = \min(\lambda_k^+, \lambda_k^-)$ for some $k \in \mathsf{Br}$ by Proposition 7.5. First assume that $\mu = \lambda_k^+ < \lambda_k^-$ for some $k \in \mathsf{Br}$. In that case, observe that

$$\arg\max_{j\in J} \, (|W_{kj}| + \xi_j) = \arg\max_{j\in J} \, (W^+_{kj} + \xi_j) \enspace .$$

Let $l \in \arg\max_{j\in J}(W^+_{lj} + \xi_j)$. Then for all $0 < \lambda < \mu$, we have $W^+_{kl} + x^\lambda_l < W^+_k \odot x^\lambda$, while $W^+_{kl} + x^\mu_l = W^+_k \odot x^\mu$. It follows that the arc (l,k) does not belong to $\vec{\mathcal{G}}_{]\xi,\xi'[}$, whereas it appears in $\vec{\mathcal{G}}_{\xi'}$, oriented from l to k. We deduce that $\vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{(l,k)\}$ is a subgraph of $\vec{\mathcal{G}}_{\xi'}$ by (i). Both are equal by Proposition 7.3. Moreover, if $\arg\max_{j\in J}(W^+_{k^*j} + \xi_j)$ contains two distincts elements l, l'. Then, by the argument above, $\vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{(l,k),(l',k)\}$ is a subgraph of $\vec{\mathcal{G}}_{\xi'}$. This contradicts Proposition 7.3.
Second, If $\mu = \lambda^-_{k^*} < \lambda^+_{k^*}$, then the arguments above show that $\vec{\mathcal{G}}_{\xi'} = \vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{(k,l)\}$, where l is the unique element in the set

$$\arg\max_{j\in J} \, (|W_{k^*j}| + \xi_j) = \arg\max_{j\in J} \, (W^-_{k^*j} + \xi_j)$$

Third, if $\lambda^-_k = \lambda^+_k$, then, by the arguments above, the hyperplane node k would have at least two incoming and two outgoing arcs in the tangent digraph at ξ', a contradiction with Proposition 7.3.
Finally, suppose that $\mu = \min(\lambda^+_k, \lambda^-_k) = \min(\lambda^+_{k'}, \lambda^-_{k'})$ for two distincts $k, k' \in \mathsf{Br}$. Then, the hyperplane nodes k and k' would both have at least three adjacent arcs in $\vec{\mathcal{G}}_{\xi'}$, again a contradiction with Proposition 7.3.

(iv) By applying i to the segment $[\xi', \xi'']$, we know that $\vec{\mathcal{G}}_{]\xi',\xi''[}$ is a subdigraph of $\vec{\mathcal{G}}_{\xi'}$. By Proposition 7.3, 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 $\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$. Indeed, the segment $[\xi', \xi'']$ is directed by J'. By Proposition 7.3, the set J' correspond to the coordinate nodes in one of the connected components of $\vec{\mathcal{G}}_{]\xi',\xi''[}$. Similarly, the set J governing the direction of $[\xi, \xi']$ correspond to a connected component in $\vec{\mathcal{G}}_{]\xi,\xi'[}$. By Proposition 7.1, we have $J \neq J'$. Consequently, the graphs $\vec{\mathcal{G}}_{]\xi,\xi'[}$ and $\vec{\mathcal{G}}_{]\xi',\xi''[}$ must be distinct. □

Proposition 7.8. *Let $[\xi, \xi'] \cup [\xi', \xi'']$ be two consecutive ordinary segments of $\mathcal{E}_K$, where $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ and $[\xi', \xi''] = \{\xi' + \lambda e^{J'} \mid 0 \leq \lambda \leq \mu'\}$. Moreover, let k^* be the unique hyperplane node of $\vec{\mathcal{G}}_{\xi'}$ of degree $(2,1)$ or $(1,2)$ and let $a_{\mathsf{old}}, a_{\mathsf{new}}$ be the two arcs incident to k^* with the same orientation. Denote by $\vec{\mathcal{D}}$ the connected component of $\vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}, a_{\mathsf{new}}\}$ that contains k^*. Then:*

(i) $J' = J \cup \{j \in [n+1] \mid j$ *is a coordinate node in* $\vec{\mathcal{D}}\}$

(ii) $\mathsf{Br}(\xi', J') = \mathsf{Br}(\xi, J) \setminus \{i \in [m] \mid i$ *is a hyperplane node in* $\vec{\mathcal{D}}\}$

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

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

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

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

Proof. (i) According to Proposition 7.3 (C2), the digraph $\vec{\mathcal{G}}_{]\xi',\xi''[}$ consists of two weakly connected components, $\vec{C}_+$ and $\vec{C}_-$, and J is the set of coordinate nodes in one of these components, say $\vec{C}_+$.
Let $l^* \in J$ be the coordinate node incident to a_{new}, as described in Proposition 7.7 (iii). The tangent digraph at ξ' is equal to $\vec{\mathcal{G}}_{\xi'} = \vec{C}_+ \cup \vec{C}_- \cup \{a_{\mathsf{new}}\}$. Since $\vec{\mathcal{G}}_{\xi'}$ is connected, and $l^* \in \vec{C}_+$, the hyperplane node k^* belongs to $\vec{C}_-$. Thus the arc a_{old} also belongs to $\vec{C}_-$. Observe that $\vec{\mathcal{D}}$ is a subgraph of $\vec{C}_-$. In fact, $\vec{C}_- \setminus \{a_{\mathsf{old}}\}$ can be decomposed into two connected components $\vec{C}'_-$ and $\vec{\mathcal{D}}$, where $\vec{\mathcal{D}}$ contains k^*.
In the next segment, the tangent digraph is $\vec{\mathcal{G}}_{]\xi',\xi''[} = \vec{C}_+ \cup \vec{C}_- \cup \{a_{\mathsf{new}}\} \setminus \{a_{\mathsf{old}}\}$. It consists of two connected components. Let $\vec{C}'_+$ denote the component that contains the hyperplane node k^*. Then observe that $\vec{C}'_+ = \vec{C}_+ \cup \vec{\mathcal{D}} \cup \{a_{\mathsf{new}}\}$. Moreover, the second connected component of $\vec{\mathcal{G}}_{]\xi',\xi''[}$ is $\vec{C}'_-$.
The two feasible directions in $]\xi', \xi''[$, are $e^{J'}$ and $-e^{J'} \equiv e^{[n+1]\setminus J}$. The set J' is set of coordinate nodes in either $\vec{C}'_+$ or $\vec{C}'_-$, by Proposition 7.3 (C2). We know that $J \subseteq J'$ by Proposition 7.1. Hence J' is the set of coordinate nodes in $\vec{C}'_+$ and thus $J' = J \cup \Delta$.

(ii) By definition of Br, we have $\min(\lambda_i^+(\xi, J), \lambda_i^-(\xi, J)) > 0$ for all $i \in \mathsf{Br}(\xi, J)$. Using (7.2), it follows that $\arg(|W_i| \odot (\xi + \lambda e^J)) = \arg(|W_i| \odot \xi)$ for all $\lambda > 0$ small enough. Consequently, $\mathsf{Br}(\xi, J) = \mathsf{Br}(\xi + \lambda e^J, J)$ for $\lambda > 0$ small enough. Hence, $\mathsf{Br}(\xi, J)$ is exactly the set of hyperplane nodes in the connected component $\vec{C}_-$ of $\vec{\mathcal{G}}_{]\xi,\xi'[}$, where $\vec{C}_-$ is defined above. Similarly, $\mathsf{Br}(\xi', J')$ is exactly the set of hyperplane nodes in the connected component $\vec{C}'_-$ of $\vec{\mathcal{G}}_{]\xi',\xi''[}$. The difference between these two sets corresponds to the hyperplane nodes in $\vec{\mathcal{D}}$.

(iii) First observe that $\mathsf{Ent}(\xi', J') \subseteq \mathsf{Ent}(\xi, J)$. Indeed, consider an $i \in K \setminus \mathsf{Ent}(\xi, J)$. Then $\arg(W_i^+ \odot \xi) \cap J \neq \emptyset$, which implies $\arg(W_i^+ \odot \xi') \subseteq J$. Using the inclusion $J \subseteq J'$, we obtain that $\arg(W_i^+ \odot \xi') \cap J' \neq \emptyset$, and therefore $i \notin \mathsf{Ent}(\xi', J')$.
Second if $i \in \mathsf{Ent}(\xi, J)$ satisfies $\mu \geq \lambda_i^+(\xi, J)$ then $\arg(W_i^+ \odot \xi')$ intersects $J \subseteq J'$, thus $i \notin \mathsf{Ent}(\xi', J')$. As a consequence:

$$\mathsf{Ent}(\xi', J') \subseteq \{i \in \mathsf{Ent}(\xi, J) \mid \mu < \lambda_i^+(\xi, J)\}. \tag{7.5}$$

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

(iv) Using (7.5) let us consider an $i \in \mathsf{Ent}(\xi, J)$ such that $\mu < \lambda_i^+(\xi, J)$. Then, as above, $\arg(W_i^+ \odot \xi') = \arg(W_i^+ \odot \xi)$. Moreover, $i \in \mathsf{Ent}(\xi, J)$ implies $\arg(W_i^+ \odot \xi) \cap J = \emptyset$. Thus $\arg(W_i^+ \odot \xi') \cap J' = \emptyset$ if and only if $\arg(W_i^+ \odot \xi) \cap (J' \setminus J) = \emptyset$.

(v) 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 (7.5). Otherwise, if $i \in \mathsf{Br}(\xi', J')$, then $i \in \mathsf{Br}(\xi, J)$ by (ii) and thus $\mu \leq \lambda_i^+(\xi, J)$ by (7.4). 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 $\lambda_i^-(\xi', J')$. □

7.1.5 Linear-time pivoting

We now present an algorithm (Algorithm 5) allowing to move along an ordinary segment $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$ 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 $]\xi, \xi'[$, the sets $\mathsf{Ent}(\xi, J)$ and $\mathsf{Br}(\xi, J)$, *etc.* We also define, for $j \in [m]$, the sets

$$\delta_j(\xi, J) := \{i \in \mathsf{Ent}(\xi, J) \mid j \in \arg(W_i^+ \odot \xi)\} \ .$$

It also uses a Boolean matrix M, such that $M_{ij} = \textbf{true}$ for the pairs $(i, j) \in \mathsf{Ent}(\xi, J) \times [n+1]$ if and only if $j \in \arg(W_i^+ \odot \xi)$. We shall see in the main pivoting algorithm that we will not need to update this matrix when pivoting over the whole tropical edge.

Algorithm 5 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 5, 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, thus iterating over its elements can be done in $O(|\Delta|)$.

Proposition 7.9. *Algorithm 5 is correct, and its time complexity is bounded by $O(n + m|J' \setminus J|)$.*

Proof. The correctness of the highlighted parts of the algorithm straightforwardly follows from the corresponding results given in annotations.

Complexity: At Lines 8 and 10, 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 9 amounts to iterate over the arcs incident to k^*, and there is exactly 3 such arcs by Proposition 7.3.

Computing the sets Δ and Ξ between Lines 11 and 14 uses $O(n)$ operations, as the graph $\vec{\mathcal{G}}_{\xi'}$ contains $O(n)$ nodes and edges. Moreover, the sets $\Delta \subseteq [n+1]$ and $\Xi \subseteq K$ are of size $O(n)$, thus updating J and Br uses $O(n)$ operations.

At Line 15, we visit the $O(m)$ elements $\mathsf{Ent}(\xi, J)$. For each $i \in \mathsf{Ent}(\xi, J)$, we first test in $O(1)$ whether $\mu < \lambda_i^+(\xi, J)$. Second, we iterate over the elements $j \in \Delta$ and test whether $j \in \arg(W_i^+ \odot \xi)$ using the Boolean matrix M. Since there is $|J' \setminus J|$ elements in Δ, and since any entry of M can be accessed in $O(1)$, we obtain an overall complexity of $O(m|J' \setminus J|)$.

Computations at Lines 18 and 19 are done by iterating over elements $j \in \Delta$ 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 to in constant time. Therefore, $\lambda_i^+(\xi', J')$ and $\lambda_i^-(\xi', J')$ are computed in time $O(|\Delta|) = O(|J' \setminus J|)$. The complexity of other operations is easily obtained. In total, the complexity of the algorithm is $O(n + m|J' \setminus J|)$. □

Theorem 7.10. *Algorithm 6 allows 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 matrix M 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 7.8, we have the inclusion $\mathsf{Ent}(\xi', J') \subseteq \mathsf{Ent}(\xi, J)$ and the equality $\arg(W_i^+ \odot \xi') = \arg(W_i^+ \odot \xi)$ for all $i \in \mathsf{Ent}(\xi', J')$. It follows that if M_{ij} determines whether $j \in \arg(W_i^+ \odot \xi)$ for all $i \in \mathsf{Ent}(\xi, J)$, it can be used as well to determine whether $j \in \arg(W_i^+ \odot \xi')$ for all $i \in \mathsf{Ent}(\xi', J')$.

Then, the correctness of the algorithm follows from Proposition 7.7 (ii) (for the computation of $\vec{\mathcal{G}}_{]\xi^1, \xi^2[}$ at Line 2), Proposition 7.3 (for the computation of J at Line 3) and Proposition 7.9.

The complexity of the operations from Lines 1 to 7 can easily be verified to be in $O(mn)$. Let $q \leq n$ be the number of iterations of the loop from Lines 8 and 11, and let

Algorithm 5: 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 $[\xi, \xi'] = \{\xi + \lambda e^J \mid 0 \leq \lambda \leq \mu\}$
- 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 $i \in \mathsf{Br}(\xi, J) \cup \mathsf{Ent}(\xi, J)$
- a Boolean matrix M such that $M_{ij} = \textbf{true}$ only for the $i \in \mathsf{Ent}(\xi, J)$ and $j \in [n+1]$ such that $j \in \arg(W_i^+ \odot \xi)$

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 $[\xi', \xi'']$
- 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 7.5 (i)–(ii)

$\mu \leftarrow \min\{\min(\lambda_i^+(\xi, J), \lambda_i^-(\xi, J)) \mid i \in \mathsf{Br}(\xi, J) \text{ or } (i \in \mathsf{Ent}(\xi, J) \text{ and } \lambda_i^-(\xi, J) \leq \lambda_i^+(\xi, J))\}$ $O(m)$
$\xi' \leftarrow \xi + \mu e^J$ $O(n)$
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 7.7(iii)–(iv)

$k^* \leftarrow$ the unique element of $\mathsf{Br}(\xi, J)$ such that $\mu = \min(\lambda_{k^*}^+(\xi, J), \lambda_{k^*}^-(\xi, J))$ (ξ' is a breakpoint)
$l^* \leftarrow$ the unique element in $\arg\max_{j \in J} |W_{k^* j}| + \xi_j$ $O(n)$
$a_{\mathsf{new}} \leftarrow$ the arc from l to k^* if $\lambda_{k^*}^+(\xi, J) < \lambda_{k^*}^-(\xi, J)$, the arc from k^* to l otherwise $O(1)$
$\vec{\mathcal{G}}_{\xi'} \leftarrow \vec{\mathcal{G}}_{]\xi,\xi'[} \cup \{a_{\mathsf{new}}\}$ $O(n)$
$a_{\mathsf{old}} \leftarrow$ the only arc incident to k^* in $\vec{\mathcal{G}}_{\xi'}$ with the same orientation as a_{new} $O(1)$
$\vec{\mathcal{G}}_{]\xi',\xi''[} \leftarrow \vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}\}$ $O(n)$

Proposition 7.8 (i)–(ii)

$\mathit{\Delta} \leftarrow$ coordinate nodes of $\vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}, a_{\mathsf{new}}\}$ connected to k^* $O(n)$
$\mathit{\Xi} \leftarrow$ hyperplane nodes of $\vec{\mathcal{G}}_{\xi'} \setminus \{a_{\mathsf{old}}, a_{\mathsf{new}}\}$ connected to k^* $O(n)$
$J' \leftarrow J \cup \mathit{\Delta}$ $O(n)$
$\mathsf{Br}(\xi', J') \leftarrow \mathsf{Br}(\xi, J) \setminus \mathit{\Xi}$ $O(n)$

Proposition 7.8 (iv)–(v)

$\mathsf{Ent}(\xi', J') \leftarrow \{i \in \mathsf{Ent}(\xi, J) \mid \mu < \lambda_i^+(\xi, J) \text{ and } \arg(W_i^+ \odot \xi) \cap \mathit{\Delta} = \emptyset\}$ $O(m|J' \setminus J|)$ using the matrix M
for $i \in \mathsf{Ent}(\xi', J') \cup \mathsf{Br}(\xi', J')$ **do** $O(m)$ iterations
  $W_i^+ \odot \xi' := W_i^+ \odot \xi$ $O(1)$
  $\lambda_i^+(\xi', J') := \min\Big(\lambda_i^+(\xi, J) - \mu \ , \ (W_i^+ \odot \xi) - \max_{j \in \mathit{\Delta}}(W_{ij}^+ + \xi_j)\Big)$ $O(|J' \setminus J|)$
  $\lambda_i^-(\xi', J') := \min\Big(\lambda_i^+(\xi, J) - \mu \ , \ (W_i^+ \odot \xi) - \max_{j \in \mathit{\Delta}}(W_{ij}^- + \xi_j)\Big)$ $O(|J' \setminus J|)$

return $\xi', J', \vec{\mathcal{G}}_{]\xi',\xi''[}, \mathsf{Ent}(\xi', J'), \mathsf{Br}(\xi', J'), (W_i^+ \odot \xi')_i, (\lambda_i^+(\xi', J'))_i, (\lambda_i^-(\xi', J'))_i$

Algorithm 6: Linear-time tropical 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 $I' = (I \setminus \{i_{\mathsf{out}}\}) \cup \{i_{\mathsf{ent}}\}$

compute $\vec{\mathcal{G}}_{x^I}$ $\quad O(mn)$

$\vec{\mathcal{G}}_{]\xi^1,\xi^2[} \leftarrow \vec{\mathcal{G}}_{x^I} \setminus \{i_{\mathsf{out}}\}$ $\quad O(n)$

$J \leftarrow$ coordinate nodes weakly connected to the element of $\arg(W^+_{i_{\mathsf{out}}} \odot x^I)$ in $\vec{\mathcal{G}}_{]\xi,\xi'[}$ $\quad O(n)$

compute $E \leftarrow \mathsf{Ent}(x^I, J)$ and $B \leftarrow \mathsf{Br}(x^I, J)$ $\quad O(mn)$

compute $W^+_i \odot x^I$, $\lambda^+_i(x^I, J)$ and $\lambda^-_i(x^I, J)$ for all $i \in E \cup B$ $\quad O(mn)$

$M \leftarrow$ a $m \times (n+1)$ matrix defined by $M_{ij} = \begin{cases} \textbf{true} & \text{if } j \in \arg(W^+_i \odot x^I) \\ \textbf{false} & \text{otherwise} \end{cases}$ $\quad O(mn)$

input $\leftarrow x^I, J, \vec{\mathcal{G}}_{]\xi^1,\xi^2[}, E, B, (W^+_i \odot x)_{i \in E \cup B}, (\lambda^+_i(x^I, J))_{i \in E \cup B}, (\lambda^-_i(x^I, J))_{i \in E \cup B}, M$

while true do $\quad$ at most n iterations

  call Algorithm 5 on (*input*, M) and stores the result in *output*

  if *output* is of the form (ξ', i_{ent}) **then return** (ξ', i_{ent})

  else *input* $\leftarrow$ *output*

$e^{J_1}, e^{J_2}, \dots, e^{J_q}$ be the directions of the ordinary segments followed during the successive calls to Algorithm 5. By Proposition 7.9, 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 $O(nm)$. □

7.2 Computing reduced costs

In this section, we introduce the concept of tropical reduced costs, which are merely the signed valuation 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 classical simplex algorithm, as this operation corresponds to the update of the inverse of the basic matrix A_I.

7.2.1 Symmetrized tropical semiring

To define the tropical reduced costs, we need a signed tropical version of the system of linear equations (3.22). To that end, we use a semiring extension of signed tropical numbers called the *symmetrized tropical semiring*, introduced in [Plu90]. It is denoted by $\mathbb{T}_\diamond$, and is defined as the union of $\mathbb{T}_\pm$ with a third copy of $\mathbb{T}$, denoted $\mathbb{T}_\bullet$. The latter is the set of *balanced tropical numbers*. Its elements are written $a^\bullet$, where $a \in \mathbb{T}$. The numbers a, $\ominus a$ and $a^\bullet$ are pairwise distinct unless $a = \mathbb{0}$. Sign and modulus are extended to $\mathbb{T}_\diamond$ by setting $\operatorname{sign}(a^\bullet) = 0$ and $|a^\bullet| = a$.

The addition of two elements $x, y \in \mathbb{T}_\diamond$, 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 $x, y \in \mathbb{T}_\diamond$ 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 are the matrices with entries $A_{ij} \oplus B_{ij}$ and $\bigoplus_k A_{ik} \odot B_{kj}$, respectively.

The set $\mathbb{T}_\diamond$ 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 $x, y \in \mathbb{T}_\diamond$ satisfy the *balance relation* $x \nabla y$ when $x \ominus y$ is a balanced number. Note that

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

The balance relation is extended entry-wise to vectors in $\mathbb{T}_\diamond^n$. In the semiring $\mathbb{T}_\diamond$, 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 $M \in \mathbb{T}_\diamond^{n \times n}$ is given by

$$\operatorname{tdet}(M) = \bigoplus_{\sigma \in \operatorname{Sym}(n)} \operatorname{tsign}(\sigma) \odot M_{1\sigma(1)} \odot \cdots \odot M_{n\sigma(n)}$$

Observe that this definition of the tropical determinant extends the definition given in Section 3.2.1 Also observe that a square matrix of $\mathbb{T}_\pm^{n \times n}$ is sign-generic for the determinant polynomial if and only if $\operatorname{tdet}(M)$ is a balanced number.

Theorem 7.11 (Signed tropical Cramer Theorem [Plu90])**.** *Let $M \in \mathbb{T}_\diamond^{n \times n}$ and $d \in \mathbb{T}_\diamond^n$. Every solution $y \in \mathbb{T}_\pm^n$ of the system of balances*

$$M \odot y \nabla d \tag{7.6}$$

satisfies

$$\operatorname{tdet}(M) \odot y_j \nabla (\ominus \mathbb{1})^{\odot n+j} \odot \operatorname{tdet}(M_{\widehat{j}} \ d), \qquad \text{for all } j \in [n] \ .$$

Conversely, if the tropical determinants $\operatorname{tdet}(M)$ and $\operatorname{tdet}(M_{\widehat{j}} \ d)$ for $j \in [n]$ are not balanced elements, then the vector with entries $y_j = (\ominus \mathbb{1})^{\odot n+j} \odot \operatorname{tdet}(M_{\widehat{j}} \ d) \odot (\operatorname{tdet}(M))^{\odot -1}$ is the unique solution of (7.6) in $\mathbb{T}_\pm^n$.

This result was proved in [Plu90]; see also [AGG09] for a more recent discussion. 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 7.12. 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".

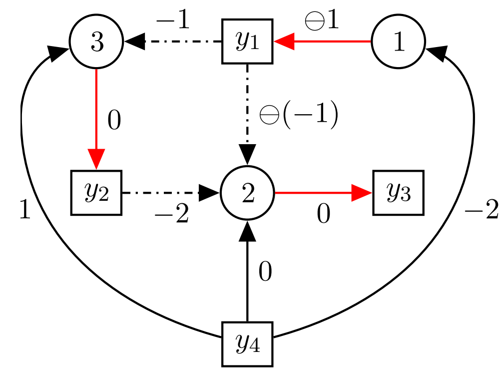

Figure 7.3: The Cramer digraph for the system of balances in (7.7). Column nodes are squares and row nodes are circles. Arcs with weight $-\infty$ are omitted. The maximizing permutation σ is given by the red arcs. Arcs of the digraph of longest paths from the column node y_4 are solid. The coordinate y_j of the signed solution y of (7.7) is obtained by the multiplication (in $\mathbb{T}_\diamond$) of the weight on the longest path from y_4 to y_j.

7.2.2 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 instrumentation of this algorithm, in the special case in which 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 $|\operatorname{tdet}(M)|$. 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 $M^{-1}_{i\sigma(i)}$, 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 7.13. The maximizing permutation for the system of balances (7.7) below is $\sigma(1) = 1, \sigma(2) = 3$ and $\sigma(3) = 2$. The Cramer digraph is represented in Figure 7.3.

$$\begin{pmatrix} \ominus(-1) & -\infty & -\infty \\ -1 & \ominus(-2) & 0 \\ \ominus(-1) & 0 & -\infty \end{pmatrix} \odot \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} \ \nabla \ \begin{pmatrix} -2 \\ 0 \\ -1 \end{pmatrix} \tag{7.7}$$

Note that all the coefficients $M_{i\sigma(i)}$ are different from $\mathbb{0}$. In the sequel, 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 every 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 7.14. *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 following recursive relations*

$$y_{\sigma(i)} = \begin{cases} d_i \odot M^{-1}_{i\sigma(i)} & \textit{when } \gamma(i) = n+1 \ , \\ \ominus M_{i\gamma(i)} \odot M^{-1}_{i\sigma(i)} \odot y_{\gamma(i)} & \textit{otherwise} \end{cases} \tag{7.8}$$

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, Equation (7.8) 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:

$$\begin{aligned} &|M_{i\sigma(i)}| + |y_{\sigma(i)}| \geq |d_i| \ , \\ &|M_{i\sigma(i)}| + |y_{\sigma(i)}| \geq |M_{ij}| + |y_j| \quad \text{for all } j \in [n] \ . \end{aligned}$$

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, (7.8) 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 $M_i \odot y \nabla d_i$. □

A digraph of longest paths for Example 7.13 is shown in Figure 7.3. From the relations (7.8), we obtain the signed solution $y = (\ominus(-1), -1, 0)$.

Complexity analysis

We now discuss the complexity of the method provided by Proposition 7.14. First, a maximizing permutation σ can be found in time $O(n^3)$ by the Hungarian method; see [Sch03, §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, §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 non-positive 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 method, which returns, along with a maximizing permutation σ, a pair of vectors $u, v \in \mathbb{T}^n$ such that

$$\begin{aligned} |M_{ij}| &\leq u_i + v_j \quad \text{for all } i, j \in [n] \ , \\ |M_{i\sigma(i)}| &= u_i + v_{\sigma(i)} \quad \text{for all } i \in [n] \ . \end{aligned} \tag{7.9}$$

The pair (u, v) is in fact an optimal solution to the dual assignment problem:

$$\begin{aligned} \min_{u,v} \quad & \sum_{i=1}^{n} u_i + \sum_{j=1}^{n} v_j \\ & |M_{ij}| \leq u_i + v_j \quad \text{for all } i, j \in [n] \ . \end{aligned}$$

Suppose we have a pair (u, v) satisfying (7.9). 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 following diagonal scaling, $M'_{ij} = \mu^{-1} \odot u_i^{-1} \odot mMij \odot v_j^{-1}$, 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 $M \odot y \ \nabla \ d$ by μ and by u_i, and performing the above change of variables, we arrive at the equivalent system $M' \odot z \ \nabla \ d'$. By choosing $\mu := \max(\max_i(|d_i| - u_i), 0)$, we get that $|d'_i| \leq 0$, and $|M'_{ij}| \leq 0$ for all $i, j \in [n]$. The longest path problem to be solved in order to apply the construction of Proposition 7.14 to $M' \odot z \nabla d'$ now involves a digraph with non-positive 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.

7.2.3 Tropical reduced costs as a solution of a tropical Cramer system

In the rest of this section, we suppose that Assumption F holds, so we only consider basic points x^I with finite entries. We also make the following assumption.

Assumption G. The matrix $(A^T \ c^T)$ is sign-generic for the minor polynomials.

Let I be a feasible basis of the tropical linear program $\mathrm{LP}(A, b, c)$. Consider the system of balances:

$$A_I^\top \odot y \ \nabla \ c^\top \ . \tag{7.10}$$

By Assumption G and Theorem 7.11, the system of balances (7.10) admits a unique solution y^I in $\mathbb{T}_\pm^I$, and this solution coincides with the tropical reduced costs by Proposition 4.34 and. So applying to this system the algorithm described in Section 7.2.2 does provides the vector reduced costs of $\mathrm{LP}(A, b, c)$ for the basis I.

Algorithm 7: 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

1	$\mathcal{G}_{x^I} \leftarrow$ tangent graph at x^I	$O(mn)$
2	$\sigma \leftarrow$ maximizing permutation in $\operatorname{tdet}(A_I)$ obtained by a traversal of $\mathcal{G}_{x^I}$	$O(n)$
3	$u \leftarrow -x^I$	$O(n)$
4	$v \leftarrow A_I^+ \odot x^I$	$O(mn)$
5	$\mu \leftarrow \max(\max_{j\in[n]}(c_j - u_j), 0)$	$O(n)$
6	$M' \leftarrow$ tropically signed matrix with entries $m'_{ij} = \mu^{-1} \odot u_i^{-1} \odot a_{ji} \odot v_j^{-1}$	$O(n^2)$
7	$d' \leftarrow$ tropically signed vector with entries $d_i = \mu^{-1} \odot u_i^{-1} \odot c_i$	$O(n)$
8	$\vec{C} \leftarrow$ Cramer digraph of the system $M' \odot y \nabla d'$ for the permutation σ	$O(n^2)$
9	apply Dijkstra's algorithm to $\vec{C}$ from column node $n+1$	$O(n^2 + n\log(n))$
10	$\gamma \leftarrow$ the tree of longest paths returned by Dijkstra's algorithm	
11	$z \leftarrow$ signed vector obtained by applying (7.8) to the tree γ	$O(n)$
12	**return** y^I the signed vector with entries $y^I_j = v_j \odot z_j$	$O(n)$

Theorem 7.15. *Algorithm 7 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 7.3, Case (i). By Lemma 4.25, 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 7.3), this traversal requires $O(n)$ operations. The complexity of the other operations of this algorithm are straightforward and are 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 $v_j \geq |A_{ji}| + x^I_i$. 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 A and F, we have v and x does not have $\mathbb{0}$ entries. Thus $u = -x^I$ and v satisfy (7.9) $M = A_I^\top$. It follows from the discussion in Section 7.2.2 that the operations between Line 3 and 12 compute the tropical reduced costs. $\square$

Chapter 8

Tropicalizing the central path

In this chapter, we apply the tropicalization process to the central path in linear programming.

We consider linear programs defined on the Hardy field $\mathbb{K}$. Since $\mathbb{K}$ is real closed, the central path of a linear program on $\mathbb{K}$ is well-defined. The elements of $\mathbb{K}$ are real-valued functions. As a result, a linear program over $\mathbb{K}$ encodes a family of linear programs over $\mathbb{R}$, and the central path on $\mathbb{K}$ describes the central paths of this family. The tropical central path is then defined as the image under the valuation map. Thus, the tropical central path is a logarithmic limit of a family of classical central paths. We establish that this convergence is uniform on closed intervals.

The tropical central path has a purely geometric characterization. We show that the tropical analytic center is the greatest element of the tropicalization of the feasible set, the tropical equivalent of a barycenter. Thus, the tropical analytic center does not depend on the external representation of the feasible set. Similarly, any point on the tropical central path is the tropical barycenter of the tropical polyhedron obtained by intersecting the values of the feasible region with a tropical sublevel set induced by the objective function. This is in stark contrast with the classical case, where the central path depends on the halfspace description of the feasible set. In this way, Deza, Nematollahi, Peyghami and Terlaky [DNPT06] bent the central path of the Klee-Minty cube by adding redundant halfspaces in its representation, so that it visits a neighborhood of every vertex of the cube.

A maybe surprising feature is that the tropical central path can degenerate to a path taken by the tropical simplex method. We can even provide a quite general sufficient condition under which the tropical central path coincides with the image of a path of the classical simplex method under the valuation map. Consequently, the tropical central path may have the same worst-case behavior as the simplex method.

A main contribution of this chapter comes from studying the total curvature of the real central paths arising from lifting tropical linear programs to the Hardy field $\mathbb{K}$. The curvature measures how far a path differs from a straight line. Intuitively, a central path with high curvature should be harder to approximate with line segments, and thus this suggests more iterations of the interior point methods. We disprove the continuous analogue of the Hirsch conjecture proposed by Deza, Terlaky and Zinchencko by constructing a family of linear programs with $3r+4$ inequalities in dimension $2r+2$ where the central path has a total curvature in $\Omega(2^r)$. This family arises by lifting tropical linear programs introduced by Bezem, Nieuwenhuis and Rodríguez-Carbonell [BNRC08] to show that an algorithm of Butkovič and Zimmermann [BZ06] has exponential running time. The tropical central path shows a fractal-like pattern, which looks like a staircase shape with $\Omega(2^r)$ steps.

Most of the contents of this chapter are covered in [ABGJ14], but it includes an improvement of the curvature analysis of the counter-example from $\Omega(2^r/r)$ to $\Omega(2^r)$.

8.1 Description of the tropical central path

In this chapter, $\mathbf{LP}(A,b,c)$ will denote linear programs of the form:

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

where $A \in \mathbb{R}^{m\times n}$, $b \in \mathbb{R}^m$, and $c \in \mathbb{R}^n$. The dual linear program reads:

$$\begin{aligned} &\text{maximize} && -b^\top y \\ &\text{subject to} && -A^\top y + c \geq 0,\ y \geq 0,\ y \in \mathbb{R}^m\ . \end{aligned}$$

In the following, we shall assume that the polyhedron $\{x \in \mathbb{R}^n \mid Ax \leq b, x \geq 0\}$ is bounded with non-empty interior. Given a positive $\mu \in \mathbb{R}$, the *barrier problem* is

$$\begin{aligned} &\text{minimize} && \frac{c^\top x}{\mu} - \sum_{j=1}^n \log(x_j) - \sum_{i=1}^m \log(w_i) \\ &\text{subject to} && Ax + b = w,\ x > 0, w > 0. \end{aligned} \tag{8.1}$$

The objective function in (8.1) is continuous, strictly convex, and it tends to infinity when (x,w) tends to the boundary of the bounded non-empty convex set $\{(x,w) \in \mathbb{R}^{n+m} \mid Ax + w = b, x > 0, w > 0\}$. Hence, the problem (8.1) admits a unique optimum (x^μ, w^μ) in the latter set. By convexity, this optimum is characterized by the first-order optimality conditions:

$$\begin{aligned} Ax + b &= w \\ -A^\top y + c &= s \\ w_i y_i &= \mu \quad \text{for all } i \in [m] \\ x_j s_j &= \mu \quad \text{for all } j \in [n] \\ x, w, y, s &> 0\ . \end{aligned} \tag{8.2}$$

Thus, for any positive real number μ, there exists a unique solution $(x^\mu, w^\mu, y^\mu, s^\mu) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^m \times \mathbb{R}^n$ to the system of polynomial equations (8.2). The *central path* is the image of the map $\mathcal{C}_{A,b,c} : \mathbb{R}_{>0} \to \mathbb{R}^{2m+2n}$ which sends a positive real number μ to the vector $(x^\mu, w^\mu, y^\mu, s^\mu)$. The *primal central path* is the projection of the central path onto the (x, w)-coordinates. Similarly, the *dual central path* is gotten by projecting onto the (y, s)-coordinates.

8.1.1 Dequantization of a definable family of central paths

Let $\mathbb{K} = H(\bar{\mathbb{R}}^{\mathbb{R}})$ be the Hardy field of the o-minimal structure $\bar{\mathbb{R}}^{\mathbb{R}}$. We consider $\boldsymbol{A} \in \mathbb{K}^{m\times n}, \boldsymbol{b} \in \mathbb{K}^m$ and $\boldsymbol{c} \in \mathbb{K}^n$. Throughout, we will make the following assumption.

Assumption H. The set $\{\boldsymbol{x} \in \mathbb{K}^n \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0, \boldsymbol{x} \geq 0\}$ is bounded with non-empty interior.

Clearly, the latter set is closed. However, in $\mathbb{K}^n$ a closed and bounded set is not necessarily compact.

Under Assumption H, the central paths of the linear programs $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ over $\mathbb{R}$ are ultimately well-defined. For a fixed real number M let us define the map $\mathcal{C} : (M, +\infty) \times \mathbb{R} \to \mathbb{R}^{2m+2n}$ which sends $t \in (M, +\infty)$ and $\lambda \in \mathbb{R}$ to $\mathcal{C}(t, \lambda) = \mathcal{C}_{\boldsymbol{A}(t),\boldsymbol{b}(t),\boldsymbol{c}(t)}(t^\lambda)$. For any t large enough, the map $\lambda \mapsto \mathcal{C}(t, \lambda)$ is a parameterization of the central path of $\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$. Our goal is to investigate the logarithmic limit

$$\mathcal{C}^{\mathbb{T}} : \lambda \mapsto \lim_{t\to+\infty} \log_t \mathcal{C}(t, \lambda) \ ,$$

where $\log_t$ is applied component-wise. The map $\mathcal{C}^{\mathbb{T}}$ is called the *tropical central path* of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. We shall prove the following theorem.

Theorem 8.1. *The family of maps* $(\log_t \mathcal{C}(t, \cdot))_t$ *converges uniformly on any closed interval* $[a, b] \subseteq \mathbb{R}$ *to the tropical central path* $\mathcal{C}^{\mathbb{T}}$.

Consider the following linear program over the ordered field $\mathbb{K}$:

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0,\ \boldsymbol{x} \geq 0,\ \boldsymbol{x} \in \mathbb{K}^n \ . \end{array} \qquad \mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$$

The problem $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ encodes the family of linear programs $(\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t)))_t$.

The next lemma shows that the central path of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ is well-defined, and that it describes the family of central paths of $(\mathbf{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t)))_t$.

Lemma 8.2. *For any* $\lambda \in \mathbb{R}$, *the map* $t \mapsto \mathcal{C}(t, \lambda)$ *is definable in* $\bar{\mathbb{R}}^{\mathbb{R}}$. *Its components are given by the unique solution* $(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}}, \boldsymbol{y}^{\boldsymbol{\mu}}, \boldsymbol{s}^{\boldsymbol{\mu}}) \in \mathbb{K}^{2m+2n}$ *of the system of polynomial equations*

$$\begin{aligned} \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} &= \boldsymbol{w} \\ \boldsymbol{A}^\top \boldsymbol{y} - \boldsymbol{c} &= \boldsymbol{s} \\ \boldsymbol{w}_i \boldsymbol{y}_i &= \boldsymbol{\mu} \quad \textit{for all } i \in [m] \\ \boldsymbol{x}_j \boldsymbol{s}_j &= \boldsymbol{\mu} \quad \textit{for all } j \in [n] \\ \boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{s} &> 0 \ , \end{aligned} \tag{8.3}$$

where $\boldsymbol{\mu} = t^\lambda$.

Proof. For an ordered field K and integers m and n, consider the following statement:

> "For any $\boldsymbol{A} \in K^{m\times n}, \boldsymbol{b} \in K^m$ and $\boldsymbol{c} \in K^n$ which satisfy Assumption H and any positive $\boldsymbol{\mu} \in K$, there exists a unique solution $(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{s}') \in K^{2m+2n}$ to the system of polynomial equations (8.3)."

This is a first-order sentence, ϕ, which is true in the structure $\bar{\mathbb{R}}$, *i.e.*, for $K = \mathbb{R}$. As $\bar{\mathbb{R}}^{\mathbb{R}}$ is an expansion of $\bar{\mathbb{R}}$, we have $\bar{\mathbb{R}}^{\mathbb{R}} \models \phi$. Thus, by Proposition 2.7, the sentence ϕ is also true in the structure $\mathfrak{H}(\bar{\mathbb{R}}^{\mathbb{R}})$. This means that the induced statement holds in the field $K = \mathbb{K} = H(\bar{\mathbb{R}}^{\mathbb{R}})$. In particular, for any $\lambda \in \mathbb{R}$, it holds for $\boldsymbol{\mu} = t^\lambda \in \mathbb{K}$.

Let $(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}}, \boldsymbol{y}^{\boldsymbol{\mu}}, \boldsymbol{s}^{\boldsymbol{\mu}}) \in \mathbb{K}^{2m+2n}$ be the unique solution of (8.3) for $\boldsymbol{\mu} = t^\lambda$. Then, for all t large enough, $(\boldsymbol{x}^{\boldsymbol{\mu}}(t), \boldsymbol{w}^{\boldsymbol{\mu}}(t), \boldsymbol{y}^{\boldsymbol{\mu}}(t), \boldsymbol{s}^{\boldsymbol{\mu}}(t)) \in \mathbb{R}^{2m+2n}$ is a solution of (8.2) for $A = \boldsymbol{A}(t)$, $b = \boldsymbol{b}(t)$, $c = \boldsymbol{c}(t)$, $\mu = \boldsymbol{\mu}(t)$. Since (8.2) admits a unique solution, we conclude that $\mathcal{C}(t,\lambda) = (\boldsymbol{x}^{\boldsymbol{\mu}}(t), \boldsymbol{w}^{\boldsymbol{\mu}}(t), \boldsymbol{y}^{\boldsymbol{\mu}}(t), \boldsymbol{s}^{\boldsymbol{\mu}}(t))$ for all t large enough. □

Since $t \mapsto \mathcal{C}(t,\lambda)$ is definable in $\bar{\mathbb{R}}^{\mathbb{R}}$, its image under the (component-wise) valuation map is well-defined, which proves the point-wise convergence of the family $(\log_t \mathcal{C}(t,\cdot))_t$. Furthermore, for any $\lambda \in \mathbb{R}$ we have

$$\lim_{t\to+\infty} \log_t \mathcal{C}(t,\lambda) = \operatorname{val}(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}}, \boldsymbol{y}^{\boldsymbol{\mu}}, \boldsymbol{s}^{\boldsymbol{\mu}}) \ ,$$

where $\boldsymbol{\mu} = t^\lambda$, and $(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}}, \boldsymbol{y}^{\boldsymbol{\mu}}, \boldsymbol{s}^{\boldsymbol{\mu}})$ is the unique solution of (8.3).

For fixed t, let z_t be a component of the map $\lambda \mapsto \log_t \mathcal{C}(t,\lambda)$. To prove uniform convergence, we will use the fact that for all large enough t, the maps z_t are "almost" 1-Lipschitz.

Lemma 8.3. *For t large enough and any $\lambda, \lambda' \in \mathbb{R}$, we have:*

$$|z_t(\lambda) - z_t(\lambda')| \leq \log_t(2n+2m) + |\lambda - \lambda'| \ .$$

Proof. Let $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{s}) \in \mathbb{K}^{2m+2n}$ and $(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{s}') \in \mathbb{K}^{2m+2n}$ be two solutions of (8.3) obtained for two parameters $\boldsymbol{\mu} = t^\lambda$ and $\boldsymbol{\mu}' = t^{\lambda'}$. As in [VY96, Lemma 16], by combining the defining equations, we obtain:

$$\sum_{j=1}^{n} \boldsymbol{x}_j \boldsymbol{s}'_j + \sum_{j=1}^{n} \boldsymbol{x}'_j \boldsymbol{s}_j + \sum_{i=1}^{m} \boldsymbol{w}_i \boldsymbol{y}'_i + \sum_{i=1}^{m} \boldsymbol{w}'_i \boldsymbol{y}_i = (n+m)(t^\lambda + t^{\lambda'}) \tag{8.4}$$

Since the summands on the left-hand side of (8.4) are all positive, every summand is smaller than $(n+m)(t^\lambda + t^{\lambda'})$. In particular, for any $j \in [n]$, we have $\boldsymbol{x}_j \boldsymbol{s}'_j \leq (n+m)(t^\lambda + t^{\lambda'})$ and $\boldsymbol{x}'_j \boldsymbol{s}_j \leq (n+m)(t^\lambda + t^{\lambda'})$. Since $\boldsymbol{x}_j \boldsymbol{s}_j = t^\lambda$ and $\boldsymbol{x}'_j \boldsymbol{s}'_j = t^{\lambda'}$, we deduce that:

$$\begin{aligned} \boldsymbol{x}_j &\leq (n+m)(1 + t^{\lambda-\lambda'}) \boldsymbol{x}'_j \\ \boldsymbol{x}'_j &\leq (n+m)(1 + t^{\lambda'-\lambda}) \boldsymbol{x}_j \ . \end{aligned}$$

To prove the lemma, it is sufficient to consider $\lambda \geq \lambda'$. In this case, $t^{\lambda-\lambda'} \geq 1$, which implies:

$$\begin{aligned} \boldsymbol{x}_j &\leq 2(n+m)t^{\lambda-\lambda'}\boldsymbol{x}'_j \\ \boldsymbol{x}'_j &\leq 2(n+m)\boldsymbol{x}_j \ . \end{aligned}$$

Applying $\log_t$ to these inequalities yields the conclusion for the components $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$. The same proof readily applies to the other components. □

Proof of Theorem 8.1. Let z be the point-wise limit of the functions z_t as t approaches infinity. Consider any closed interval $[a, b] \subseteq \mathbb{R}$. Let $\varepsilon > 0$, and choose a partition $a = a_1 < a_2 < \cdots < a_k < a_{k+1} = b$ such that $a_{i+1} - a_i \leq \varepsilon$ for all $i \in [k]$. Now let $\lambda \in [a, b]$ and let i be the index such that $\lambda \in [a_i, a_{i+1}]$. Then,

$$|z_t(\lambda) - z(\lambda)| \leq |z_t(\lambda) - z_t(a_i)| + |z_t(a_i) - z(a_i)| + |z(a_i) - z(\lambda)| \ .$$

By Lemma 8.3, we have:

$$|z_t(\lambda) - z_t(a_i)| \leq \log_t(2n + 2m) + \lambda - a_i \leq \log_t(2n + 2m) + \varepsilon \ .$$

Thus, there exists a t_ε such that $|z_t(\lambda) - z_t(a_i)| \leq 2\varepsilon$ for all $t \geq t_\varepsilon$. Furthermore, Lemma 8.3 also shows that:

$$|z(\lambda) - z(a_i)| \leq \lambda - a_i \leq \varepsilon \ .$$

Finally, since the functions z_t converge pointwise to z, there exists a t'_ε such that $|z_t(a_i) - z(a_i)| \leq \varepsilon$ for all $t \geq t'_\varepsilon$ and all $i \in [k]$. We conclude that $(z_t)_t$ converges uniformly on $[a, b]$. □

8.1.2 Geometric description of the tropical central path

We now use barrier functions on the Hardy field $H(\bar{\mathbb{R}}^{\mathbb{R}})$ to characterize the central path. In order to obtain definable barrier functions, we use the structure $\bar{\mathbb{R}}_{\exp}$ which expands the ordered real field structure $\bar{\mathbb{R}}$ by adding the exponential function. The structure $\bar{\mathbb{R}}_{\exp}$ is o-minimal [vdDMM94]. Note that every power function is definable in $\bar{\mathbb{R}}_{\exp}$, thus the definable functions of $\bar{\mathbb{R}}^{\mathbb{R}}$ are also definable in $\bar{\mathbb{R}}_{\exp}$. As a consequence, the Hardy field $H(\bar{\mathbb{R}}_{\exp})$ contains $\mathbb{K} = H(\bar{\mathbb{R}}^{\mathbb{R}})$. The exponential is definable in the structure $\mathfrak{H}(\bar{\mathbb{R}}_{\exp})$ of the Hardy field $H(\bar{\mathbb{R}}_{\exp})$, and thus the logarithm is also definable in this structure. Hence, if $\boldsymbol{f} \in \mathbb{K}$ is positive, $\log(\boldsymbol{f})$ belongs to the ordered field $H(\bar{\mathbb{R}}_{\exp})$. Consequently, given $\boldsymbol{A} \in \mathbb{K}^{m\times n}, \boldsymbol{b} \in \mathbb{K}^m, \boldsymbol{c} \in \mathbb{K}$ and $\boldsymbol{\mu} \in \mathbb{K}, \boldsymbol{\mu} > 0$, the following optimization problem on $(\boldsymbol{x}, \boldsymbol{w}) \in \mathbb{K}^n \times \mathbb{K}^m$ is well-defined if the objective function is interpreted in $H(\bar{\mathbb{R}}_{\exp})$.

$$\begin{aligned} &\text{minimize} \quad && \frac{\boldsymbol{c}^\top \boldsymbol{x}}{\boldsymbol{\mu}} - \sum_{j=1}^{n} \log(\boldsymbol{x}_j) - \sum_{i=1}^{m} \log(\boldsymbol{w}_i) \\ &\text{subject to} \quad && \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} = \boldsymbol{w}, \ \boldsymbol{x} > 0, \ \boldsymbol{w} > 0 \ . \end{aligned} \tag{8.5}$$

Lemma 8.4. *Let $(\boldsymbol{x^\mu}, \boldsymbol{w^\mu}, \boldsymbol{y^\mu}, \boldsymbol{s^\mu})$ be the unique solution of* (8.3)*. The point $(\boldsymbol{x^\mu}, \boldsymbol{w^\mu})$ is the unique solution of* (8.5)*.*

Proof. Let $\mathfrak{R}$ be the expansion of the structure $\bar{\mathbb{R}}_{\exp}$ in which we added a symbol log. The latter is interpreted as the map $x \mapsto \log(x)$ for positive elements x and $x \mapsto 0$ for non-positive elements. The structure $\mathfrak{R}$ is still o-minimal, since the sets definable in $\mathfrak{R}$ and $\bar{\mathbb{R}}_{\exp}$ are the same. Given n, m, the following statement is a sentence in the language of $\mathfrak{R}$.

> "For any $\boldsymbol{A} \in K^{m\times n}, \boldsymbol{b} \in K^m$ and $\boldsymbol{c} \in K^n$ which satisfy Assumption H and any positive $\boldsymbol{\mu} \in K$, the optimization problem (8.5) has a unique solution. It is given by the point $(\boldsymbol{x}', \boldsymbol{w}')$, where $(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{s}')$ is the unique solution of (8.3)."

We already noted that this sentence is true when $K = \mathbb{R}$, *i.e.*, in the structure $\mathfrak{R}$. Since the latter is o-minimal, by Proposition 2.7, this sentence is also true in $\mathfrak{H}(\mathfrak{R})$, *i.e.*, when $K = H(\bar{\mathbb{R}}_{\exp})$. Now if $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}$ and $\boldsymbol{\mu}$ have entries in $\mathbb{K} \subseteq H(\bar{\mathbb{R}}_{\exp})$, the system (8.3) admits a unique solution with entries in $\mathbb{K}$ by Lemma 8.2. □

Let $\mathcal{P}$ be a non-empty bounded tropical polyhedron in $\mathbb{T}^n$. Then, there is a unique element in $\mathcal{P}$ which is the coordinate-wise maximum of all elements in $\mathcal{P}$. We call it the *tropical barycenter* of $\mathcal{P}$. Indeed, $\mathcal{P} = \operatorname{tconv}(V)$ for some finite set $V \subseteq \mathbb{T}^n$ by Theorem 4.11. Hence, $\mathcal{P}$ contains the point $\bigoplus_{v\in V} v$, which is greater than any other point in $\mathcal{P}$ with respect to the partial order of $\mathbb{T}^n$. In particular if $\boldsymbol{\mathcal{P}}$ is a non-empty bounded Hardy polyhedron included in the positive orthant, then $\operatorname{val}(\boldsymbol{\mathcal{P}})$ is a bounded tropical polyhedron. So $\operatorname{val}(\boldsymbol{\mathcal{P}})$ has a well-defined tropical barycenter.

Theorem 8.5. *Let $(\boldsymbol{x^\mu}, \boldsymbol{w^\mu})$ be the point on the primal central path of the Hardy linear program* $\operatorname{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ *at $\boldsymbol{\mu} \in \mathbb{K}$ with $\boldsymbol{\mu} > 0$, and let $\boldsymbol{\nu}$ be that LP's optimal value. Then* $\operatorname{val}(\boldsymbol{x^\mu}, \boldsymbol{w^\mu})$ *is the tropical barycenter of* $\operatorname{val}(\boldsymbol{\mathcal{P}^\mu})$ *where*

$$\boldsymbol{\mathcal{P}^\mu} := \{(\boldsymbol{x}, \boldsymbol{w}) \in \mathbb{K}^{n+m} \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} = \boldsymbol{w},\ \boldsymbol{c}\boldsymbol{x} \leq \boldsymbol{\nu} + (n+m)\boldsymbol{\mu},\ \boldsymbol{x} \geq 0,\ \boldsymbol{w} \geq 0\} .$$

Proof. Let $(\boldsymbol{x^\mu}, \boldsymbol{w^\mu}, \boldsymbol{y^\mu}, \boldsymbol{s^\mu})$ be a point on the central path. By (8.3), we have

$$\begin{aligned}
\boldsymbol{c}^\top \boldsymbol{x^\mu} &= (\boldsymbol{s^\mu})^\top \boldsymbol{x^\mu} + (\boldsymbol{y^\mu})^\top \boldsymbol{A}\boldsymbol{x^\mu} = (\boldsymbol{s^\mu})^\top \boldsymbol{x^\mu} + (\boldsymbol{y^\mu})^\top (\boldsymbol{w^\mu} - \boldsymbol{b}) \\
&= \sum_{j=1}^n \boldsymbol{s}_j^{\boldsymbol{\mu}} \boldsymbol{x}_j^{\boldsymbol{\mu}} + \sum_{i=1}^m \boldsymbol{y}_i^{\boldsymbol{\mu}} \boldsymbol{w}_i^{\boldsymbol{\mu}} - \boldsymbol{b}^\top \boldsymbol{y^\mu} = (n+m)\boldsymbol{\mu} - \boldsymbol{b}^\top \boldsymbol{y^\mu} .
\end{aligned}$$

Furthermore, $\boldsymbol{y^\mu}$ is a feasible solution of the dual linear program:

$$\begin{aligned}
&\text{maximize} && -\boldsymbol{b}^\top \boldsymbol{y} \\
&\text{subject to} && -\boldsymbol{A}^\top \boldsymbol{y} + \boldsymbol{c} \geq 0,\ \boldsymbol{y} \geq 0,\ \boldsymbol{y} \in \mathbb{K}^m .
\end{aligned}$$

By weak duality (Theorem 3.6), we have $-\boldsymbol{b}^\top \boldsymbol{y^\mu} \leq \boldsymbol{\nu}$. Consequently, $\boldsymbol{c}^\top \boldsymbol{x^\mu} \leq \boldsymbol{\nu} + (n+m)\boldsymbol{\mu}$.

Now by Lemma 8.4, $(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}})$ is the unique solution of the barrier problem (8.5). By the discussion above, we can add the constraint $\boldsymbol{c}^\top \boldsymbol{x} \leq \boldsymbol{\nu} + (n+m)\boldsymbol{\mu}$ to the problem (8.5) without changing its optimal solution. Moreover, adding the constant $-\boldsymbol{\nu}/\boldsymbol{\mu}$ to the objective function still does not change the solution of the problem. Thus $(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}})$ is the unique solution of

$$\begin{aligned} &\text{minimize} && \frac{\boldsymbol{c}^\top \boldsymbol{x} - \boldsymbol{\nu}}{\boldsymbol{\mu}} - \sum_{j=1}^{n} \log(\boldsymbol{x}_j) - \sum_{i=1}^{m} \log(\boldsymbol{w}_i) \\ &\text{subject to} && \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} = \boldsymbol{w},\ \boldsymbol{c}^\top \boldsymbol{x} \leq \boldsymbol{\nu} + (n+m)\boldsymbol{\mu},\ \boldsymbol{x} > 0,\ \boldsymbol{w} > 0\ . \end{aligned} \tag{8.6}$$

Let $\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}_{>0}$ be the feasible set of (8.6) and consider a feasible solution $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}_{>0}$. Since $\boldsymbol{c}^\top \boldsymbol{x} - \boldsymbol{\nu} \leq (n+m)\boldsymbol{\mu}$, the term $(\boldsymbol{c}^\top \boldsymbol{x} - \boldsymbol{\nu})/\boldsymbol{\mu}$ is the germ of a function which is asymptotically ct^α for some $\alpha, c \in \mathbb{R}$ with $\alpha \leq 0$. On the other hand, $\log(\boldsymbol{x}_j)$ is asymptotically $\operatorname{val}(\boldsymbol{x}_j)\log(t)$ for any $j \in [n]$. Since $t^\alpha = o(\log(t))$ when $\alpha \leq 0$, the objective value of (8.6) is asymptotically

$$-\left(\sum_{j=1}^{n} \operatorname{val}(\boldsymbol{x}_j) + \sum_{i=1}^{m} \operatorname{val}(\boldsymbol{w}_i)\right) \log(t)\ .$$

As a consequence, $\operatorname{val}(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}})$ is the supremum of $\sum_{j=1}^n x_j + \sum_{i=1}^n w_i$ as (x, w) ranges over the set $\operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}_{>0})$. Now, let (x^*, w^*) be the tropical barycenter of $\operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}})$. Then, $x^*_j \geq \operatorname{val}(\boldsymbol{x}^{\boldsymbol{\mu}}_j)$ and $w^*_i \geq \operatorname{val}(\boldsymbol{w}^{\boldsymbol{\mu}}_i)$ for all $i \in [m]$, $j \in [n]$. In particular, $x^*_j > -\infty$ and $w^*_i > -\infty$. It follows that $(x^*, w^*) \in \operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}_{>0})$, and:

$$\sum_{j=1}^{n} \operatorname{val}(\boldsymbol{x}^{\boldsymbol{\mu}}_j) + \sum_{i=1}^{m} \operatorname{val}(\boldsymbol{w}^{\boldsymbol{\mu}}_i) \geq \sum_{j=1}^{n} x^*_j + \sum_{i=1}^{m} w^*_i\ .$$

We conclude that $\operatorname{val}(\boldsymbol{x}^{\boldsymbol{\mu}}, \boldsymbol{w}^{\boldsymbol{\mu}}) = (x^*, w^*)$. □

The *analytic center* of the polyhedron

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

can be defined as the unique minimum point $(\boldsymbol{x}, \boldsymbol{w})$ of (8.6), when $\boldsymbol{c} = 0$. Then, the *tropical analytic center* is defined as the image of the analytic center by the valuation map. By specializing the characterization of the tropical central path to $\boldsymbol{c} = 0$, we get:

Corollary 8.6. *The tropical analytic center of the polyhedron $\boldsymbol{\mathcal{P}}$ coincides with the tropical barycenter of the image of this polyhedron by the valuation map.* □

Hence, even if the analytic center is an algebraic notion (it depends on the external representation of the set $\boldsymbol{\mathcal{P}}$), the tropical analytic center is, suprisingly, completely determined by the set $\boldsymbol{\mathcal{P}}$. We shall see that the whole tropical central path also has a purely geometric description. We begin with a case where the geometric description can be obtained explicitely from $\operatorname{val}(\boldsymbol{\mathcal{P}})$ and $\operatorname{val}(\boldsymbol{c})$.

Corollary 8.7. *Suppose that the optimal value of* $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *is* $\boldsymbol{\nu}=0$ *and that* $\boldsymbol{c}$ *has nonnegative entries. Then, the tropical central path at* $\lambda \in \mathbb{R}$ *is the tropical barycenter of the set*

$$\mathcal{P}^\lambda := \{(x,w) \in \operatorname{val}(\boldsymbol{\mathcal{P}}) \mid \max(x_1 + \operatorname{val}(\boldsymbol{c}_1), \dots, x_n + \operatorname{val}(\boldsymbol{c}_n)) \leq \lambda\} \ .$$

Proof. Let $\boldsymbol{\mu} = t^\lambda$. By Theorem 8.5, the tropical central path at λ is the tropical barycenter of $\operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}})$. Clearly, $\operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}) \subseteq \mathcal{P}^\lambda$. Thus, we only need to prove that the tropical barycenter (x^λ, w^λ) of $\mathcal{P}^\lambda$ admits a pre-image by the valuation map which belongs to $\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}$.

By definition, there exists $(\boldsymbol{x}^\lambda, \boldsymbol{w}^\lambda) \in \boldsymbol{\mathcal{P}}$ such that $\operatorname{val}(\boldsymbol{x}^\lambda, \boldsymbol{w}^\lambda) = (x^\lambda, w^\lambda)$. If $\boldsymbol{c}\boldsymbol{x}^\lambda = 0$, then $\boldsymbol{c}\boldsymbol{x}^\lambda \leq (n+m)\boldsymbol{\mu}$ and thus $(\boldsymbol{x}^\lambda, \boldsymbol{w}^\lambda) \in \boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}$. Otherwise, the germ $\boldsymbol{c}\boldsymbol{x}^\lambda$ is asymptotically αt^β for some $\alpha, \beta \in \mathbb{R}$ with $\alpha \neq 0$. Since $\boldsymbol{c}$ and $\boldsymbol{x}^\lambda$ has nonnegative entries, $\alpha > 0$ and we have

$$\beta = \operatorname{val}(\boldsymbol{c}\boldsymbol{x}^\lambda) = \max(x^\lambda_1 + \operatorname{val}(\boldsymbol{c}_1), \dots, x^\lambda_n + \operatorname{val}(\boldsymbol{c}_n)) \leq \lambda \ .$$

If $\alpha < n+m$, then clearly $\boldsymbol{c}\boldsymbol{x}^\lambda < (n+m)t^\lambda = (n+m)\boldsymbol{\mu}$ and thus $\boldsymbol{x}^\lambda \in \boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}$.

We now treat the case $\alpha \geq n+m$. Let $(\boldsymbol{x}^*, \boldsymbol{w}^*)$ be an optimal solution of $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$. Consider the point:

$$(\boldsymbol{x}, \boldsymbol{w}) = \frac{1}{\alpha}(\boldsymbol{x}^\lambda, \boldsymbol{w}^\lambda) + \left(1 - \frac{1}{\alpha}\right)(\boldsymbol{x}^*, \boldsymbol{w}^*) \ .$$

As $\alpha > 1$, we have $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}$ by convexity. Moreover, $\boldsymbol{c}\boldsymbol{x} = \frac{1}{\alpha}\boldsymbol{c}\boldsymbol{x}^\lambda$ since $\boldsymbol{c}\boldsymbol{x}^* = 0$ by assumption. Thus $\boldsymbol{c}\boldsymbol{x}$ is asymptotically t^β. Since $\beta \leq \lambda$, we obtain that $\boldsymbol{c}\boldsymbol{x} \leq (n+m)t^\lambda = (n+m)\boldsymbol{\mu}$, hence that $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}$. It remains to show that $\operatorname{val}(\boldsymbol{x}, \boldsymbol{w}) = (x^\lambda, w^\lambda)$. To this end, observe that $\operatorname{val}(\boldsymbol{x}, \boldsymbol{w}) \geq (x^\lambda, w^\lambda)$ since $(\boldsymbol{x}^\lambda, \boldsymbol{w}^\lambda)$ and $(\boldsymbol{x}^*, \boldsymbol{w}^*)$ both have nonnegative entries and $\alpha > 1$. Furthermore, $\operatorname{val}(\boldsymbol{x}, \boldsymbol{w}) \leq (x^\lambda, w^\lambda)$ as $\operatorname{val}(\boldsymbol{x}, \boldsymbol{w}) \in \operatorname{val}(\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}}) \subseteq \mathcal{P}^\lambda$. This concludes the proof. □

In the general case, the tropical central path still admits a geometric description, but this description involves an optimal solution of the dual of $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$.

Corollary 8.8. *There exists a pair* $(y^*, s^*) \in \mathbb{T}^m \times \mathbb{T}^n$ *such that the tropical central path at any* $\lambda \in \mathbb{R}$ *is given by the tropical barycenter of the set:*

$$\{(x,w) \in \operatorname{val}(\boldsymbol{\mathcal{P}}) \mid \max(x_1 + s^*_1, \dots, x_n + s^*_n, w_1 + y^*_1, \dots, w_m + y^*_m) \leq \lambda\} \ .$$

Proof. Let $(\boldsymbol{y}^*, \boldsymbol{s}^*)$ be an optimal dual solution and $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}$. Then, we have:

$$\boldsymbol{c}^\top \boldsymbol{x} = -\boldsymbol{b}^\top \boldsymbol{y}^* + (\boldsymbol{s}^*)^\top \boldsymbol{x} + (\boldsymbol{y}^*)^\top \boldsymbol{w} \ .$$

Furthermore, $-\boldsymbol{b}^\top \boldsymbol{y}^* = \boldsymbol{\nu}$ by strong duality. Thus,

$$\boldsymbol{\mathcal{P}}^{\boldsymbol{\mu}} = \{(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}} \mid (\boldsymbol{s}^*)^\top \boldsymbol{x} + (\boldsymbol{y}^*)^\top \boldsymbol{w} \leq (n+m)\boldsymbol{\mu}\} \ .$$

Since $(\boldsymbol{y}^*, \boldsymbol{s}^*) \geq 0$, applying the arguments of the proof of Corollary 8.7 provides the result. □

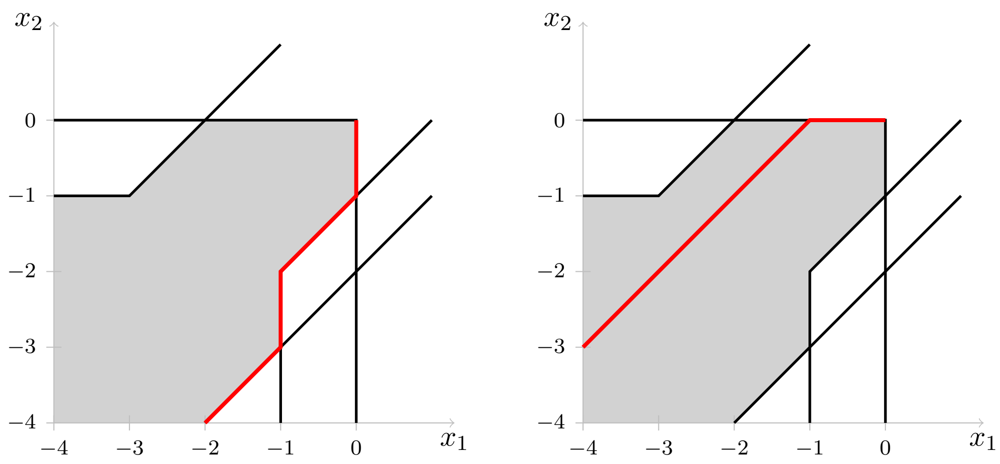

Figure 8.1: Tropical central paths on the Hardy polyhedron (8.7) for the objective function $\min \boldsymbol{x}_1$ (left) and $\min t\boldsymbol{x}_1 + \boldsymbol{x}_2$ (right).

Example 8.9. Consider the Hardy polyhedron of $\mathbb{K}^2$ defined by:

$$\begin{aligned} \boldsymbol{x}_1 + \boldsymbol{x}_2 &\leq 2 \\ t\boldsymbol{x}_1 &\leq 1 + t^2\boldsymbol{x}_2 \\ t\boldsymbol{x}_2 &\leq 1 + t^3\boldsymbol{x}_1 \\ \boldsymbol{x}_1 &\leq t^2\boldsymbol{x}_2 \\ \boldsymbol{x}_1, \boldsymbol{x}_2 &\geq 0 \ . \end{aligned} \tag{8.7}$$

Its value $\operatorname{val}(\boldsymbol{\mathcal{P}})$ is the tropical set described by the inequalities:

$$\begin{aligned} \max(x_1, x_2) &\leq 0 \\ 1 + x_1 &\leq \max(0, 2 + x_2) \\ 1 + x_2 &\leq \max(0, 3 + x_1) \\ x_1 &\leq 2 + x_2 \ . \end{aligned} \tag{8.8}$$

Tropical central paths on the polyhedron (8.7), for two objective functions, are depicted in Figure 8.1. The hyperplanes associated with the first four halfspaces in (8.7) induce an arrangement in the positive orthant $\mathbb{K}^2_+$. Figure 8.2 depicts the tropical central paths on the cells of this arrangement for the objective functions $\min t\boldsymbol{x}_1 + \boldsymbol{x}_2$ and $\max t\boldsymbol{x}_1 + \boldsymbol{x}_2$. Observe that the central paths trace the arrangement of tropical hyperplanes associated with the tropical halfspaces in (8.8), as well as the line $\{(-1+\gamma, \gamma) \mid \gamma \in \mathbb{R}\}$ associated with the objective function.

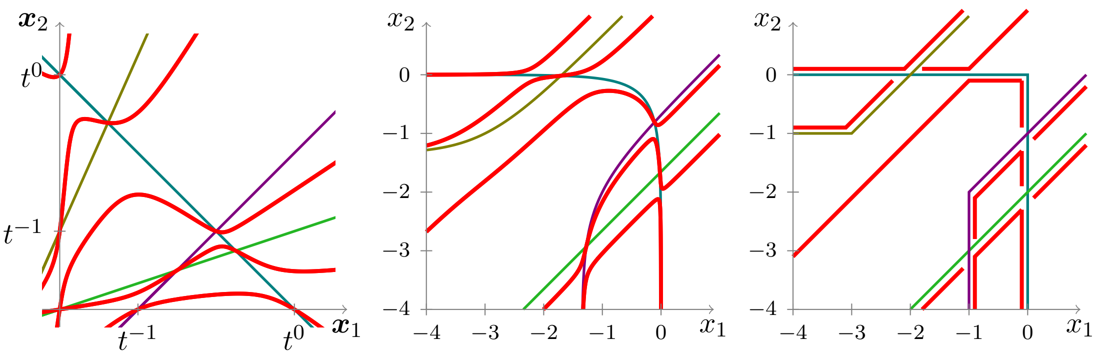

Figure 8.2: Illustration of the tropicalization of the central paths of the full-dimensional cells included in the positive orthant induced by the arrangement of hyperplanes associated with (8.7), for the objective function $\min t\boldsymbol{x}_1 + \boldsymbol{x}_2$ and $\max t\boldsymbol{x}_1 + \boldsymbol{x}_2$. Right: the real central paths for a real parameter t. Middle: the image of real central paths under the logarithmic map in base t. Right: the corresponding tropical central paths, where the parts of the paths that lie on the boundaries are slightly shifted inside their respective cell.

8.2 A tropical central path can degenerate to a tropical simplex path

In this section, we will restrict our attention to the x components of the tropical central path. To fix the notation, we consider a Hardy linear program $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, and the polyhedron $\boldsymbol{\mathcal{P}} = \{(\boldsymbol{x}, \boldsymbol{w}) \in \mathbb{K}^n \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{w} = \boldsymbol{b},\ \boldsymbol{x} \geq 0,\ \boldsymbol{w} \geq 0\}$. From this viewpoint, the tropical central path may visit the boundary of the (projection on the x-space) of the set $\operatorname{val}(\boldsymbol{\mathcal{P}})$. We will show that under some assumptions, the tropical central path lies on the image by the valuation map of the graph of the polyhedron $\boldsymbol{\mathcal{P}}$.

When the signed valuation of $(\boldsymbol{A}\ \boldsymbol{b})$ is sign-generic for the minor polynomials, we have a purely tropical description of the set $\operatorname{val}(\boldsymbol{\mathcal{P}})$ by Theorem 4.22. Furthermore, that result also shows that the images of the faces of $\boldsymbol{\mathcal{P}}$ under the valuation map also have a tropical description. In particular, this holds for the basic points and the edges of $\boldsymbol{\mathcal{P}}$, see Section 4.3. Using the notation of Theorem 4.22, $x \in \operatorname{val}(\boldsymbol{\mathcal{P}})$ is the value of a basic point (hence of a vertex by Proposition 3.14) if and only if it satisfies a system of n equalities $A_I^+ \odot x \oplus b_I^+ = A_I^- \odot x \oplus b_I^-$ where $\operatorname{tdet}(A_I) \neq \mathbb{0}$.

Proposition 8.10. *Consider a Hardy polyhedron, $\boldsymbol{\mathcal{P}} = \{\boldsymbol{x} \in \mathbb{K}^n \mid \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b} \geq 0\}$, contained in the positive orthant such that* $\operatorname{sval}(\boldsymbol{A}\ \boldsymbol{b})$ *is sign-generic for the minor polynomials, and $\boldsymbol{A}^- = (\min(\boldsymbol{A}_{ij}, 0))$ has at most one non-zero coefficient in each row. Then the tropical analytic center of $\boldsymbol{\mathcal{P}}$ coincides with the value of a vertex of $\boldsymbol{\mathcal{P}}$.*

Proof. By Theorem 4.22, the tropical polyhedron $\operatorname{val}(\boldsymbol{\mathcal{P}})$ is described by $\{x \in \mathbb{T}^n \mid A^+ \odot x \oplus b^+ \geq A^- \odot x \oplus b^-\}$. Since $\boldsymbol{A}^-$ has at most one non-zero coefficient in each row, for every $i \in [m]$ the tropical inequality $A_i^+ \odot x \oplus b_i^+ \geq A_i^- \odot x \oplus b_i^-$ is of the form;

$$\max(A_{i1}^+ + x_1, \dots, A_{in}^+ + x_n, b_i^+) \geq \max(A_{ik}^- + x_k, b_i^-) \ ,$$

for some $k \in [n]$.

Let x^* be the tropical analytic center of $\boldsymbol{\mathcal{P}}$. By Corollary 8.6, x^* is the tropical barycenter of $\operatorname{val}(\boldsymbol{\mathcal{P}})$. By Assumption H, x_j^* finite for each $j \in [n]$. Thus, there must exist an $i \in [m]$ such that $A_{ij}^- \neq -\infty$ and

$$\max(A_{i1}^+ + x_1^*, \dots, A_{in}^+ + x_n^*, b_i^+) = A_{ij}^- + x_j^* \ . \tag{8.9}$$

Consequently, x^* satisfies a set I of n equalities, one for each coordinate $j \in [n]$. By construction we have $\operatorname{tdet}(A_I^-) \neq \mathbb{0}$, thus $\operatorname{tdet}(A_I) \neq \mathbb{0}$. Consequently, x^* is the value of a vertex by Theorem 4.22. □

Vertices of $\boldsymbol{\mathcal{P}}$ are connected by edges, which are sets of the form $\{\boldsymbol{x} \in \boldsymbol{\mathcal{P}} \mid \boldsymbol{A}_K \boldsymbol{x} + \boldsymbol{b}_K = 0\}$ where $K \subseteq [m]$ is of cardinality $n-1$ and $\boldsymbol{A}_K$ is of rank $n-1$. Under the conditions of Theorem 4.22, the image of the edges under the valuation map are exactly the sets described by $\{x \in \operatorname{val}(\boldsymbol{\mathcal{P}}) \mid A_K^+ \odot x \oplus b_K^+ = A_K^- \odot x \oplus b_K^-\}$ where $K \subseteq [m]$ is of cardinality $n-1$ and A_K has a maximal square submatrix with non $\mathbb{0}$ tropical determinant.

Proposition 8.11. *Let $\boldsymbol{\mathcal{P}}$ be a Hardy polyhedron which satisfies the conditions of Proposition 8.10. Consider a linear program of the form:*

$$\min \boldsymbol{x}_k \ \text{s.t.}\ \boldsymbol{x} \in \boldsymbol{\mathcal{P}} \ , \tag{\textbf{LP}}$$

for some $k \in [n]$. If the optimal value of **LP** *is $\boldsymbol{\nu} = 0$, then the tropical central path of* **LP** *is contained in the image by the valuation map of the graph of $\boldsymbol{\mathcal{P}}$.*

Proof. By Corollary 8.7, the point x^λ on the tropical central path at $\lambda \in \mathbb{R}$ is the tropical barycenter of the tropical polyhedron $\{x \in \operatorname{val}(\boldsymbol{\mathcal{P}}) \mid x_k \leq \lambda\}$. As in the proof of Proposition 8.10, for each $j \in [n] \setminus \{k\}$ the point x^λ must satisfy an equality of the form (8.9). Thus, x^λ satisfies a set K of $n-1$ equalities and it is straightforward to check that the minor of A_K^- formed with the columns indexed by $[n] \setminus \{k\}$ has a finite tropical determinant. □

The latter proposition is illustrated in Figure 8.1 (left).

8.3 Central paths with high curvature

Bezem, Nieuwenhuis and Rodríguez-Carbonell [BNRC08] constructed a class of tropical linear programs for which an algorithm of Butkovič and Zimmermann [BZ06] exhibits an exponential running time. We lift each of these tropical linear programs to the Hardy field $\mathbb{K} = H(\bar{\mathbb{R}}^{\mathbb{R}})$ which then gives rise to a one-parameter family of ordinary linear programs over the reals. The latter are interesting as their central paths have an unusually high total curvature.

Let r be any positive integer. We define a linear program, $\mathbf{LP}_r$, over the Hardy field $\mathbb{K}$ in the $2r+2$ variables $\boldsymbol{u}_0, \boldsymbol{v}_0, \boldsymbol{u}_1, \boldsymbol{v}_1, \ldots, \boldsymbol{u}_r, \boldsymbol{v}_r$ as follows.

$$\begin{array}{lll} \min & \boldsymbol{v}_0 & \\ \text{s.t.} & \boldsymbol{u}_0 \leq t & \\ & \boldsymbol{v}_0 \leq t^2 & \\ & \boldsymbol{v}_i \leq t^{1-\frac{1}{2^i}}(\boldsymbol{u}_{i-1} + \boldsymbol{v}_{i-1}) & \text{for } 1 \leq i \leq r \\ & \boldsymbol{u}_i \leq t\boldsymbol{u}_{i-1} & \text{for } 1 \leq i \leq r \\ & \boldsymbol{u}_i \leq t\boldsymbol{v}_{i-1} & \text{for } 1 \leq i \leq r \\ & \boldsymbol{u}_r \geq 0, \ \boldsymbol{v}_r \geq 0 & \end{array} \qquad \mathbf{LP}_r$$

Clearly, the optimal value of $\mathbf{LP}_r$ is $\boldsymbol{\nu} = 0$, and an optimal solution is $\boldsymbol{u} = \boldsymbol{v} = 0$. It is straightforward to verify that the feasible set is bounded with a non-empty interior. Moreover, the feasible set is contained in the positive orthant and the $3r+4$ inequalities listed define facets. In particular, the remaining non-negativity constraints $\boldsymbol{u}_i \geq 0$ and $\boldsymbol{v}_i \geq 0$ for $0 \leq i < r$ are satisfied but redundant. We will denote the feasible region of $\mathbf{LP}_r$ as $\boldsymbol{\mathcal{P}}_r$.

Replacing t in $\mathbf{LP}_r$ by any positive real number gives rise to an ordinary linear program. For t sufficiently large the polytope of feasible points is combinatorially equivalent to the polytope of feasible points of the Hardy linear program. Figure 8.3 shows an example for $r = 1$ and $t \geq 2$, which *is* sufficiently large in this case.

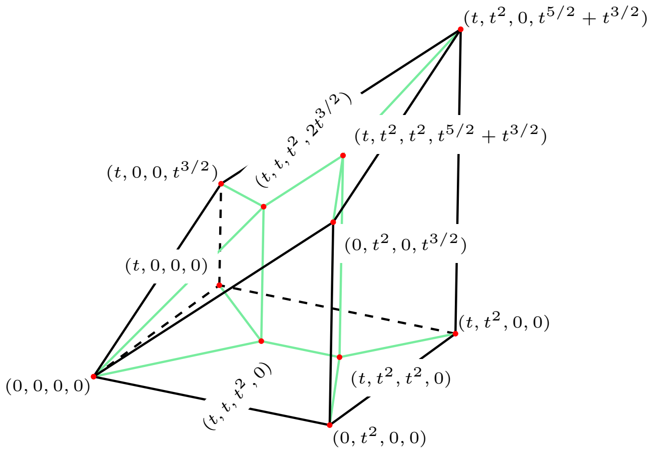

Figure 8.3: Schlegel diagram for $r = 1$ (and $t \geq 2$), projected onto the facet $u_1 = 0$; the points are written in (u_0, v_0, u_1, v_1)-coordinates

8.3.1 Tropical central path

We next compute the tropical central path arising from the linear program $\mathbf{LP}_r$ over the Hardy field $\mathbb{K}$. To this end, we introduce slack variables:

$$
\begin{array}{lll}
\min & \boldsymbol{v}_0 & \\
\text{s.t.} & \boldsymbol{u}_0 + \boldsymbol{z}_0 = t & \\
& \boldsymbol{v}_0 + \boldsymbol{h}_0 = t^2 & \\
& \boldsymbol{v}_i + \boldsymbol{h}_i = t^{1-\frac{1}{2^i}}(\boldsymbol{u}_{i-1} + \boldsymbol{v}_{i-1}) & \text{for } 1 \le i \le r \\
& \boldsymbol{u}_i + \boldsymbol{z}_i = t\boldsymbol{u}_{i-1} & \text{for } 1 \le i \le r \\
& \boldsymbol{u}_i + \boldsymbol{z}'_i = t\boldsymbol{v}_{i-1} & \text{for } 1 \le i \le r \\
& \boldsymbol{z}_0 \ge 0,\ \boldsymbol{h}_0 \ge 0 & \\
& \boldsymbol{u}_i \ge 0,\ \boldsymbol{v}_i \ge 0,\ \boldsymbol{z}_i \ge 0,\ \boldsymbol{z}'_i \ge 0,\ \boldsymbol{h}_i \ge 0 & \text{for } 1 \le i \le r\ .
\end{array}
\qquad \mathbf{LP}'_r
$$

For each positive parameter $\boldsymbol{\mu} \in \mathbb{K}$, we denote by $(\boldsymbol{u}^{\boldsymbol{\mu}}, \boldsymbol{v}^{\boldsymbol{\mu}}, \boldsymbol{z}^{\boldsymbol{\mu}}, (\boldsymbol{z}')^{\boldsymbol{\mu}}, \boldsymbol{h}^{\boldsymbol{\mu}})$ the point of the primal central path with parameter $\boldsymbol{\mu}$. Recall that the tropical central path $\mathcal{C}^{\mathbb{T}}$ is such that $\mathcal{C}^{\mathbb{T}}(\lambda)$ is the image by the valuation of $(\boldsymbol{u}^{\boldsymbol{\mu}}, \boldsymbol{v}^{\boldsymbol{\mu}}, \boldsymbol{z}^{\boldsymbol{\mu}}, (\boldsymbol{z}')^{\boldsymbol{\mu}}, \boldsymbol{h}^{\boldsymbol{\mu}})$ for $\boldsymbol{\mu} = t^{\lambda}$. The valuation of every point of the feasible set $\boldsymbol{\mathcal{P}}'_r$ of the program $\mathbf{LP}'_r$ satisfies the following equalites:

$$
\begin{array}{ll}
\max(u_0, z_0) = 1 & \\
\max(v_0, h_0) = 2 & \\
\max(v_i, h_i) = 1 - \dfrac{1}{2^i} + \max(u_{i-1}, v_{i-1}) & \text{for } 1 \le i \le r \\
\max(u_i, z_i) = 1 + u_{i-1} & \text{for } 1 \le i \le r \\
\max(u_i, z'_i) = 1 + v_{i-1} & \text{for } 1 \le i \le r \\
u_i \in \mathbb{T},\ v_i \in \mathbb{T},\ z_i \in \mathbb{T},\ h_i \in \mathbb{T} & \text{for } 0 \le i \le r \\
z'_i \in \mathbb{T} & \text{for } 1 \le i \le r.
\end{array}
\tag{8.10}
$$

Proposition 8.12. *For all $\lambda \in \mathbb{R}$, the tropical central path at λ, coincides with the maximal point $(u(\lambda), v(\lambda), z(\lambda), z'(\lambda), h(\lambda))$ satisfying the constraints* (8.10) *and $v_0 \le \lambda$. It is determined by:*

$$
\begin{array}{ll}
u_0 = z_0 = 1 & \\
h_0 = 2 & \\
v_0 = \min(2, \lambda) & \\
v_i = h_i = 1 - \dfrac{1}{2^i} + \max(u_{i-1}, v_{i-1}) & \textit{for } 1 \le i \le r \\
u_i = 1 + \min(u_{i-1}, v_{i-1}) & \textit{for } 1 \le i \le r
\end{array}
\tag{8.11}
$$

Proof. By Corollary 8.7, $\mathcal{C}^{\mathbb{T}}(\lambda)$ is the maximal point of the intersection of $\mathrm{val}(\boldsymbol{\mathcal{P}}'_r)$ with the tropical half-space

$$
\mathcal{H}^{\lambda} := \{(u, v, z, z', h) \in (\mathbb{T}^{r+1})^3 \times \mathbb{T}^r \times \mathbb{T}^{r+1} \mid v_0 \le \lambda\}\ .
$$

Using the homomorphism property of the valuation map, every point of $\mathcal{H}^\lambda \cap \operatorname{val}(\boldsymbol{\mathcal{P}}'_r)$ satisfies $v_0 \leq \lambda$ as well as (8.10).

It is straightforward to verify that (8.11) defines the maximal vector satisfying $v_0 \leq \lambda$ and (8.10). Therefore, $\mathcal{C}^{\mathbb{T}}(\lambda) \leq (u(\lambda), v(\lambda), z(\lambda), z'(\lambda), h(\lambda))$. To show that the opposite inequality holds, using Corollary 8.7 again, it suffices to lift $(u(\lambda), v(\lambda), z(\lambda), z'(\lambda), h(\lambda))$ to an element of $\boldsymbol{\mathcal{P}}'_r$. Such a lift can be obtained as the unique solution of the following system:

$$
\begin{aligned}
\boldsymbol{u}_0 = \boldsymbol{z}_0 &= \frac{1}{2}t \\
\boldsymbol{v}_0 &= \frac{1}{2}\min(t^2, t^\lambda) \\
\boldsymbol{h}_0 &= t^2 - \boldsymbol{v}_0 \\
\boldsymbol{v}_i = \boldsymbol{h}_i &= \frac{1}{2}t^{1-\frac{1}{2^i}}(\boldsymbol{u}_{i-1} + \boldsymbol{v}_{i-1}) && \text{for } 1 \leq i \leq r \\
\boldsymbol{u}_i &= \frac{1}{2}\min(t\boldsymbol{u}_{i-1}, t\boldsymbol{v}_{i-1}) && \text{for } 1 \leq i \leq r \\
\boldsymbol{z}_i &= t\boldsymbol{u}_{i-1} - \boldsymbol{u}_i && \text{for } 1 \leq i \leq r \\
\boldsymbol{z}'_i &= t\boldsymbol{v}_{i-1} - \boldsymbol{u}_i && \text{for } 1 \leq i \leq r \square
\end{aligned}
$$

It follows from Proposition 8.12 that $(u_i(\lambda), v_i(\lambda))_{0 \leq i \leq r}$ completely determine the other components of the tropical central path. Observe that $v_i(\lambda)$ is equal to the maximum of $u_{i-1}(\lambda)$ and $v_{i-1}(\lambda)$ translated by $1 - \frac{1}{2^i}$, while $u_i(\lambda)$ follows the minimum of these two variables shifted by 1; see Figure 8.4. Since the translation offsets differ by $\frac{1}{2^i}$, the components u_i and v_i cross each other $\Omega(2^i)$ times. More precisely, our next result shows that the curve $(u_i(\lambda), v_i(\lambda))$ has the shape of a staircase with $\Omega(2^i)$ steps.

Proposition 8.13. *Let $i \in [r]$ and $k \in \{0, \dots, 2^{i-1} - 1\}$. Then, for all λ in the interval $[\frac{4k}{2^i}, \frac{4k+2}{2^i}]$, we have*

$$u_i(\lambda) = i + \lambda - \frac{2k}{2^i} \text{ and } v_i(\lambda) = i + \frac{2k+1}{2^i} \,,$$

while for all $\lambda \in [\frac{4k+2}{2^i}, \frac{4k+4}{2^i}]$ we have

$$u_i(\lambda) = i + \frac{2k+2}{2^i} \text{ and } v_i(\lambda) = i + \lambda - \frac{2k+1}{2^i} \,.$$

Proof. We proceed by a bounded induction on $i \in [r]$. Starting with $i = 1$ and $k = 0$ we consider the tropical central path point at any $\lambda \in [0, 2]$. Our goal is to determine the tropical analytic center. It follows from (8.11) that

$$u_1 = 1 + \min(1, \lambda), \qquad v_1 = \frac{1}{2} + \max(1, \min(2, \lambda)) \,.$$

Thus for $\lambda \in [0, 1]$, $u_1 = 1 + \lambda$ and $v_1 = 1 + \frac{1}{2}$. For $\lambda \in [1, 2]$ we have $u_1 = 1 + \frac{2}{2}$ and $v_1 = 1 + \lambda - \frac{1}{2}$. Consequently, the claim holds for $i = 1$.

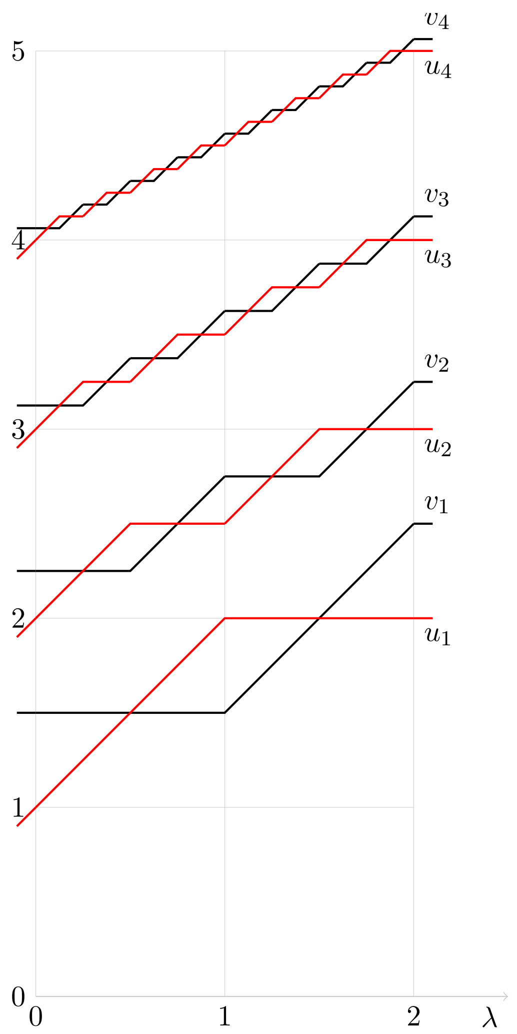

Figure 8.4: Evolution of the components of the tropical central path of $\mathbf{LP}_4$ with λ.

By induction, suppose the result is verified for $i < r$. We will show that it is also true for $i+1$. Consider any integer, k, in $\{0, \dots, 2^i - 1\}$. If k is even, let $k' = k/2$. Then, for all λ in the interval $[\frac{4k}{2^{i+1}}, \frac{4k+4}{2^{i+1}}] = [\frac{4k'}{2^i}, \frac{4k'+2}{2^i}]$, we have by induction:

$$u_i = i + \lambda - \frac{2k'}{2^i} = i + \lambda - \frac{k}{2^i} \quad \text{and} \quad v_i = i + \frac{2k'+1}{2^i} = i + \frac{k+1}{2^i} \ .$$

Thus,

$$u_{i+1} = i+1+\min\left(\frac{k+1}{2^i}, \lambda - \frac{k}{2^i}\right) \quad \text{and} \quad v_{i+1} = i+1+\max\left(\frac{k+1}{2^i}, \lambda - \frac{k}{2^i}\right) - \frac{1}{2^{i+1}} \ .$$

Separating the cases $\lambda \leq \frac{4k+2}{2^{i+1}}$ and $\lambda \geq \frac{4k+2}{2^{i+1}}$ leads to the desired conclusion.

If k is odd, $k = 2k' + 1$, then for any $\lambda \in [\frac{4k}{2^{i+1}}, \frac{4k+4}{2^{i+1}}] = [\frac{4k'+2}{2^i}, \frac{4k'+4}{2^i}]$ we have:

$$\begin{aligned} u_i &= i + \frac{2k'+2}{2^i} = i + \frac{k+2}{2^i} \\ v_i &= i + \lambda - \frac{2k'+1}{2^i} = i + \lambda - \frac{k+1}{2^i} \ . \end{aligned}$$

Thus,

$$\begin{aligned} u_{i+1} &= i + 1 + \min\left(\lambda - \frac{k+1}{2^i}, \frac{k+2}{2^i}\right) \\ v_{i+1} &= i + 1 + \max\left(\lambda - \frac{k+1}{2^i}, \frac{k+2}{2^i}\right) - \frac{1}{2^{i+1}} \ . \end{aligned}$$

As above, by separating the cases $\lambda \leq \frac{4k+2}{2^{i+1}}$ and $\lambda \geq \frac{4k+2}{2^{i+1}}$ we conclude that the inductive claim holds for $i + 1$. □

Remark 8.14. A similar induction shows that for $\lambda \geq 2$ the tropical central path is at the tropical analytic center, defined by $u_0 = 1, v_0 = 2$ and

$$u_i = i + 1 \quad \text{and} \quad v_i = i + 1 + \frac{1}{2^i} \qquad \text{for all } 1 \leq i \leq r \ .$$

For $\lambda \leq 0$, the tropical central path is a tropical half-line towards an optimum. We have $u_0(\lambda) = 1, v_0(\lambda) = \lambda$ as well as

$$u_i(\lambda) = i + \lambda \quad \text{and} \quad v_i(\lambda) = i + \frac{1}{2^i} \qquad \text{for all } 1 \leq i \leq r \ .$$

We will now show that the tropical central path of $\mathbf{LP}_r$ coincides with the image of a path of the simplex method under the valuation map. Our proof is elementary and independent of Proposition 8.11.

Proposition 8.15. *Under projection on the (u, v)-components, the tropical central path of $\mathbf{LP}_r$ is contained in the image of the vertex-edge graph of $\boldsymbol{\mathcal{P}}_r$ under the valuation map. The tropical central path at $\lambda \in \mathbb{R}$ is the value of a vertex if and only if $\lambda \geq 2$ or $\lambda = \frac{2k}{2^r}$ for some $k \in \{1, \dots, 2^r\}$.*

Proof. We prove the claim by induction on r. Suppose that $r = 1$. This situation in four dimensions is depicted in Figure 8.3. For $\lambda \geq 2$, the tropical central path is at the tropical analytic center of $\mathbf{LP}_1$:

$$u_0 = 1, \ v_0 = 2, \ u_1 = 2, \ v_1 = 5/2 \ .$$

This is the value of the vertex $(t, t^2, t^2, t^{5/2} + t^{3/2})$ of the Hardy polyhedron $\boldsymbol{\mathcal{P}}_1$ which is uniquely defined by the conditions

$$\boldsymbol{u}_0 = t, \boldsymbol{v}_0 = t^2, \boldsymbol{u}_1 = t\boldsymbol{u}_0, \boldsymbol{v}_1 = t^{1/2}(\boldsymbol{u}_0 + \boldsymbol{v}_0) \ . \tag{8.12}$$

For $\lambda = 1$ the tropical central path is at the point with coordinates

$$u_0 = 1,\ v_0 = 1,\ u_1 = 2,\ v_1 = 3/2\ ,$$

which corresponds to the vertex $(t, t, t^2, 2t^{3/2})$ of $\boldsymbol{\mathcal{P}}_1$, the unique solution of:

$$\boldsymbol{u}_0 = t,\ \boldsymbol{v}_0 = t,\ \boldsymbol{u}_1 = t,\ \boldsymbol{v}_0,\ \boldsymbol{v}_1 = t^{1/2}(\boldsymbol{u}_0 + \boldsymbol{v}_0)\ . \tag{8.13}$$

It is straighforward to check that the tropical central path for $\lambda \in [1, 2]$ is the image by the valuation map of the edge between the vertices (8.12) and (8.13). Similarly, for $\lambda \in]-\infty, 1]$, the tropical central path:

$$u_0 = 1,\ v_0 = \lambda,\ u_1 = 1 + \lambda,\ v_1 = 3/2\ .$$

is the value of the edge between the vertices (8.13) and $(t, 0, 0, t^{3/2})$ of $\boldsymbol{\mathcal{P}}_1$ defined by

$$\boldsymbol{u}_0 = t,\ \boldsymbol{v}_0 = 0,\ \boldsymbol{u}_1 = t\boldsymbol{v}_0,\ \boldsymbol{v}_1 = t^{1/2}(\boldsymbol{u}_0 + \boldsymbol{v}_0)\ . \tag{8.14}$$

Now suppose that the claim holds for $r \geq 1$. For $\lambda \geq 2$, the tropical central path of $\mathbf{LP}_{r+1}$ is at the analytic center $(u_0, v_0, \dots, u_r, v_r, u_{r+1}, v_{r+1})$. By Proposition 8.13, we have $v_{r+1} = 1 - 1/2^{r+1} + \max(u_r, v_r)$ and $u_{r+1} = 1 + u_r$. By induction, $(u_0, v_0, \dots, u_r, v_r)$ is the value of the vertex $(\boldsymbol{u}_0, \boldsymbol{v}_0, \dots, \boldsymbol{u}_r, \boldsymbol{v}_r)$ of $\mathbf{LP}_r$. The system defining this vertex of $\boldsymbol{\mathcal{P}}_r$, along with the equalities $\boldsymbol{v}_{r+1} = t^{1-1/2^{r+1}}(\boldsymbol{u}_r + \boldsymbol{v}_r)$ and $\boldsymbol{u}_{r+1} = t\boldsymbol{u}_r$ clearly have a unique solution which is feasible for $\mathbf{LP}_{r+1}$. Thus it defines a vertex of $\boldsymbol{\mathcal{P}}_{r+1}$. It is straightforward to verify that the valuation map applied to this vertex yields the tropical analytic center. Similarly, the argument above shows that the tropical central path of $\mathbf{LP}_{r+1}$ is the value of a vertex when $\lambda = \frac{2k}{2^r} = \frac{4k}{2^{r+1}}$ for some $k \in \{1, \dots, 2^r\}$.

Fix a $k \in \{1, \dots, 2^r - 1\}$. Then central path of $\mathbf{LP}_r$ at $\lambda \in [\frac{2k}{2^r}, \frac{2k+2}{2^r}] = [\frac{4k}{2^{r+1}}, \frac{4k+4}{2^{r+1}}]$ is the value of a point on an edge of $\boldsymbol{\mathcal{P}}_r$. This edge in $\mathbb{K}^{2r+2}$ defines a 3-dimensional face $\boldsymbol{\mathcal{F}}$ of $\boldsymbol{\mathcal{P}}_{r+1}$ in $\mathbb{K}^{2r+4}$. The intersection of $\boldsymbol{\mathcal{F}}$ with the three hyperplanes

$$\boldsymbol{v}_{r+1} = t^{1-1/2^{r+1}}(\boldsymbol{u}_r + \boldsymbol{v}_r),\ \boldsymbol{u}_{r+1} = t\boldsymbol{u}_r \quad \text{and} \quad \boldsymbol{u}_{r+1} = t\boldsymbol{u}_r \tag{8.15}$$

yields a vertex of $\boldsymbol{\mathcal{P}}_{r+1}$, and it can be checked that the value of this vertex is on the tropical central path of $\mathbf{LP}_{r+1}$ for $\lambda = \frac{4k+2}{2^{r+1}}$. It follows that the tropical central path of $\mathbf{LP}_{r+1}$ at $\lambda \in [\frac{4k}{2^{r+1}}, \frac{4k+2}{2^{r+1}}]$ and $\lambda \in [\frac{4k+2}{2^{r+1}}, \frac{4k+4}{2^{r+1}}]$ corresponds to points on two distinct edges of $\boldsymbol{\mathcal{P}}_{r+1}$. These two edges are obtained by intersecting $\boldsymbol{\mathcal{F}}$ with $\boldsymbol{v}_{r+1} = t^{1-1/2^{r+1}}(\boldsymbol{u}_r + \boldsymbol{v}_r)$ and either $\boldsymbol{u}_{r+1} = t\boldsymbol{u}_r$ or $\boldsymbol{u}_{r+1} = t\boldsymbol{v}_r$. It remains to consider $\lambda \leq \frac{4}{2^{r+1}}$. By induction, the tropical central path of $\mathbf{LP}_r$ for $\lambda \leq \frac{4}{2^{r+1}} = \frac{2}{2^r}$ is the set of values of an edge of $\boldsymbol{\mathcal{P}}_r$. As above, this edge yields a 3-face $\boldsymbol{\mathcal{F}}$ of $\boldsymbol{\mathcal{P}}_{r+1}$. Intersecting $\boldsymbol{\mathcal{F}}$ with the three hyperplanes (8.15) yields a vertex whose value is the tropical central path of $\mathbf{LP}_{r+1}$ at $\lambda = \frac{2}{2^{r+1}}$. Intersecting $\boldsymbol{\mathcal{F}}$ with $\boldsymbol{v}_{r+1} = t^{1-1/2^{r+1}}(\boldsymbol{u}_r + \boldsymbol{v}_r)$ and $\boldsymbol{u}_{r+1} = t\boldsymbol{v}_r$ yields an edge of $\boldsymbol{\mathcal{P}}_{r+1}$ whose set of values is the tropical central path at $\lambda \in [\frac{2}{2^{r+1}}, \frac{4}{2^{r+1}}]$. For $\lambda \leq \frac{2}{2^{r+1}}$, the tropical central path is the set of values of the edge obtained as the intersection of $\boldsymbol{\mathcal{F}}$ with $\boldsymbol{v}_{r+1} = t^{1-1/2^{r+1}}(\boldsymbol{u}_r + \boldsymbol{v}_r)$ and $\boldsymbol{u}_{r+1} = t\boldsymbol{u}_r$. □

8.3.2 Curvature analysis

Let $[a,b]$ be an interval of $\mathbb{R}$, and $\Phi : [a,b] \to \mathbb{R}^d$ be the parametrization of a path in $\mathbb{R}^d$. Assume that Φ is twice continuously differentiable. For any $\lambda \in [a,b]$, the arc length of the path Φ between $\Phi(a)$ and $\Phi(\lambda)$ is $\ell(\lambda) := \int_a^\lambda \|\dot{\Phi}(\gamma)\| \mathrm{d}\gamma$. Let $\Phi : [0, \ell(b)] \to \mathbb{R}^d$ be the parameterization of $\Phi([a,b])$ by its arc length, *i.e.*, $\Phi(\ell(\lambda)) := \Phi(\lambda)$ for all $\lambda \in [a,b]$. As a consequence, $\dot{\Phi}(\ell(\lambda)) = \dot{\Phi}(\lambda)/\|\dot{\Phi}(\lambda)\|$. Thus, $\dot{\Phi}$ describes a path on the unit sphere $\mathcal{S}^{d-1} \subseteq \mathbb{R}^d$. The length of the latter path, $\int_0^{\ell(b)} \|\ddot{\Phi}(\tau)\| \mathrm{d}\tau$, is the *total curvature* of Φ between $\Phi(a)$ and $\Phi(b)$.

The total curvature can also be defined in terms of angles. Given points $U, V, W \in \mathbb{R}^d$, we shall denote by $\angle UVW$ the measure $\alpha \in [0, \pi]$ of the angle between the vectors $V - U$ and $W - V$, so that

$$\cos \alpha = \langle \frac{(V-U)}{\|V-U\|}, \frac{(W-V)}{\|W-V\|} \rangle \ ,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard scalar product of $\mathbb{R}^d$, and $\| \cdot \|$ denotes the associated Euclidean norm.

If $\tau : [a,b] \to \mathbb{R}^d$ parametrizes a polygonal line $[X^0, X^1] \cup [X^1, X^2] \cup \dots [X^q, X^{q+1}]$, the total curvature $\kappa(\tau, [a,b])$ is defined as the sum of angles between consecutive segments:

$$\kappa(\tau, [a,b]) := \sum_{k=1}^{q} \angle X^{k-1} X^k X^{k+1} \ .$$

A polygonal line $\tau : [a,b] \to \mathbb{R}^d$ is *inscribed* in a path $\Phi : [a,b] \to \mathbb{R}^d$ if there exists a subdivision $a = \lambda^0 < \lambda^1 < \cdots < \lambda^{q+1} = b$ such that $X^k = \Phi(\lambda^k)$ for all $0 \leq k \leq q+1$.

The total curvature $\kappa(\Phi, [a,b])$ can be defined for an arbitrary curve Φ, as the supremum of $\kappa(\tau, [a,b])$ over all polygonal curves τ inscribed in Φ. When Φ is twice continuously differentiable, this coincides with the previous definition of the total curvature, see Chapter V of [AR89] for more background.

Tropical lower bounds on the curvature of a definable family of paths

Now consider an interval $[a,b]$ of $\mathbb{R}$ and $\boldsymbol{\rho} : [a,b] \to \mathbb{K}^d$ a path in $\mathbb{K}^d$. Since the elements of $\mathbb{K}$ are real valued functions, the path $\boldsymbol{\rho}$ parametrizes a family of paths in $\mathbb{R}^d$. For a fixed real number M let us define the map $\rho : (M, +\infty) \times [a,b] \to \mathbb{R}^d$ by $\rho(t, \lambda) = \boldsymbol{\rho}(\lambda)(t)$. For each t large enough, $\rho(t, \cdot)$ parametrizes a path in $\mathbb{R}^d$. We now derive lower bounds on the total curvature of the paths $\rho(t, \cdot)$ using $\rho^{\mathbb{T}} = \operatorname{val}(\boldsymbol{\rho})$.

Lemma 8.16. *Let $\boldsymbol{\rho} : [a,b] \to \mathbb{K}^d$ be a path in $\mathbb{K}^d$ and $\lambda_1 < \lambda_2 < \lambda_3$ three scalars in the interval $[a,b]$. Suppose that $\rho^{\mathbb{T}} = \operatorname{val}(\boldsymbol{\rho})$ satisfies :*

$$\max_{1 \leq i \leq d} \rho_i^{\mathbb{T}}(\lambda_1) < \max_{1 \leq i \leq d} \rho_i^{\mathbb{T}}(\lambda_2) < \max_{1 \leq i \leq d} \rho_i^{\mathbb{T}}(\lambda_3), \quad \textit{and} \quad \operatorname*{arg\,max}_{1 \leq i \leq d} \rho_i^{\mathbb{T}}(\lambda_2) \cap \operatorname*{arg\,max}_{1 \leq i \leq d} \rho_i^{\mathbb{T}}(\lambda_3) = \emptyset.$$

Then,

$$\lim_{t \to \infty} \angle \rho(t, \lambda_1) \rho(t, \lambda_2) \rho(t, \lambda_3) = \frac{\pi}{2} \ .$$

Proof. By definition of the valuation map, for all $\varepsilon > 0$ small enough, we have the inequalities

$$t^{\rho_i^{\mathbb{T}}(\lambda)-\varepsilon} \leq |\boldsymbol{\rho}_i(\lambda)| \leq t^{\rho_i^{\mathbb{T}}(\lambda)+\varepsilon} \quad \text{for all } i \in [d] \text{ and } \lambda \in [a,b] \ .$$

For $k = 1,2,3$, let m_k be the maximum of the entries of $\rho^{\mathbb{T}}(\lambda_k)$, and denote $I_k := \arg\max_{1\leq i\leq d} \rho_i^{\mathbb{T}}(\lambda_k)$. Also denote by

$$m'_k := \max\{\rho_i^{\mathbb{T}}(\lambda_k) \mid i \in [d] \setminus I_k\}$$

the value of the second maximal coordinate of $\rho^{\mathbb{T}}(\lambda_k)$. We can choose $\varepsilon > 0$ such that

$$\begin{aligned} m_2 &> m_1 + 3\varepsilon \\ m_2 &> m'_2 + 3\varepsilon \ . \end{aligned}$$

Consider the vector $\boldsymbol{\xi} := \boldsymbol{\rho}(\lambda^2) - \boldsymbol{\rho}(\lambda^1)$, and any $i \in I_2$. By our choice of ε, we have $m_2 > \rho_i^{\mathbb{T}}(\lambda_1) + 3\varepsilon$. Consequently, we can bound the norm of $\boldsymbol{\xi}$ as follows:

$$\begin{aligned} \|\boldsymbol{\xi}\| \geq |\boldsymbol{\rho}_i(\lambda_2)| - |\boldsymbol{\rho}_i^{\mathbb{T}}(\lambda_1)| &\geq t^{\rho_i^{\mathbb{T}}(\lambda_2)-\varepsilon} - t^{\rho_i^{\mathbb{T}}(\lambda_1)+\varepsilon} \\ &\geq t^{m_2-\varepsilon}(1 - t^{-m_2+\rho_i^{\mathbb{T}}(\lambda_1)+2\varepsilon}) \\ &\geq t^{m_2-\varepsilon}(1 - t^{-\varepsilon}) \ . \end{aligned}$$

Now consider the normalized vector $\bar{\boldsymbol{\xi}} := \boldsymbol{\xi}/\|\boldsymbol{\xi}\|$. By our choice of ε, for any $j \in [d] \setminus I_2$, we have:

$$\begin{aligned} m_2 &> \rho_j^{\mathbb{T}}(\lambda_1) + 3\varepsilon \\ m_2 &> \rho_j^{\mathbb{T}}(\lambda_2) + 3\varepsilon \ . \end{aligned}$$

Consequently, for any $j \in [d] \setminus I_2$ the component $\bar{\boldsymbol{\xi}}_j$ of the normalized vector satisfies:

$$|\bar{\boldsymbol{\xi}}|_j \leq \frac{t^{\rho_j^{\mathbb{T}}(\lambda_2)+\varepsilon} + t^{\rho_j^{\mathbb{T}}(\lambda_1)+\varepsilon}}{t^{m_2-\varepsilon}(1-t^{-\varepsilon})} \leq \frac{t^{-m_2+\rho_j^{\mathbb{T}}(\lambda_2)+2\varepsilon} + t^{-m_2+\rho_j^{\mathbb{T}}(\lambda_1)+2\varepsilon}}{1-t^{-\varepsilon}} \leq \frac{2t^{-\varepsilon}}{1-t^{-\varepsilon}} = \frac{2}{t^{\varepsilon}-1}$$

Consequently, $\bar{\boldsymbol{\xi}}_j(t)$ tends to 0 as t tends to infinity for $j \in [d] \setminus I_2$. Observe that the map $t \to \bar{\boldsymbol{\xi}}(t)$ is definable in the polynomially bounded structure $\mathbb{R}^{\mathbb{R}}$. Since $\|\bar{\boldsymbol{\xi}}\| = 1$, we deduce that $\bar{\boldsymbol{\xi}}(t)$ has a limit $\bar{\boldsymbol{\xi}}(\infty)$ as t tends to infinity. Clearly, $\bar{\boldsymbol{\xi}}_j(\infty) = 0$ for all $j \in [d] \setminus I_2$.

Similarly, let $\boldsymbol{\eta} = \boldsymbol{\rho}(\lambda_3) - \boldsymbol{\rho}(\lambda_2)$ and $\bar{\boldsymbol{\eta}} = \boldsymbol{\eta}/\|\boldsymbol{\eta}\|$. We deduce that $\bar{\boldsymbol{\eta}}_j(\infty) = 0$ for all $j \in [d] \setminus I_3$, where $\bar{\boldsymbol{\eta}}(\infty)$ denotes the limit of $t \mapsto \bar{\boldsymbol{\eta}}(t)$ as $t \to \infty$. As $I_2 \cap I_3 = \emptyset$, we have $\bar{\boldsymbol{\xi}}(\infty) \cdot \bar{\boldsymbol{\eta}}(\infty) = 0$. We deduce that

$$\lim_{t\to\infty} \angle\rho(t,\lambda_1)\rho(t,\lambda_2)\rho(t,\lambda_3) = \lim_{t\to\infty} \arccos(\bar{\boldsymbol{\xi}}(t) \cdot \bar{\boldsymbol{\eta}}(t)) = \arccos(\bar{\boldsymbol{\xi}}(\infty) \cdot \bar{\boldsymbol{\eta}}(\infty)) = \frac{\pi}{2} \ . \ \square$$

We define the *combinatorial angle* $\angle^c \rho^{\mathbb{T}}(\lambda_1)\rho^{\mathbb{T}}(\lambda_2)\rho^{\mathbb{T}}(\lambda_3)$ of the points $\rho^{\mathbb{T}}(\lambda_1)$, $\rho^{\mathbb{T}}(\lambda_2)$ and $\rho^{\mathbb{T}}(\lambda_2)$ to be 1 if the conditions of Lemma 8.16 are satisfied. Otherwise, the combinatorial angle is defined to be 0. Given a subdivision $a = \lambda_0 < \cdots < \lambda_{q+1} = b$ of an interval $[a,b] \subseteq \mathbb{R}$, we denote by $\kappa^c(\rho^{\mathbb{T}}; \lambda_0, \ldots, \lambda_{q+1})$ the sum of combinatorial angles

$$\sum_{k \in [q]} \angle^c \rho^{\mathbb{T}}(\lambda_{k-1})\rho^{\mathbb{T}}(\lambda_k)\rho^{\mathbb{T}}(\lambda_{k+1}) \ .$$

Finally, we define the *total combinatorial curvature* of $\rho^{\mathbb{T}}$ over the interval $[a,b]$, denoted by $\kappa^c(\rho^{\mathbb{T}}, [a,b])$, to be the supremum of $\kappa^c(\rho^{\mathbb{T}}; \lambda_0, \ldots, \lambda_{q+1})$ over all subdivisions of the interval $[a,b]$.

Theorem 8.17. *For all real numbers $a < b$, we have*

$$\lim_{t \to \infty} \kappa(\rho(t, \cdot), [a,b]) \geq \kappa^c(\rho^{\mathbb{T}}, [a,b]) \frac{\pi}{2} \ .$$

Proof. Consider any subdivision $a = \lambda_0 < \cdots < \lambda_{q+1} = b$. By Lemma 8.16, for all $k \in [q]$, we have:

$$\lim_{t \to \infty} \angle \rho(t, \lambda_{k-1})\rho(t, \lambda_k)\rho(t, \lambda_{k+1}) \geq \angle^c \rho^{\mathbb{T}}(\lambda_{k-1})\rho^{\mathbb{T}}(\lambda_k)\rho^{\mathbb{T}}(\lambda_{k+1}) \frac{\pi}{2} \ .$$

It follows that,

$$\begin{aligned} \lim_{t \to \infty} \kappa(\rho(t, \cdot), [a,b]) &\geq \sum_{k \in [q]} \lim_{t \to \infty} \angle \rho(t, \lambda_{k-1})\rho(t, \lambda_k)\rho(t, \lambda_{k+1}) \\ &\geq \sum_{k \in [q]} \angle^c \rho^{\mathbb{T}}(\lambda_{k-1})\rho^{\mathbb{T}}(\lambda_k)\rho^{\mathbb{T}}(\lambda_{k+1}) \ . \end{aligned}$$

Finally, the conclusion of the theorem is obtained by taking the maximum over all subdivisions. □

In general, the information provided by the valuation may not be enough to infer the total curvature, and so, the bound of Theorem 8.17 is not expected to be tight in general.

8.3.3 Application to the counter-example

Given any integer $r \geq 1$, the Hardy linear program $\mathbf{LP}_r$ gives rise to a family real linear programs $\mathbf{LP}_r(t)$ for t large enough, that are parametrized by $\mathcal{C}(t, \cdot)$. With the notation of Lemma 8.2, we define a path

$$\begin{aligned} \boldsymbol{\mathcal{C}} : \mathbb{R} &\to (\mathbb{K}^{r+1})^3 \times \mathbb{K}^r \times \mathbb{K}^{r+1} \\ \lambda &\mapsto \big(\boldsymbol{u}^{\boldsymbol{\mu}}, \boldsymbol{v}^{\boldsymbol{\mu}}, \boldsymbol{z}^{\boldsymbol{\mu}}, (\boldsymbol{z}')^{\boldsymbol{\mu}}, \boldsymbol{h}^{\boldsymbol{\mu}}\big) \qquad \text{where } \boldsymbol{\mu} = t^{\lambda} \ . \end{aligned}$$

Hence, $\mathcal{C}(t, \cdot) = \boldsymbol{\mathcal{C}}(\cdot)(t)$ parametrize the central path of $\mathbf{LP}_r(t)$.

We first analyze the curvature of the $(\boldsymbol{u}, \boldsymbol{v})$ components of the central paths. We define Φ to be the projection of $\mathcal{C}$ on the $(\boldsymbol{u}, \boldsymbol{v})$ components.

Theorem 8.18. *We have*

$$\lim_{t\to\infty} \kappa(\Phi(t,\cdot),[0,2]) \geq (2^r-1)\frac{\pi}{2} \enspace .$$

Proof. Consider the subdivision $0 = \lambda_0 < \cdots < \lambda_r = 2$ given by $\lambda_k = 4k/2^r$ for $k = 0,\dots,2^r$. can readily check from Proposition 8.13 that the combinatorial angles $\angle^c \Phi^{\mathbb{T}}(\lambda_0)\Phi^{\mathbb{T}}(\lambda_1)\Phi^{\mathbb{T}}(\lambda_2)$, ..., $\angle^c \Phi^{\mathbb{T}}(\lambda_{2^r-2})\Phi^{\mathbb{T}}(\lambda_{2^r-1})\Phi^{\mathbb{T}}(\lambda_{2^r})$ are all equal to one. Actually, the maximum of the coordinates of $\Phi^{\mathbb{T}}(\lambda_k)$ is attained alternatively by the components u_r and v_r, depending on the parity of k, and it is a strictly increasing function of k. Then, the conclusion follows from Theorem 8.17. □

We now turn to the whole central path $\mathcal{C}$ of $\mathbf{LP}_r(t)$.

Theorem 8.19. *We have*

$$\lim_{t\to\infty} \kappa(\mathcal{C}(t,\cdot),[0,2]) \geq (2^{r-1}-1)\frac{\pi}{2} \enspace .$$

Proof. Define now $\lambda_k = 4k/2^r$, for $k = 0,\dots,2^{r-1}$. It easily follows from Propositions 8.12 and 8.13 that the combinatorial angles

$$\angle^c \mathcal{C}^{\mathbb{T}}(\lambda_0)\mathcal{C}^{\mathbb{T}}(\lambda_1)\mathcal{C}^{\mathbb{T}}(\lambda_2),\dots,\angle^c \mathcal{C}^{\mathbb{T}}(\lambda_{2^{r-1}-2})\mathcal{C}^{\mathbb{T}}(\lambda_{2^{r-1}-1})\mathcal{C}^{\mathbb{T}}(\lambda_{2^{r-1}})$$

are all equal to one. The maximizing variables of the tropical central path at all these points $\lambda_0,\dots,\lambda_{2^{r-1}}$ are alternatively z_r and z'_r. Then, the conclusion follows from Theorem 8.17. □

Chapter 9

Conclusion and perspectives

In Chapter 3, we tropicalized the simplex method. The key idea is to compute the sign of a polynomial by tropical means. This idea could lead to the tropicalitation of other kinds of algorithms, even unrelated to linear programming. More precisely, one could tropicalize in this way any *semi-algebraic* algorithm, *i.e.*, that rely only the signs of polynomials evaluated on the input. However, in order to obtain a tropical algorithm which runs in polynomial time, the polynomials must satisfy some conditions. In particular, the "size" of the polynomials, measured by the magnitude of their exponents, should not be too large.

It would also be interesting to consider the quantization of tropical algorithm, *i.e.*, to apply tropical algorithms to classical problems. Under which conditions does a tropical semi-algebraic algorithm provide an algorithm for arbitrary classical problems? For example, the policy iteration algorithm for mean payoff games could provide a new algorithm for classical linear programming. This question is related to the realizability of classical polyhedra as tropical polyhedra discussed below.

In Chapter 4, we used the tropicalization of the simplex method to solve arbitrary tropical linear program. Our main tool is a perturbation scheme that rely on groups of higher order rank. Our perturbation transforms an arbitrary problem into a problem which is generic for *any* polynomial. Hence, this approach could be used with the tropicalization of other algorithms than the simplex method. This perturbation scheme could have further applications in tropical geometry. In particular, it would be worthwhile to compare it to the concept of stable intersection.

In Chapter 5, we obtain a transfer principle from classical linear programming to tropical linear programming via the simplex method. We showed that a polynomial time pivoting rule for the simplex method could yield a polynomial time algorithm for tropical linear programming. The most natural question is whether the converse statement holds. From our point of view, this question boils down to the realizability of classical polyhedra as tropical polyhedra.

Question 9.1. *Is any (non-degenerate) classical polyhedra combinatorially realizable as a tropical polyhedra?*

A positive answer to the this question would entail a transfer principle from tropical linear programming to classical linear programming. This could show that Smale's problem on the existence of a strongly polynomial algorithm for classical linear programming somehow reduces to the NP $\cap$ co-NP problem of tropical linear programming and mean-payoff games. Indeed, Theorem 5.4 indicates that polynomial algorithms for tropical linear programming could provide strongly polynomial algorithms for classical linear programming.

Chapter 5 also presents a class of classical linear programs on which the simplex method is polynomial in the bit model. This class is obtained by quantization of edge-improving tropical linear programs. However, it does not seem easy to decide whether a classical linear program belongs to this class. It would be interesting to study alternative characterizations of these problems. Since the simplex method is polynomial on this class of instances, this suggests that polyhedra with large diameter do *not* belong to it. Moreover, given such a classical instance, one can ask for a way to compute the corresponding tropical problem. Indeed, this would permit to use the tropical simplex method to solve these classical instances.

The tropicalization of the shadow-vertex rule in Chapter 6 allowed us to derive the first algorithm with a polynomial average-case complexity for mean payoff games. The shadow-vertex rule is used in several significant results. Can we tropicalize the randomized polynomial-time algorithm of Kelner and Spielman [KS06]? Or the smoothed-complexity result of Spielman and Teng [ST04]?

In Chapter 7, we proposed an efficient implementation for the tropical pivoting operation and the computation of tropical reduced costs. These procedures use $O(n(m+n))$ tropical operations for a linear program described by m inequalities on n variables. It would be interesting to take advantage of sparsity. Preliminary results indicate that these procedures could be implemented in $O(k + m\log(m) + n)$ operations, where k is the number of non $\mathbb{0}$ entries of the input.

Finally, in Chapter 8 we studied the tropicalization of the central path. We showed that the tropical central path has a geometric description, and that it may coincide with a run of the tropical simplex method. This could lead to a "central path" pivoting rule for the simplex method. We also disproved the continuous analogue of the Hirsch conjecture by exhibiting a family of real linear programs constrained by $3r+4$ inequalities in dimension $2r+2$ with a total curvature of $\Omega(2^r)$. This family is parametrized by a real number t that must be large enough. A necessary next step is to bound the minimal value of t for which the total curvature is $\Omega(2^r)$. Preliminary results indicate that $t = 2^{r2^r}$ is enough. An interesting question is also to what extent the total curvature can be worse than $\Omega(2^r)$? Can we obtain a total curvature of $\Omega(2^{2^r})$, or even of arbitrary tower of exponentials? A step in this direction would be to carry the idea underlying the tropical linear program used in Chapter 8 over to tropical semirings of higher rank, and then lift to the Hardy field of the structure $\bar{\mathbb{R}}_{\exp}$.

Bibliography

[ABG14] X. Allamigeon, P. Benchimol, and S. Gaubert. The tropical shadow-vertex algorithm solves mean payoff games in polynomial time on average. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 89–100. Springer, 2014.

[ABGJ13a] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Combinatorial simplex algorithms can solve mean payoff games. To appear in SIAM J. on Optimization. E-print arXiv:1309.5925, 2013.

[ABGJ13b] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Tropicalizing the simplex algorithm. To appear in SIAM Disc. Math. E-print arXiv:1308.0454, 2013.

[ABGJ14] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Long and winding central paths. E-print arXiv:1405.4161, 2014.

[AC78] D. Avis and V. Chvátal. Notes on Bland's pivoting rule. In *Polyhedral Combinatorics*, pages 24–34. Springer, 1978.

[AF13] D. Avis and O. Friedmann. An exponential lower bound for Cunningham's rule. E-print arXiv:1305.3944, 2013.

[AGG09] M. Akian, S. Gaubert, and A. Guterman. Linear independence over tropical semirings and beyond. In G. Litvinov and S. Sergeev, editors, *Proceedings of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemporary Mathematics*, pages 1–38. AMS, 2009.

[AGG10] X. Allamigeon, S. Gaubert, and É. Goubault. The tropical double description method. In J.-Y. Marion and T. Schwentick, editors, *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science (STACS 2010)*, volume 5 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47–58, Dagstuhl, Germany, 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[AGG12] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(1):125001, 2012.

[AGG13] X. Allamigeon, S. Gaubert, and E. Goubault. Computing the vertices of tropical polyhedra using directed hypergraphs. *Discrete & Computational Geometry*, 49(2):247–279, 2013. E-print arXiv:0904.3436v4.

[AGG14] M. Akian, S. Gaubert, and A. Guterman. Tropical Cramer determinants revisited. In G. Litvinov and S. Sergeev, editors, *Proceedings of the International Conference on Tropical and Idempotent Mathematics*, volume 616 of *Contemporary Mathematics*, pages 1–45. American Mathematical Society, 2014.

[AGK11] X. Allamigeon, S. Gaubert, and R. Katz. The number of extreme points of tropical polyhedra. *Journal of Combinatorial Theory, Series A*, 118(1):162 – 189, 2011.

[AGNS11] M. Akian, S. Gaubert, V. Nitica, and I. Singer. Best approximation in max-plus semimodules. *Linear Algebra and its Applications*, 435(12):3261–3296, 2011.

[AK13] X. Allamigeon and R. Katz. Minimal external representations of tropical polyhedra. *Journal of Combinatorial Theory, Series A*, 120(4):907–940, 2013.

[AKS86] I. Adler, R. Karp, and R. Shamir. A family of simplex variants solving an $m \times d$ linear program in expected number of pivot steps depending on d only. *Mathematics of Operations Research*, 11(4):570–590, 1986.

[AKS87] I. Adler, R. M. Karp, and R. Shamir. A simplex variant solving an $m \times d$ linear program in $O(\min(m^2, d^2))$ expected number of pivot steps. *Journal of Complexity*, 3(4):372 – 387, 1987.

[AKW06] F. Ardila, C. Klivans, and L. Williams. The positive Bergman complex of an oriented matroid. *European Journal of Combinatorics*, 27(4):577–591, 2006.

[Ale13] D. Alessandrini. Logarithmic limit sets of real semi-algebraic sets. *Adv. Geometry*, 13(1):155–190, 2013.

[Ans91] K. Anstreicher. On the performance of Karmarkar's algorithm over a sequence of iterations. *SIAM Journal on Optimization*, 1(1):22–29, 1991.

[AR89] A. D. Alexandrov and Y. G. Reshetnyak. *General theory of irregular curves.* Kluwer, 1989.

[AZ96] N. Amenta and G. M. Ziegler. Deformed products and maximal shadows of polytopes. In *Advances in Discrete and Computational Geometry*, Contemporary Mathematics 223, Providence, R.I., 1996. Amer. Math. Soc.

[BA08] P. Butkovič and A. Aminu. Introduction to max-linear programming. *IMA Journal of Management Mathematics*, 20(3):233–249, 2008.

[BCOQ92] F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and linearity: an algebra for discrete event systems*. Wiley, 1992.

[BDSE+12] N. Bonifas, M. Di Summa, F. Eisenbrand, N. Hähnle, and M. Niemeier. On sub-determinants and the diameter of polyhedra. In *Proceedings of the 2012 symposuim on Computational Geometry*, pages 357–362. ACM, 2012.

[Ben14] P. Benchimol. Simplet, a solver for tropical linear programming. Available at http://simplet.gforge.inria.fr, 2014.

[Ber71] G. Bergman. The logarithmic limit-set of an algebraic variety. *Transactions of the American Mathematical Society*, 157:459–469, 1971.

[BH84] P. Butkovič and G. Hegedüs. An elimination method for finding all solutions of the system of linear equations over an extremal algebra. *Ekonomicko-matematicky Obzor*, 20:203–215, 1984.

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

[BHR05] W. Briec, C. Horvath, and A. Rubinov. Separation in $\mathbb{B}$-convexity. *Pacific Journal of Optimization*, 1:13–30, 2005.

[BJS+07] T. Bogart, A. N. Jensen, D. Speyer, B. Sturmfels, and R. R. Thomas. Computing tropical varieties. *J. Symbolic Comput.*, 42(1-2):54–73, 2007.

[BL97] D. Bertsimas and X. Luo. On the worst case complexity of potential reduction algorithms for linear programming. *Mathematical Programming*, 77(2):321–333, May 1997.

[Bla77] R. G. Bland. New finite pivoting rules for the simplex method. *Mathematics of Operations Research*, 2(2):103–107, 1977.

[BM14a] P. Butkovic and M. MacCaig. A strongly polynomial method for solving integer max-linear optimization problems in a generic case. *Journal of Optimization Theory and Applications*, pages 1–23, 2014.

[BM14b] P. Butkovič and M. MacCaig. On the integer max-linear programming problem. *Discrete Applied Mathematics*, 162:128–141, 2014.

[BNRC08] M. Bezem, R. Nieuwenhuis, and E. Rodríguez-Carbonell. Exponential behaviour of the Butkovič-Zimmermann algorithm for solving two-sided linear systems in max-algebra. *Discrete Appl. Math.*, 156(18):3506–3509, 2008.

[Bor87] K. H. Borgwardt. *The simplex method: a probabilistic analysis*, volume 1 of *Algorithms and Combinatorics*. Springer Verlag, 1987.

[BR14] T. Brunsch and H. Röglin. Finding short paths on polytopes by the shadow vertex algorithm. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2014.

[Bru12] E. Brugallé. Some aspects of tropical geometry. *Newsletter of the European Mathematical Society,(83)*, pages 23–28, 2012.

[BSS07] P. Butkovič, H. Schneider, and S. Sergeev. Generators, extremals and bases of max cones. *Linear Algebra Appl.*, 421(2-3):394–406, 2007.

[But03] P. Butkovič. Max-algebra: the linear algebra of combinatorics? *Linear Algebra and its applications*, 367:313–335, 2003.

[But10] P. Butkovič. *Max-linear systems: theory and algorithms.* Springer, 2010.

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

[BY06] F. Block and J. Yu. Tropical convexity via cellular resolutions. *J. Algebraic Combin.*, 24(1):103–114, 2006.

[BZ06] P. Butkovič and K. Zimmermann. A strongly polynomial algorithm for solving two-sided linear systems in max-algebra. *Discrete Applied Mathematics*, 154(3):437–446, 2006.

[CC58] A. Charnes and W. Cooper. The strong Minkowski-Farkas-Weyl theorem for vector spaces over ordered fields. *Proceedings of the National Academy of Sciences of the United States of America*, 44(9):914, 1958.

[CD69] P. Conrad and J. Dauns. An embedding theorem for lattice-ordered fields. *Pacific Journal of Mathematics*, 30(2):385–398, 1969.

[CG79] R. Cuninghame-Green. *Minimax algebra*, volume 166 of *Lecture Notes in Economics and Mathematical Systems.* Springer-Verlag, Berlin, 1979.

[CGB03] R. Cuninghame-Green and P. Butkovič. The equation $a \otimes x = b \otimes y$ over (max,+). *Theoretical Computer Science*, 293:3–12, 2003.

[CGQ99] G. Cohen, S. Gaubert, and J.-P. Quadrat. Max-plus algebra and system theory: where we are and where to go now. *Annual Reviews in Control*, 23:207–219, 1999.

[CGQ01] G. Cohen, S. Gaubert, and J. P. Quadrat. Hahn-Banach separation theorem for max-plus semimodules. In J. L. Menaldi, E. Rofman, and A. Sulem, editors, *Optimal Control and Partial Differential Equations*, pages 325–334. IOS Press, 2001.

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

[CGQS05] G. Cohen, S. Gaubert, J.-P. Quadrat, and I. Singer. Max-plus convex sets and functions. In G. Litvinov and V. Maslov, editors, *Idempotent Mathematics and Mathematical Physics*, volume 377 of *Contemporary Mathematics*, pages 105–129. Amer. Math. Soc., 2005.

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

[CMQV89] G. Cohen, P. Moller, J.-P. Quadrat, and M. Viot. Algebraic tools for the performance evaluation of discrete event systems. *Proceedings of the IEEE*, 77(1):39–85, 1989.

[Cos00] M. Coste. An introduction to o-minimal geometry, 2000. Dip. Mat. Univ. Pisa, Dottorato di Ricerca in Matematica, Istituti Editoriali e Poligrafici Internazionali, Pisa.

[CTGG99] J. Cochet-Terrasson, S. Gaubert, and J. Gunawardena. A constructive fixed-point theorem for min-max functions. *Dynamics and Stability of Systems*, 14(4):407–433, 1999.

[Dan98] G. B. Dantzig. *Linear programming and extensions.* Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, corrected edition, 1998.

[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)*, volume 180, Pisa, Italy, 2006. article No. 12.

[DL11] J. A. De Loera. New insights into the complexity and geometry of linear optimization. *OPTIMA, Newsletter of the Math. Optim. Soc.*, 87:1–16, 2011.

[DLGKL10] M. Di Loreto, S. Gaubert, R. D. Katz, and J.-J. Loiseau. Duality between invariant spaces for max-plus linear discrete event systems. *SIAM Journal on Control and Optimization*, 48(8):5606–5628, 2010.

[DLKOS09] J. A. De Loera, E. D. Kim, S. Onn, and F. Santos. Graphs of transportation polytopes. *Journal of Combinatorial Theory, Series A*, 116(8):1306 – 1325, 2009.

[DLSV10] J. De Loera, B. Sturmfels, and C. Vinzant. The central curve in linear programming. *Arxiv preprint arXiv:1012.3978*, 2010.

[DMS05] J.-P. Dedieu, G. Malajovich, and M. Shub. On the curvature of the central path of linear programming theory. *Foundations of Computational Mathematics*, 5(2):145–171, 2005.

[DNPT06] A. Deza, E. Nematollahi, R. Peyghami, and T. Terlaky. The central path visits all the vertices of the Klee–Minty cube. *Optimisation Methods and Software*, 21(5):851–865, 2006.

[Dre86] A. Dress. Duality theory for finite and infinite matroids with coefficients. *Adv. in Math.*, 59(2):97–123, 1986.

[DS04] M. Develin and B. Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27 (electronic), 2004. correction: ibid., pp. 205–206.

[DS05] J.-P. Dedieu and M. Shub. Newton flow and interior point methods in linear programming. *International Journal of Bifurcation and Chaos*, 15(03):827–839, 2005.

[DTZ08] A. Deza, T. Terlaky, and Y. Zinchenko. Polytopes and arrangements: diameter and curvature. *Operations Research Letters*, 36(2):215–222, 2008.

[DTZ09] A. Deza, T. Terlaky, and Y. Zinchenko. Central path curvature and iteration-complexity for redundant Klee-Minty cubes. In *Advances in applied mathematics and global optimization*, volume 17 of *Adv. Mech. Math.*, pages 223–256. Springer, New York, 2009.

[DW96] H. G. Dales and W. H. Woodin. *Super-real fields. Totally ordered fields with additional structure*, volume 14 of *London Mathematical Society Monographs. New Series.* The Clarendon Press, Oxford University Press, New York, 1996.

[DY07] M. Develin and J. Yu. Tropical polytopes and cellular resolutions. *Experiment. Math.*, 16(3):277–291, 2007.

[EKL06] M. Einsiedler, M. Kapranov, and D. Lind. Non-Archimedean amoebas and tropical varieties. *J. Reine Angew. Math.*, 601:139–157, 2006.

[EM79] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.

[ER89] B. C. Eaves and U. G. Rothblum. A theory on extending algorithms for parametric problems. *Mathematics of Operations Research*, 14(3):502–533, 1989.

[FAA02] J. Filar, E. Altman, and K. Avrachenkov. An asymptotic simplex method for singularly perturbed linear programs. *Oper. Res. Lett.*, 30(5):295–307, 2002.

[FHZ11] O. Friedmann, T. D. Hansen, and U. Zwick. Subexponential lower bounds for randomized pivoting rules for the simplex algorithm. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 283–292, New York, NY, USA, 2011. ACM.

[Fos10] T. Foster. *Power functions and exponentials in o-minimal expansions of fields*. PhD thesis, Oxford University, 2010.

[Fri11] O. Friedmann. A subexponential lower bound for Zadeh's pivoting rule for solving linear programs and games. In *Integer Programming and Combinatoral Optimization - 15th International Conference, IPCO 2011, New York, NY, USA, June 15-17, 2011. Proceedings*, pages 192–206, 2011.

[FT87] A. Frank and É. Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987.

[Fuc63] L. Fuchs. *Partially ordered algebraic systems*. International series of monographs in pure and applied mathematics. Pergamon Press, 1963.

[GGK04] J. Gilbert, C. Gonzaga, and E. Karas. Examples of ill-behaved central paths in convex optimization. *Mathematical Programming*, 103(1):63–94, December 2004.

[GK09] S. Gaubert and R. D. Katz. The tropical analogue of polar cones. *Linear Algebra Appl.*, 431(5-7):608–625, 2009. E-print arXiv:0805.3688.

[GK11] S. Gaubert and R. Katz. Minimal half-spaces and external representation of tropical polyhedra. *Journal of Algebraic Combinatorics*, 33(3):325–348, 2011.

[GKS12] S. Gaubert, R. Katz, and S. Sergeev. Tropical linear-fractional programming and parametric mean payoff games. *Journal of symbolic computation*, 47(12):1447–1478, December 2012.

[GKZ94] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Mathematics: Theory & Applications. Birkhäuser Boston, 1994.

[GLS88] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 1988.

[GM84] M. Gondran and M. Minoux. Linear algebra in dioids: a survey of recent results. *Annals of Discrete Mathematics*, 19:147–164, 1984.

[GM10] S. Gaubert and F. Meunier. Carathéodory, Helly and the others in the max-plus world. *Discrete and Computational Geometry*, 43(3):648–662, 2010.

[GMH14] V. M. Gonçalves, C. A. Maia, and L. Hardouin. On tropical fractional linear programming. *Linear Algebra and its Applications*, 459:384–396, 2014.

[Gol94] D. Goldfarb. On the complexity of the simplex method. In *Advances in optimization and numerical analysis*, pages 25–38. Springer, 1994.

[Gon12] J. Gondzio. Interior point methods 25 years later. *European Journal of Operational Research*, 218(3):587–601, 2012.

[GP97] S. Gaubert and M. Plus. Methods and applications of $(\max, +)$ linear algebra. In *Proceedings of the 14th Annual Symposium on Theoretical Aspects of Computer Science, Lybeck, Germany February 17 – March 1, 1997 (STACS 97)*, volume 1200 of *Lecture Notes in Comput. Sci.*, pages 261–282. Springer-Verlag, 1997.

[GS55] S. Gass and T. Saaty. The computational algorithm for the parametric objective function. *Naval research logistics quarterly*, 2(1-2):39–45, 1955.

[GS79] D. Goldfarb and W. Y. Sit. Worst case behavior of the steepest edge simplex method. *Discrete Applied Mathematics*, 1(4):277 – 285, 1979.

[GS07] S. Gaubert and S. Sergeev. Cyclic projectors and separation theorems in idempotent convex geometry. *Fundamentalnaya i prikladnaya matematika*, 13(4):33–52, 2007. Engl. translation in Journal of Mathematical Sciences (Springer, New-York), Vol. 155, No. 6, pp.815–829, 2008.

[Hah07] H. Hahn. Über die nichtarchimedischen Größensysteme. *Wien. Ber.*, 116:601–655, 1907.

[HKZ14] T. D. Hansen, H. Kaplan, and U. Zwick. Dantzig's pivoting rule for shortest paths, deterministic mdps, and minimum cost to time ratio cycles. In *SODA*, pages 847–860. SIAM, 2014.

[HOvdW06] B. Heidergott, G. J. Olsder, and J. W. van der Woude. *Max Plus at work: modeling and analysis of synchronized systems: a course on Max-Plus algebra and its applications*, volume 13. Princeton University Press, 2006.

[IKS03] I. Itenberg, V. Kharlamov, and E. Shustin. Welschinger invariant and enumeration of real rational curves. *International Mathematics research notices*, 2003(49):2639–2653, 2003.

[IMS07] I. Itenberg, G. Mikhalkin, and E. Shustin. *Tropical algebraic geometry*, volume 35 of *Oberwolfach Seminars*. Birkhäuser Verlag, Basel, 2007.

[IV96] I. Itenberg and O. Viro. Patchworking algebraic curves disproves the ragsdale conjecture. *The Mathematical Intelligencer*, 18(4):19–28, 1996.

[Jer73a] R. Jeroslow. Asymptotic linear programming. *Operations Research*, 21(5):1128–1141, 1973.

[Jer73b] R. G. Jeroslow. The simplex algorithm with the pivot rule of maximizing criterion improvement. *Discrete Mathematics*, 4(4):367–377, 1973.

[Jos05] M. Joswig. Tropical halfspaces. In *Combinatorial and computational geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 409–431. Cambridge Univ. Press, Cambridge, 2005.

[Jos09] M. Joswig. Tropical convex hull computations. In G. L. Litvinov and S. N. Sergeev, editors, *Proceedings of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemporary Mathematics*, pages 193–212. American Mathematical Society, 2009.

[JPZ06] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA 2006)*, January 2006.

[JSY07] M. Joswig, B. Sturmfels, and J. Yu. Affine buildings and tropical convexity. *Albanian J. Math.*, 1(4):187–211, 2007. E-print arXiv:0706.1918.

[JY94] J. Ji and Y. Yinyu. A complexity analysis for interior-point algorithms based on Karmarkar's potential function. *SIAM Journal on Optimization*, 4(3):512–520, 1994.

[Kap42] I. Kaplansky. Maximal fields with valuations. *Duke Mathematical Journal*, 9(2):303–321, 1942.

[Kar84] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, 1984.

[Kat07] R. D. Katz. Max-plus (A, B)-invariant spaces and control of timed discrete event systems. *IEEE Trans. Aut. Control*, 52(2):229–241, 2007. E-print arXiv:math.OC/0503448.

[Kha80] L. G. Khachiyan. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics*, 20(1):53–72, 1980.

[KK92] G. Kalai and D. J. Kleitman. A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bull. Amer. Math. Soc*, 26(2):315–316, 1992.

[KM72] V. Klee and G. J. Minty. How good is the simplex method? In *Inequalities III (Proceedings of the Third Symposium on Inequalities held at the University of California, Los Angeles, Calif., September 1–9, 1969, dedicated to the memory of Theodore S. Motzkin)*, pages 159–175, 1972.

[KM13a] T. Kitahara and S. Mizuno. A bound for the number of different basic solutions generated by the simplex method. *Mathematical Programming*, 137(1-2):579–586, 2013.

[KM13b] T. Kitahara and S. Mizuno. An upper bound for the number of different solutions generated by the primal simplex method with any selection rule of entering variables. *Asia-Pacific Journal of Operational Research*, 30(03), 2013.

[KOT13] S. Kakihara, A. Ohara, and T. Tsuchiya. Information geometry and interior-point algorithms in semidefinite programs and symmetric cone programs. *Journal of Optimization Theory and Applications*, 157(3):749–780, 2013.

[KS06] J. A. Kelner and D. A. Spielman. A randomized polynomial-time simplex algorithm for linear programming. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 51–60. ACM, 2006.

[KS10] E. D. Kim and F. Santos. An update on the Hirsch conjecture. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 112(2):73–98, 2010.

[KY91] J. Kaliski and Y. Ye. Convergence behavior of Karmarkar's projective algorithm for solving a simple linear program. *Operations research letters*, 10(7):389–393, 1991.

[Lit07] G. Litvinov. Maslov dequantization, idempotent and tropical mathematics: a brief introduction. *Journal of Mathematical Sciences*, 140(3):426–444, 2007.

[LL69] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Rev.*, 11:604–607, 1969.

[LMS01] G. L. Litvinov, V. P. Maslov, and G. B. Shpiz. Idempotent functional analysis. An algebraic approach. *Mat. Zametki*, 69(5):758–797, 2001.

[Mar02] D. Marker. *Model theory*, volume 217 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002. An introduction.

[Mar10] T. Markwig. A field of generalised Puiseux series for tropical geometry. *Rend. Semin. Mat., Univ. Politec. Torino*, 68(1):79–92, 2010.

[Meg87] N. Megiddo. On the complexity of linear programming. *Advances in economic theory*, pages 225–268, 1987.

[Mik05] G. Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *Journal of the American Mathematical Society*, 18(2):313–377, 2005.

[Mil94a] C. Miller. Expansions of the real field with power functions. *Annals of Pure and Applied Logic*, 0072(93), 1994.

[Mil94b] C. Miller. Exponentiation is hard to avoid. *Proceedings of the American Mathematical Society*, 122(1):257, September 1994.

[Mil12] C. Miller. Basics of o-minimality and Hardy fields. In *Lecture Notes on O-minimal Structures and Real Analytic Geometry*, pages 43–69. Springer, 2012.

[MS] D. Maclagan and B. Sturmfels. Introduction to tropical geometry. Draft of a book, `http://homepages.warwick.ac.uk/staff/D.Maclagan/papers/TropicalBook.html`.

[MS89] N. Megiddo and M. Shub. Boundary behavior of interior point algorithms in linear programming. *Mathematics of Operations Research*, 14(1):97–146, 1989.

[MT08] R. D. Monteiro and T. Tsuchiya. A strong bound on the integral of the central path curvature and its relationship with the iteration-complexity of primal-dual path-following LP algorithms. *Mathematical Programming*, 115(1):105–149, 2008.

[MT13a] M. Mut and T. Terlaky. A tight iteration-complexity upper bound for the MTY predictor-corrector algorithm via redundant Klee-Minty cubes. 2013.

[MT13b] M. Mut and T. Terlaky. An analogue of the Klee-Walkup result for Sonnevend's curvature of the central path. 2013.

[Mur80] K. G. Murty. Computational complexity of parametric linear programming. *Mathematical programming*, 19(1):213–219, 1980.

[Nad89] D. Naddef. The Hirsch conjecture is true for (0, 1)-polytopes. *Mathematical Programming*, 45(1):109–110, 1989.

[Orl97] J. B. Orlin. A polynomial time primal network simplex algorithm for minimum cost flows. *Mathematical Programming*, 78(2):109–129, 1997.

[Pas77] D. S. Passman. *The algebraic structure of group rings*. Pure and Applied Mathematics. Wiley-Interscience [John Wiley & Sons], New York-London-Sydney, 1977.

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

[Poo93] B. Poonen. Maximally complete fields. *Enseign. Math.(2)*, 39(1-2):87–106, 1993.

[Pow93] M. Powell. On the number of iterations of Karmarkar's algorithm for linear programming. *Mathematical Programming*, 62(1-3):153–197, 1993.

[Pur95] A. Puri. *Theory of hybrid systems and discrete event systems*. PhD thesis, University of California at Berkeley, Berkeley, CA, USA, 1995.

[RGST05] J. Richter-Gebert, B. Sturmfels, and T. Theobald. First steps in tropical geometry. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 2005.

[Rob77] A. Robinson. *Complete theories*. North Holland, 1977.

[RS02] Q. I. Rahman and G. Schmeisser. *Analytic theory of polynomials*. London Mathematical Society monographs. Oxford University Press, Oxford, 2002.

[San12] F. Santos. A counterexample to the Hirsch conjecture. *Annals of mathematics*, 176(1):383–412, 2012.

[Sch03] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Paths, flows, matchings, Chapters 1–38.

[Sei54] A. Seidenberg. A new decision method for elementary algebra. *Annals of Mathematics*, pages 365–374, 1954.

[Sin97] I. Singer. *Abstract convex analysis*. Wiley, 1997.

[Sma83] S. Smale. On the average number of steps of the simplex method of linear programming. *Mathematical Programming*, 27(3):241–262, 1983.

[Sma98] S. Smale. Mathematical problems for the next century. *Math. Intelligencer*, 20(2):7–15, 1998.

[SSZ91] G. Sonnevend, J. Stoer, and G. Zhao. On the complexity of following the central path of linear programs by linear extrapolation II. *Mathematical Programming*, 52(1-3):527–553, 1991.

[ST04] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.

[Ste10] S. A. Steinberg. *Lattice-ordered rings and modules.* Springer New York, 2010.

[SW05] D. Speyer and L. Williams. The tropical totally positive grassmannian. *Journal of Algebraic Combinatorics*, 22(2):189–210, 2005.

[SZ93] B. Sturmfels and A. Zelevinsky. Maximal minors and their leading terms. *Adv. Math.*, 98(1):65–112, 1993.

[Tab13] L. Tabera. On real tropical bases and real tropical discriminants. E-print arXiv:1311.2211, 2013.

[Tar51] A. Tarski. A decision method for elementary algebra and geometry. Research report R-109, Rand Corporation, 1951.

[The06] T. Theobald. On the frontiers of polynomial computations in tropical geometry. *Journal of Symbolic Computation*, 41(12):1360–1375, 2006.

[Tod14] M. J. Todd. An improved Kalai-Kleitman bound for the diameter of a polyhedron. E-print arXiv:1402.3579, 2014.

[TY96] M. Todd and Y. Ye. A lower bound on the number of iterations of long-step primal-dual linear programming algorithms. *Annals of Operations Research*, 62(1):233–252, 1996.

[Val79] L. G. Valiant. The complexity of computing the permanent. *Theoretical computer science*, 8(2):189–201, 1979.

[vdD98] L. van den Dries. *Tame topology and o-minimal structures*, volume 248. Cambridge University Press, 1998.

[vdDMM94] L. van den Dries, A. Macintyre, and D. Marker. The elementary theory of restricted analytic fields with exponentiation. *Annals of Mathematics*, 140(1):183–205, 1994.

[vdDS98] L. van den Dries and P. Speissegger. The real field with convergent generalized power series. *Trans. Amer. Math. Soc.*, 350(11):4377–4421, 1998.

[Vin12] C. Vinzant. Real radical initial ideals. *Journal of Algebra*, 352(1):392–407, 2012.

[Vir01] O. Viro. Dequantization of real algebraic geometry on logarithmic paper. In *European Congress of Mathematics*, pages 135–146. Springer, 2001.

[VY96] S. Vavasis and Y. Ye. A primal-dual interior point method whose running time depends only on the constraint matrix. *Mathematical Programming*, 74(1):79–120, 1996.

[Wri05] M. Wright. The interior-point revolution in optimization: history, recent developments, and lasting consequences. *Bulletin of the American mathematical society*, 42(1):39–56, 2005.

[Ye11] Y. Ye. The simplex and policy-iteration methods are strongly polynomial for the markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4):593–603, 2011.

[Zad80] N. Zadeh. What is the worst case behavior of the simplex algorithm? Technical Report 27, Departement of Operations Research, Stanford, 1980.

[Zim77] K. Zimmermann. A general separation theorem in extremal algebras. *Ekonom.-Mat. Obzor*, 13(2):179–201, 1977.

[ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoret. Comput. Sci.*, 158(1-2):343–359, 1996.

[ZS93] G. Zhao and J. Stoer. Estimating the complexity of a class of path-following methods for solving linear programs by curvature integrals. *Applied Mathematics and Optimization*, 27(1):85–103, 1993.