

What Tropical Geometry Tells Us about the Complexity of Linear Programming

Xavier Allamigeon, Pascal Benchimol, Stéphane Gaubert, Michael Joswig

▶ **To cite this version:**

Xavier Allamigeon, Pascal Benchimol, Stéphane Gaubert, Michael Joswig. What Tropical Geometry Tells Us about the Complexity of Linear Programming. SIAM Review, 2021, 63 (1), pp.123-164. 10.1137/20M1380211 . hal-03505719

HAL Id: hal-03505719

https://inria.hal.science/hal-03505719

Submitted on 31 Dec 2021

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.

WHAT TROPICAL GEOMETRY TELLS US ABOUT THE COMPLEXITY OF LINEAR PROGRAMMING *

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

Abstract. Tropical geometry has been recently used to obtain new complexity results in convex optimization and game theory. In this paper, we present an application of this approach to a famous class of algorithms for linear programming, i.e., log-barrier interior point methods. We show that these methods are not strongly polynomial, by constructing a family of linear programs with $3r + 1$ inequalities in dimension $2r$ for which the number of iterations performed is in $\Omega(2^r)$. The total curvature of the central path of these linear programs is also exponential in r, disproving a continuous analogue of the Hirsch conjecture proposed by Deza, Terlaky, and Zinchenko. These results are obtained by analyzing the tropical central path, which is the piecewise linear limit of the central paths of parameterized families of classical linear programs viewed through "logarithmic glasses." This allows us to provide combinatorial lower bounds for the number of iterations and the total curvature, in a general setting.

Key words. linear programming, central path, strongly polynomial complexity, continuous analogue of the Hirsch conjecture, tropical geometry

AMS subject classifications. 90C51, 14T05

1. Introduction.

An open question in computational optimization asks whether (the feasibility problem of) linear programming can be solved with a strongly polynomial algorithm. Strong polynomiality requires the algorithm to be polynomial time in the bit model (meaning that the execution time is polynomial in the number of bits of the input) *and* the number of arithmetic operations to be bounded by a polynomial in the number of numerical entries of the input, uniformly in their bit length.

This problem has been stated by Smale among his list of problems for the 21^{st} century [Sma00]. It comes as a generalization of a question that goes back to the development of the simplex method by Dantzig, and which asks whether the number of iterations of the simplex method is bounded by a polynomial in the number of variables and constraints. This has motivated substantial research on the complexity of the simplex method as well as on the diameter of polyhedra in relation with the Hirsch conjecture. We refer to [KS10, Zie12, DL19, DH20] for recent surveys on this topic. Moreover, relevant classes of linear programs with strongly polynomial complexity have been identified [Tar86, Ye11].

The development of interior point methods after the work of Karmarkar has brought an additional perspective to Smale's problem. Interior point methods are

* During this work, the second author was affiliated with INRIA Saclay Île-de-France and CMAP, École Polytechnique, CNRS UMR 7641.

Funding: The first and third authors were partially supported by the PGMO program of EDF and Fondation Mathématique Jacques Hadamard. The second author was supported by a Ph.D. fellowship from DGA and École Polytechnique. The fourth author was partially supported by Einstein Foundation Berlin; by DFG (SPP 1489, SFB/TRR 109, SFB/TRR 195); by a CNRS INSMI visiting professorship at CMAP, École Polytechnique, UMR 7641 and at IMJ, Université Pierre et Marie Curie, UMR 7586; and by the National Science Foundation under grant 1440140 while at the Mathematical Sciences Research Institute in Berkeley, California, during the fall of 2017.

†INRIA and CMAP, CNRS, École Polytechnique, Institut Polytechnique de Paris, 91128 Palaiseau Cedex, France (xavier.allamigeon@inria.fr, stephane.gaubert@inria.fr).

‡Kisio Digital, 20 rue Hector Malot, 75012 Paris, France (pascal.benchimol@polytechnique.edu).

§Technische Universität Berlin, Chair of Discrete Mathematics/Geometry, and MPI MiS, Leipzig, Germany (joswig@math.tu-berlin.de).

extensively used as they are efficient both in theory (i.e., polynomial time in the bit model [Kar84]) and in practice. They are driven to an optimal solution along a trajectory called the *central path*. Early on, Bayer and Lagarias recognized that the central path is "a fundamental mathematical object underlying Karmarkar's algorithm and that the good convergence properties of Karmarkar's algorithm arise from good geometric properties of the set of trajectories" [BL89, p. 500]. Renegar showed that the number of iterations performed by path-following interior point methods for linear programming is $O(\sqrt{m}L)$, where m is the number of inequality constraints and L is the total bit size of all coefficients [Ren88]. Understanding the precise dependence in the factor L is a key issue towards strong polynomiality. To this extent, Vavasis and Ye showed that the number of iterations can be bounded by $O(\sqrt{m}\log\chi)$, where χ is a condition number of the constraint matrix [VY96]. This approach has recently been refined by Dadush et al. in [DHNV20] where χ is replaced by a scale-invariant condition number. However, this does not lead to algorithms whose number of iterations is bounded by a polynomial that only depends on the dimensions of the constraint matrix.

Dedieu and Shub proposed a more geometric approach to the complexity of linear programming, considering the total curvature as a complexity measure of the central path which is independent of the type of path-following algorithm. They conjectured that the total curvature of the central path is linearly bounded in the dimension of the ambient space [DS05]. Subsequently, Dedieu, Malajovich, and Shub showed that this property is valid in an average sense [DMS05]. However, Deza, Terlaky, and Zinchenko provided a counter-example by constructing a redundant Klee–Minty cube [DTZ09]. This led them to propose a revised conjecture, the "continuous analogue of the Hirsch conjecture" which says that the total curvature of the central path is linearly bounded in the number of constraints.

Contributions. We show that a significant class of polynomial time interior point methods, namely the primal-dual path-following methods with respect to a log-barrier function, are not strongly polynomial. At the same time, we disprove the conjecture of Deza, Terlaky, and Zinchenko.

To this end, given $r \geqslant 1$, we consider the linear program

$$\mathbf{LW}_r(t) \qquad \begin{aligned} \text{minimize} \quad & x_1 \\ \text{subject to} \quad & x_1 \leqslant t^2, \\ & x_2 \leqslant t, \\ & \left.\begin{aligned} & x_{2j+1} \leqslant t\, x_{2j-1}\,,\ x_{2j+1} \leqslant t\, x_{2j} \\ & x_{2j+2} \leqslant t^{1-1/2^j}(x_{2j-1} + x_{2j}) \end{aligned}\right] \ 1 \leqslant j < r, \\ & x_{2r-1} \geqslant 0\,,\ x_{2r} \geqslant 0, \end{aligned}$$

which depends on a parameter $t > 1$. The linear program $\mathbf{LW}_r(t)$ has $2r$ variables and $3r + 1$ constraints. The notation $\mathbf{LW}_r(t)$ for these linear programs refers to the "long and winding" nature of their central paths.

Our first main result provides a lower bound for the iteration-complexity of interior point methods which is exponential in the number of variables and constraints. This result is based on the assumption that the iterations stay in the so-called "wide" neighborhood of the central path. Relevant examples of such methods include short- and long-step methods, such as those of Kojima, Mizuno, and Yoshise [KMY89a, KMY89b] and Monteiro and Adler [MA89], as well as predictor-corrector methods, such as those of Mizuno, Todd, and Ye [MTY93] and Vavasis and

Ye [VY96]. We refer to Section 3 for a thorough discussion.

Theorem A (see Corollary 32). *The number of iterations of any primal-dual path-following interior point algorithm with a log-barrier function which iterates in the wide neighborhood of the central path is exponential in r on the linear programs $\mathbf{LW}_r(t)$, provided that $t > 1$ is sufficiently large.*

Our second main result disproves the conjecture of [DTZ09].

Theorem B (see Theorem 25). *The total curvature of the central path of the linear programs $\mathbf{LW}_r(t)$ is exponential in r, provided that $t > 1$ is sufficiently large.*

The proofs of these theorems rely on tropical geometry. The latter can be seen as the (algebraic) geometry on the tropical (max-plus) semifield $(\mathbb{T}, \oplus, \odot)$ where the set $\mathbb{T} = \mathbb{R} \cup \{-\infty\}$ is endowed with the operations $a \oplus b = \max(a, b)$ and $a \odot b = a + b$. A tropical variety can be obtained as the limit at infinity of a sequence of classical algebraic varieties depending on one real parameter t and drawn on logarithmic paper, with t as the logarithmic base. This process is known as Maslov's *dequantization* [Lit07], or Viro's method [Vir01]. It can be traced back to the work of Bergman [Ber71]. In a way, dequantization yields a piecewise linear image of classical algebraic geometry. For instance, the log-limit of an algebraic curve is a piecewise linear curve. This approach has been applied, in particular, by Itenberg and Viro to construct counterexample to Ragsdale's conjecture [IV96], and by Mikhalkin to compute enumerative invariants of curves [Mik05]. See [IMS07, MS15] for an introduction.

The paper is organized as follows. In Section 2, we provide an introductory description of the tropical approach to the complexity of interior point methods in linear programming. Section 3 provides background on the central path and path-following interior point methods. In Section 4, we deal with the relation between the analogues of polyhedra in the tropical semifield, polyhedra over non-archimedean fields and log-limits of parametric families of real polyhedra. While some of the results in this section are already known, we establish explicit metric estimates on the convergence of log-limits of polyhedra. Section 5 deals with the tropical central path, defined as the log-limit of the central paths of a parametric family of linear programs, like $(\mathbf{LW}_r(t))_{t>0}$. We show that it is a piecewise linear curve with an explicit geometric characterization, in terms of a tropical analogue of the barycenter of a sublevel set of the feasible set induced by the duality gap; see Section 5.1. Interestingly, it turns out that the tropical central path does not depend on the external representation of the feasible set. This is in stark contrast with the classical case; see [DTZ09] for an example. We study the convergence properties of the classical central path to the tropical one in Section 5.2. In Section 5.3, we show that, in the case of the linear programs $\mathbf{LW}_r(t)$, the tropical central path has an inductive construction, leading to a self-similar pattern, resulting in a number of linear pieces exponential in r.

An interest of tropical geometry is to have a combinatorial nature and yet to retain a lot of information about classical objects. In this vein, we show in Section 6 that, when t is specialized to a suitably large real value, the total curvature of the central path of the linear program $\mathbf{LW}_r(t)$ is bounded below by a combinatorial curvature of the tropical central path. This leads to the lower bound of Theorem B. Similarly, we show in Section 7 that the number of iterations performed by interior point methods is bounded from below by the number of tropical segments forming the tropical central path, which yields Theorem A. We provide an explicit lower bound for the value of the parameter t. For the linear programs $\mathbf{LW}_r(t)$ that bound is doubly exponential

in r. This makes the bit length of t exponential in r, which is consistent with the fact that interior point methods are polynomial in the bit model.

Related work. The redundant Klee–Minty cube of [DTZ09] and the "snake" in [DTZ08] are instances which show that the total curvature of the central path can be in $\Omega(m)$ for a polytope described by m inequalities. The central path has been studied by Dedieu, Malajovich, and Shub [DMS05] via the multihomogeneous Bézout theorem and by De Loera, Sturmfels, and Vinzant [DLSV12] using matroid theory. These two papers provide an upper bound of $O(n)$ on the total curvature averaged over all regions of an arrangement of hyperplanes in dimension n.

However, to the best of our knowledge, there is no explicit relation between the geometric curvature and the iteration-complexity of interior point methods. Zhao and Stoer [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 *Sonnevend's curvature*, was introduced in [SSZ91]. The tight relationship between the total Sonnevend curvature and the iteration-complexity of interior points methods have been extended to semidefinite and symmetric cone programs [KOT13]. Note that Sonnevend's curvature is different from the *geometric* curvature we study in this paper. We circumvent this difficulty here by showing directly that the geometric curvature *and* the number of iterations are exponential for the family $\mathbf{LW}_r(t)$.

The present work relies on the considerations of amoebas (log-limits) of algebraic and semialgebraic sets; see [EKL06, DY07, RGST05, Ale13] for background. Another ingredient is a construction of Bezem, Nieuwenhuis, and Rodríguez-Carbonell [BNRC08]. Their goal was to show that an algorithm of Butkovič and Zimmermann [BZ06] has exponential running time. We arrive at our family of linear programs by lifting a variant of this construction to a nonarchimedean field of Puiseux series.

The applicability of tropical geometry to the complexity of optimization is not limited to interior point methods. Recent examples include connections between convex optimization and repeated zero-sum games [ABGJ14a, AGS18] or multicriteria optimization [JL20].

Theorem B initially appeared in our preprint [ABGJ14b], which additionally discusses aspects of model theory not included here. The present paper is an extended version of [ABGJ18].

2. An overview of the tropical approach to linear programming. Before we start to prove our main results in this section we informally explain how tropical geometry is relevant for linear programming. To motivate our approach, we consider the following linear program in two variables, x_1 and x_2:

$$\begin{aligned} \text{minimize} \quad & x_1 + t^3 x_2 \\ \text{subject to} \quad & x_1 + x_2 \leqslant 2 \\ & t x_1 \leqslant 1 + t^2 x_2 \\ & t x_2 \leqslant 1 + t^3 x_1 \\ & x_1 \leqslant t^2 x_2 \\ & x_1, x_2 \geqslant 0\,. \end{aligned} \tag{1}$$

Some of the entries of the objective function and the constraints depend on a parameter $t > 1$. By picking any fixed value of t, we arrive at a usual linear program. There is also a way to "keep" t as a parameter, which then allows to study how the optimal value and any optimal solutions of (1) depend on t. The crucial step is to think of (1)

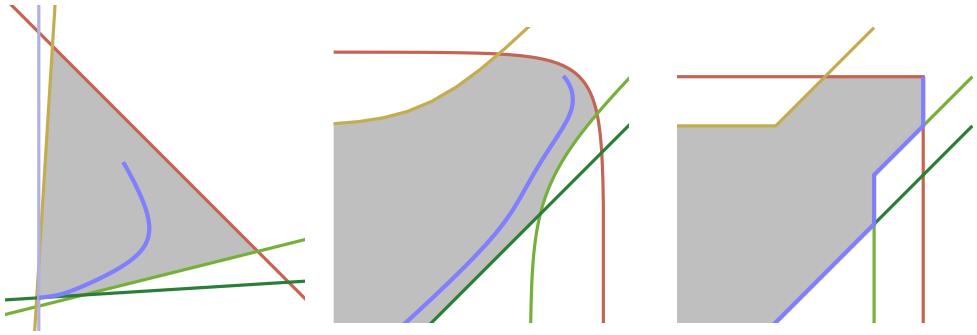


FIG. 1. *Left: the ordinary linear program* (1) *for* $t = 4$*, and its central path (in blue). Center: their logarithmic deformation, i.e., their image by the map* $\log_4$*. Right: the tropicalization* (3)*, and the tropical central path.*

as a linear program over a field $\mathcal{K}$ of real-valued functions of the parameter t. We suppose that the functions of $\mathcal{K}$ are well defined for large enough positive values of t, and totally ordered by the relation $\leqslant$ where for $f, g \in \mathcal{K}$, $f \leqslant g$ if $f(t) \leqslant g(t)$ for all t large enough. An elementary example of such a field is $\mathbb{R}(t)$, the field of rational functions with real coefficients.

In this way, the properties of a parametric family of linear programs, like (1), can be studied by means of the corresponding linear program over the field $\mathcal{K}$. For instance, the latter is feasible if and only if the former is feasible for all large enough values of t. Such ideas go back to van der Waerden [vdW93, §147] and have been exploited in optimization for a while [Jer73].

The tropical approach applies when assuming that, in addition, for all $f \in \mathcal{K}$, the *valuation* of f does exist. The latter is defined as the log-limit

$$\operatorname{val} f := \lim_{t \to \infty} \log_t |f(t)| \in \mathbb{R} \cup \{-\infty\},$$

where $\log_t \cdot := \frac{\log \cdot}{\log t}$ stands for the logarithm with respect to the base $t > 1$. Tropical geometry makes sense over any field satisfying these properties. However, the classical central path is the real positive part of an algebraic curve. Therefore, it is convenient to choose a field that is real-closed, which ensures that the points of the central path are defined over the same field. A wide class of such fields arises from germs of functions definable in a polynomially bounded o-minimal structure [Ale13]. In this paper, we will work with a concrete example of these fields, consisting of absolutely convergent generalized Puiseux series with real exponents [vdDS98]. It has the additional advantage that the map val is surjective.

For all functions $f, g \in \mathcal{K}$ that take nonnegative values for sufficiently large t, we have:

$$\operatorname{val}(f \cdot g) = \operatorname{val}(f) + \operatorname{val}(g) \quad \text{and} \quad \operatorname{val}(f + g) = \max(\operatorname{val}(f), \operatorname{val}(g)) \,. \tag{2}$$

That is to say, the valuation map is a homomorphism from the semifield of nonnegative elements of $\mathcal{K}$ to the tropical semifield $\mathbb{T}$, i.e., to the set $\mathbb{R} \cup \{-\infty\}$, with max as the addition and the usual addition as the multiplication. Observe that, for f and g not necessarily positive, cancellation may occur, which results in the weaker relation $\operatorname{val}(f + g) \leqslant \max(\operatorname{val}(f), \operatorname{val}(g))$. This is the reason why we wrote the constraints in (1) such that all nonzero coefficients are positive.

We now apply a syntactical transformation, called *tropicalization*, to the linear program (1), replacing coefficients by their valuations, the addition by the maximum, and the multiplication by addition. This leads to the following optimization problem:

$$\begin{aligned} &\text{minimize} && \max(x_1, 3 + x_2) \\ &\text{subject to} && \max(x_1, x_2) \leqslant 0 \\ &&& 1 + x_1 \leqslant \max(0, 2 + x_2) \\ &&& 1 + x_2 \leqslant \max(0, 3 + x_1) \\ &&& x_1 \leqslant 2 + x_2 \,, \end{aligned} \tag{3}$$

which is an example of a *tropical linear program*. A few more remarks are in order. First, the valuation map restricted to nonnegative functions is order preserving, whence it makes sense that the tropicalization of a minimization problem is again a minimization problem. Second, the zero function has valuation $-\infty$ so that the nonnegativity constraints $x_1, x_2 \geqslant 0$, become $x_1, x_2 \geqslant -\infty$ in the tropicalization. These are redundant and hence omitted.

What makes a tropicalized linear program useful is that its connection with the original optimization problem is more than merely syntactical. First, given a polyhedron $\mathcal{P}$ over $\mathcal{K}$, like the feasible set of (1), we can consider the ordinary polyhedron $\mathcal{P}(t)$ (over $\mathbb{R}$), obtained by specializing t to a suitably large positive value. We will see that if $\mathcal{P}$ is included in the nonnegative orthant, then the log-limit of $\mathcal{P}(t)$, i.e., the limit in the sense of Hausdorff metric of the images of $\mathcal{P}(t)$ under the map $\log_t$ (understood entrywise) does exist and coincides with the image of $\mathcal{P}$ by the valuation (understood entrywise); see Theorem 9 and Remark 10 below. Moreover, under certain genericity conditions, which hold in our example, this Hausdorff limit is precisely the set obtained by the syntactical procedure, see [ABGJ15, AGS20]. This entails for instance that the original program (1) is feasible for all t large enough if and only if the tropicalized version is feasible. Moreover, this correspondence also applies to optimal values and optimal solutions. In fact, the entire simplex method can be tropicalized [ABGJ15].

In the present work we employ the tropical approach to analyze the interior point method of linear programming, where we restrict our attention to the standard logarithmic barrier function. The idea is to take an ordinary linear program, its constraints and the central path, through the process of logarithmic deformation to the tropical limit. We apply the same machinery, i.e., studying the image under the valuation map, or the log-limit, of the central path. We show in Theorem 15 and Theorem 19 that this image and the log-limit coincide with a piecewise linear curve that we call the tropical central path. Moreover, the tropical central path has a simple geometric characterization: for instance, in the case of (1), the tropical central path is the map $\lambda \mapsto x(\lambda)$ where $x(\lambda)$ is the maximal point of the feasible set of (3) such that the tropical objective function $\max(x_1, 3+x_2)$ does not exceed λ. As shown in Figure 1, this geometric characterization entails that the tropical central path may stay in the boundary of the tropical feasible set. Moreover, for sufficiently large t, the $\log_t$ image of the classical central path is close to its limit. In other words, interior point methods are no longer interior in the tropical limit. Since the combinatorics of the boundary of the tropical feasible set can be as complex as the boundary of an ordinary polyhedron, all these ingredients lead to the lower bounds in Theorem A and Theorem B.

We illustrate the approach in Figure 2. The latter provides the successive iterations of the predictor-corrector interior point method of Mizuno, Todd, and Ye [MTY93],

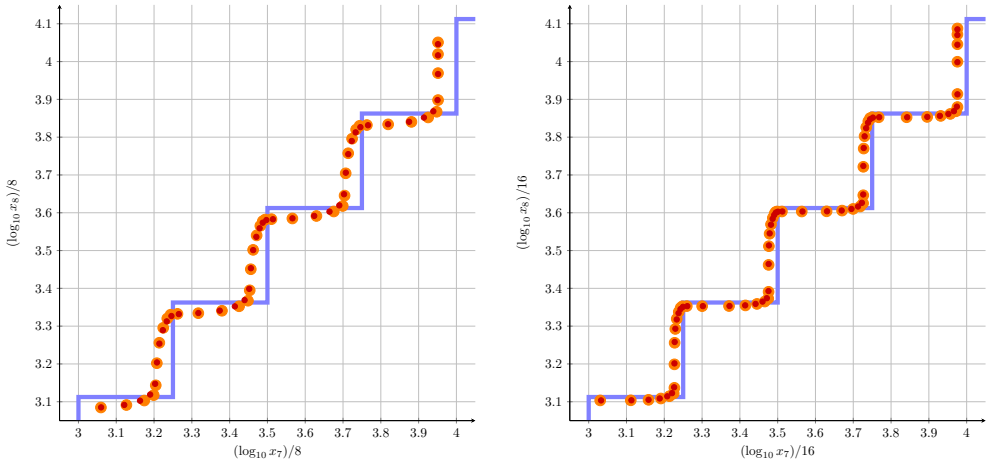


FIG. 2. *The successive iterations performed by the predictor-corrector interior point method of [MTY93] on the linear program* $\mathbf{LW}_4(t)$ *to reduce the duality measure* μ *from* t^2 *to* 1*, when* t *is equal to* 10^8 *(left) and* 10^{16} *(right). The points depict the projection of the iterations on the last two coordinates* (x_7, x_8) *in logarithmic scale (where the logarithm is taken in base* t*). Orange and red points respectively correspond to prediction and corrections steps; see Section 3 for background on predictor-corrector methods. The tropical central path is depicted in blue.*

on the linear program $\mathbf{LW}_4(t)$, when t is equal to 10^8 or 10^{16}. We use the implementation of [BEM18] which performs computation using arbitrary precision floating point numbers. As expected from Theorem 19, we observe that the tropical central path, depicted in blue, is an accurate prediction of the trajectory of the algorithm (orange/red points).

3. The primal-dual central path and its neighborhood. In this section, we recall the definition of the central path and introduce the notions related to path-following interior point methods which we will use in the rest of the paper.

In what follows, we consider a linear program of the form

$$\mathrm{LP}(A,b,c) \qquad \text{minimize} \quad \langle c, x\rangle \quad \text{subject to} \quad Ax + w = b\,,\ (x,w) \in \mathbb{R}^{n+m}_+,$$

in which the slack variables w are explicit. Here and below, A is an $m \times n$ matrix, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and we denote by $\langle \cdot, \cdot \rangle$ the standard scalar product and by $\mathbb{R}_+$ the set of nonnegative reals. The dual linear program takes a similar form

$$\mathrm{DualLP}(A,b,c) \qquad \text{maximize} \quad \langle -b, y\rangle \quad \text{subject to} \quad s - A^\top y = c\,,\ (s,y) \in \mathbb{R}^{n+m}_+\,,$$

where $\cdot^\top$ denotes the transposition. For the sake of brevity, we set $N \coloneqq n + m$ to represent the total number of variables in both linear programs. Let $\mathcal{F}^\circ$ be the set of strictly feasible primal-dual elements, i.e.,

$$\mathcal{F}^\circ \coloneqq \left\{ z = (x,w,s,y) > 0 \colon Ax + w = b\,,\ s - A^\top y = c \right\},$$

which we assume to be nonempty. In this situation, for any given $\mu > 0$, the system

of equations and inequalities

$$\begin{aligned} Ax + w &= b, \\ s - A^\top y &= c, \\ \begin{pmatrix} xs \\ wy \end{pmatrix} &= \mu e, \\ x, w, y, s &> 0 \end{aligned} \tag{4}$$

is known to have a unique solution $(x^\mu, w^\mu, s^\mu, y^\mu) \in \mathbb{R}^{2N}$; here e stands for the all-1-vector in $\mathbb{R}^N$; further, xs and wy denote the Hadamard products of x by s and w by y, respectively. The *central path* of the dual pair $\mathrm{LP}(A, b, c)$ and $\mathrm{DualLP}(A, b, c)$ of linear programs is defined as the function which maps $\mu > 0$ to the point $(x^\mu, w^\mu, s^\mu, y^\mu)$. The latter shall be referred to as the *point of the central path with parameter* μ. The equality constraints in (4) define a real algebraic curve, the *central curve* of the dual pair of linear programs, which has been studied in [BL89] and [DLSV12]. The central curve is the Zariski closure of the central path.

The *primal* and *dual central paths* are defined as the projections of the central path onto the (x, w)- and (s, y)-coordinates, respectively. Equivalently, given $\mu > 0$, the points (x^μ, w^μ) and (s^μ, y^μ) on the primal and dual central paths can be defined as the unique optimal solutions of the following pair of *logarithmic barrier* problems:

$$\begin{aligned} \text{minimize} \quad & \langle c, x \rangle - \mu\Big(\textstyle\sum_{j=1}^n \log(x_j) + \sum_{i=1}^m \log(w_i)\Big) \\ \text{subject to} \quad & Ax + w = b\,,\ x > 0\,,\ w > 0\,, \end{aligned}$$

and

$$\begin{aligned} \text{maximize} \quad & \mu\Big(\textstyle\sum_{j=1}^n \log(s_j) + \sum_{i=1}^m \log(y_i)\Big) + \langle -b, y \rangle \\ \text{subject to} \quad & s - A^\top y = c\,,\ s > 0\,,\ y > 0\,. \end{aligned}$$

The uniqueness of the optimal solutions follows from the fact that the objective functions are strictly convex and concave, respectively. The equivalence to (4) results from the optimality conditions of the logarithmic barrier problems. The main property of the central path is that the sequences (x^μ, w^μ) and (s^μ, y^μ) converge to optimal solutions (x^*, w^*) and (s^*, y^*) of the linear programs $\mathrm{LP}(A, b, c)$ and $\mathrm{DualLP}(A, b, c)$ when μ tends to 0.

The *duality measure* $\bar{\mu}(z)$ of an arbitrary point $z = (x, w, s, y) \in \mathbb{R}^{2N}_+$ is defined by

$$\bar{\mu}(z) \coloneqq \frac{1}{N}\big(\langle x, s \rangle + \langle w, y \rangle\big)\,. \tag{5}$$

With this notation, observe that the point z belongs to the central path if and only if we have

$$\begin{pmatrix} xs \\ wy \end{pmatrix} = \bar{\mu}(z) e\,. \tag{6}$$

In other words, the difference $\binom{xs}{wy} - \bar{\mu}(z)e$ indicates how far the point $z = (x, w, s, y)$ is from the central path. This leads to introducing the neighborhood

$$\mathcal{N}_\theta \coloneqq \left\{ z \in \mathcal{F}^\circ \colon \left\| \begin{pmatrix} xs \\ wy \end{pmatrix} - \bar{\mu}(z)e \right\| \leqslant \theta \bar{\mu}(z) \right\} \tag{7}$$

of the central path by bounding some norm of the deviation in terms of a precision parameter $0 < \theta < 1$. Clearly, this neighborhood depends on the choice of the norm $\|\cdot\|$. In the context of interior point methods, common choices include the ℓ_2-norm and the ℓ_∞-norm. However, here we focus on the *wide neighborhood*

$$\mathcal{N}_\theta^{-\infty} \coloneqq \left\{ z \in \mathcal{F}^\circ \colon \begin{pmatrix} xs \\ wy \end{pmatrix} \geqslant (1-\theta)\bar{\mu}(z)e \right\}. \tag{8}$$

This arises from replacing $\|\cdot\|$ in (7) by the *one-sided ℓ_∞-norm*. The latter is the map sending a vector v to $\max(0, \max_i(-v_i))$. This is a weak norm in the sense of [PT14] (it is positively homogeneous and subadditive), but it vanishes on some nonzero vectors.

Our first observation is that the map $\bar{\mu}$ commutes with affine combinations.

PROPOSITION 1. *Let $z = (x, w, s, y)$ and $z' = (x', w', s', y')$ be two points in $\mathcal{F}$. Then, for all $\alpha \in \mathbb{R}$, we have*

$$\bar{\mu}((1-\alpha)z + \alpha z') = (1-\alpha)\bar{\mu}(z) + \alpha\bar{\mu}(z').$$

Proof. We set $\Delta z \coloneqq z' - z$ and act similarly for the components, i.e., $\Delta z = (\Delta x, \Delta w, \Delta s, \Delta y)$. Since $z, z' \in \mathcal{F}$, we have $A\Delta x + \Delta w = 0$ and $\Delta s - A^\top \Delta y = 0$. Employing these equalities, it can be verified that $\langle \Delta x, \Delta s\rangle + \langle \Delta w, \Delta y\rangle = 0$. Therefore, the function $\alpha \mapsto \bar{\mu}(z + \alpha\Delta z)$ is affine, which completes the proof. ∎

Interior point methods follow the central path by computing a sequence of points in a prescribed neighborhood $\mathcal{N}_\theta$ of the central path, in such a way that the duality measure decreases. Here, we do not precisely specify $\mathcal{N}_\theta$, but we only assume that it arises from the choice of some weak norm $\|\cdot\|$. The basic step of the algorithm can be summarized as follows. At iteration k, given a current point $z^k = (x^k, w^k, s^k, y^k) \in \mathcal{N}_\theta$ with duality measure $\mu^k = \bar{\mu}(z^k)$, and a positive parameter $\mu < \mu^k$, the algorithm aims at solving the system (4) up to a small error in order to get an approximation of the point of the central path with parameter μ. To this end, it starts from the point z^k and exploits the Newton direction $\Delta z = (\Delta x, \Delta w, \Delta s, \Delta y)$, which satisfies

$$\begin{aligned} A\Delta x + \Delta w &= 0, \\ \Delta s - A^\top \Delta y &= 0, \\ \begin{pmatrix} x^k \Delta s \\ w^k \Delta y \end{pmatrix} + \begin{pmatrix} \Delta x\, s^k \\ \Delta w\, y^k \end{pmatrix} &= \mu e - \begin{pmatrix} x^k s^k \\ w^k y^k \end{pmatrix}. \end{aligned} \tag{9}$$

Then, the algorithm follows the direction Δz and iterates to a point of the form

$$z(\alpha) \coloneqq z^k + \alpha\Delta z,$$

where $0 < \alpha \leqslant 1$. The correctness and the convergence of the approach are based on the conditions that, first, the point $z^{k+1} \coloneqq z(\alpha)$ still belongs to the neighborhood $\mathcal{N}_\theta$ and, second, the ratio of the new value $\mu^{k+1} \coloneqq \bar{\mu}(z(\alpha))$ of the duality measure with μ^k is sufficiently small. The following lemma shows that, in fact, the whole line segment between z^k and $z(\alpha)$ is contained in $\mathcal{N}_\theta$.

LEMMA 2. *If z^k and $z(\alpha)$ are contained in $\mathcal{N}_\theta$, then $z(\beta) \in \mathcal{N}_\theta$ for all $\beta \in [0, \alpha]$.*

Proof. The point $z(\beta)$ lies in $\mathcal{F}^\circ$ since the latter set is convex. We use the notation $z(\beta) = (x(\beta), w(\beta), s(\beta), y(\beta))$. Using the last equality in (9), we can write

$$\left(\begin{smallmatrix} x(\beta)s(\beta) \\ w(\beta)y(\beta) \end{smallmatrix}\right) = \beta\mu e + (1-\beta)\left(\begin{smallmatrix} x^k s^k \\ w^k y^k \end{smallmatrix}\right) + \beta^2 \left(\begin{smallmatrix} \Delta x \Delta s \\ \Delta w \Delta y \end{smallmatrix}\right).$$

The first two equalities in (9) entail $\langle \Delta x, \Delta s\rangle + \langle \Delta w, \Delta y\rangle = 0$. Exploiting the last equality, we get $\bar{\mu}(z(1)) = \mu$. Then, it follows from Proposition 1 that $\bar{\mu}(z(\beta)) = (1-\beta)\mu^k + \beta\mu$. We deduce that

$$\begin{pmatrix} x(\beta)s(\beta) \\ w(\beta)y(\beta)\end{pmatrix} - \bar{\mu}(z(\beta))e = (1-\beta)\Big[\begin{pmatrix} x^k s^k \\ w^k y^k\end{pmatrix} - \mu^k e\Big] + \beta^2 \begin{pmatrix} \Delta x \Delta s \\ \Delta w \Delta y\end{pmatrix}.$$

The same relation holds when $\beta = \alpha$. In this way, we eliminate the term $\begin{pmatrix} \Delta x \Delta s \\ \Delta w \Delta y\end{pmatrix}$ to write

$$\begin{pmatrix} x(\beta)s(\beta) \\ w(\beta)y(\beta)\end{pmatrix} - \bar{\mu}(z(\beta))e = \left((1-\beta) - \frac{\beta^2}{\alpha^2}(1-\alpha)\right)\Big[\begin{pmatrix} x^k s^k \\ w^k y^k\end{pmatrix} - \mu^k e\Big] + \frac{\beta^2}{\alpha^2}\Big[\begin{pmatrix} x(\alpha)s(\alpha) \\ w(\alpha)y(\alpha)\end{pmatrix} - \bar{\mu}(z(\alpha))e\Big].$$

Since $\beta \leqslant \alpha \leqslant 1$, the term $(1-\beta) - \frac{\beta^2}{\alpha^2}(1-\alpha)$ is nonnegative. Using the fact that z^k and $z(\alpha)$ belong to $\mathcal{N}_\theta$, and using the subadditivity and positive homogeneity of the weak norm $\|\cdot\|$, we deduce that

$$\begin{aligned}\left\|\begin{pmatrix} x(\beta)s(\beta) \\ w(\beta)y(\beta)\end{pmatrix} - \bar{\mu}(z(\beta))e\right\| &\leqslant \left((1-\beta) - \frac{\beta^2}{\alpha^2}(1-\alpha)\right)\theta\mu^k + \frac{\beta^2}{\alpha^2}\theta\bar{\mu}(z(\alpha)) \\ &= \theta\left((1-\beta)\mu^k + \frac{\beta^2}{\alpha}\mu\right) \\ &\leqslant \theta\big((1-\beta)\mu^k + \beta\mu\big) = \theta\bar{\mu}(z(\beta)) \quad \text{as } \beta \leqslant \alpha\,. \qquad \square\end{aligned}$$

The implementation of the basic iteration step which we have previously described varies from one interior point method to another. In particular, there exist several strategies for the choice of the neighborhood, the parameter μ with respect to the current value of the duality measure μ^k, and the step length α, in order to achieve a polynomial time complexity. Let us describe the main ones in more detail. Considering the large variety of existing path-following interior point methods in the literature, we stick to the classification of [Wri97, Chapter 5], to which we refer the reader for a complete account of the topic.

Short-step interior point methods, such as in [KMY89a, MA89], use an ℓ_2-neighborhood of prescribed size θ, set μ to $\sigma\mu^k$ where $\sigma < 1$ is constant throughout the method (chosen in a careful way to ensure the convergence), and set α to 1. In contrast, long-step interior point methods, such as in [KMY89b], exploit the wider neighborhood $\mathcal{N}_\theta^{-\infty}$, allow more freedom for the choice of μ at every iteration (μ is set to $\sigma\mu^k$ where $\sigma < 1$ is chosen in prescribed interval $[\sigma^{\min}, \sigma^{\max}]$), and take $\alpha \in [0,1]$ as large as possible to ensure that $z(\alpha) \in \mathcal{N}_\theta^{-\infty}$. Another important class of methods, the so-called *predictor-corrector* ones, make use of two nested ℓ_2-neighborhoods $\mathcal{N}_{\theta'}$ and $\mathcal{N}_\theta$ ($\theta' < \theta$), and alternate between predictor and corrector steps. In the former, μ is optimistically set to 0 (the duality measure of optimal solutions), while α is chosen as the largest value in $[0,1]$ such that $z(\alpha) \in \mathcal{N}_\theta$. The next corrector step aims at "centering" the trajectory by performing one Newton step in the direction of the point of the central path with parameter $\mu^{k+1} = \bar{\mu}(z(\alpha))$. This means that the duality measure is kept to μ^{k+1}, and the step length is set to 1. A careful choice of θ depending on θ' ensures that we obtain in this way a point in the narrower neighborhood $\mathcal{N}_{\theta'}$.

The predictor-corrector scheme, initially introduced in [MTY93], inspired several subsequent articles. As we mentioned in the introduction, Vavasis and Ye took a step towards a strongly polynomial complexity by arriving at an iteration-complexity upper bound that depends only on a condition number of the matrix A [VY96]. Their technique was later refined into more practical algorithms; see [MMT98, MT03, KT13,

DHNV20]. The difference between these methods and the original predictor-corrector method is that they sometimes exploit a direction other than the Newton one, called the *layered least squares direction*. However, in such iterations, the step length α is always chosen so that for all $0 \leqslant \beta \leqslant \alpha$, the point $z(\beta)$ lies in the neighborhood $\mathcal{N}_\theta$ (see [VY96, Theorem 9]).

As a consequence of Lemma 2 and the previous discussion, the aforementioned interior point methods all share the property of describing a piecewise linear trajectory entirely included in a certain neighborhood $\mathcal{N}_\theta^{-\infty}$ of the central path, where θ is a prescribed value.[1] We stress that this property is the only assumption made in our complexity result, Theorem A, on interior point methods. More formally, this trajectory is a *polygonal curve* in $\mathbb{R}^{2N}$, i.e., a union of finitely many segments $[z^0, z^1], [z^1, z^2], \ldots, [z^{p-1}, z^p]$. Since polygonal curves play an important role in this paper, we introduce some terminology. We say that a polygonal curve is *supported* by the vectors $v^1, \ldots, v^p$ when the latter correspond to the direction vectors of the successive segments $[z^0, z^1], [z^1, z^2], \ldots, [z^{p-1}, z^p]$. In the case where we equip the curve with an orientation, we assume that the direction vectors are oriented consistently.

4. Ingredients from tropical geometry. Tropical geometry provides a combinatorial approach to studying algebraic varieties defined over a field with a non-Archimedean valuation. To deal with optimization issues, we need some valued field which is ordered. We restrict our attention to one such field, which is particularly convenient for our application, to keep our exposition elementary. Our field of choice, which we denote as $\mathbb{K}$, are the absolutely convergent generalized real Puiseux series. Here "generalized" means that we allow arbitrary real numbers as exponents as in [Mar10]. Note that the ordinary Puiseux series have value group $\mathbb{Q}$, leading to restrictions which are artificial from a tropical perspective. In some sense, $\mathbb{K}$ is the "simplest" real closed valued field for which we can obtain our results.

4.1. Fields of real Puiseux series and Puiseux polyhedra. The field $\mathbb{K}$ of *absolutely convergent generalized real Puiseux series* consists of elements of the form

$$\boldsymbol{f} = \sum_{\alpha \in \mathbb{R}} a_\alpha t^\alpha \,, \tag{10}$$

where $a_\alpha \in \mathbb{R}$ for all α, and such that (i) the support $\{\alpha \in \mathbb{R} \colon a_\alpha \neq 0\}$ either is finite or has $-\infty$ as the only accumulation point, and (ii) there exists $\rho > 0$ such that the series absolutely converges for all $t > \rho$. Note that the null series is obtained by taking an empty support. When $\boldsymbol{f} \neq 0$, the first requirement ensures that the support has a greatest element $\alpha_0 \in \mathbb{R}$. We say that the element $\boldsymbol{f}$ is *positive* when the associated coefficient a_{α_0} is positive. This extends to a total ordering of $\mathbb{K}$, defined by $\boldsymbol{f} \leqslant \boldsymbol{g}$ if $\boldsymbol{g} - \boldsymbol{f}$ is the null series or positive. Equivalently, the relation $\boldsymbol{f} \leqslant \boldsymbol{g}$ holds if and only if $\boldsymbol{f}(t) \leqslant \boldsymbol{g}(t)$ for all sufficiently large t. We write $\mathbb{K}_+$ for the set of nonnegative elements of $\mathbb{K}$.

The *valuation* map $\mathrm{val} \colon \mathbb{K} \to \mathbb{R} \cup \{-\infty\}$ is given as follows. For $\boldsymbol{f} \in \mathbb{K}$, the valuation $\mathrm{val}(\boldsymbol{f})$ is defined as the greatest element α_0 of the support of $\boldsymbol{f}$ if $\boldsymbol{f} \neq 0$, and $-\infty$ otherwise. Recall that we denote $\log_t(\cdot) := \frac{\log(\cdot)}{\log t}$. With this notation, we have

$$\mathrm{val}(\boldsymbol{f}) = \lim_{t \to +\infty} \log_t |\boldsymbol{f}(t)| \,,$$

[1] We point out that the ℓ_2-neighborhood with size θ is obviously contained in $\mathcal{N}_\theta^{-\infty}$.

with the convention $\log_t 0 = -\infty$. Observe that, for all $\boldsymbol{f}, \boldsymbol{g} \in \mathbb{K}$, this yields

$$\operatorname{val}(\boldsymbol{f}+\boldsymbol{g}) \leqslant \max(\operatorname{val}(\boldsymbol{f}), \operatorname{val}(\boldsymbol{g})) \quad \text{and} \quad \operatorname{val}(\boldsymbol{f}\boldsymbol{g}) = \operatorname{val}(\boldsymbol{f}) + \operatorname{val}(\boldsymbol{g})\,. \tag{11}$$

The inequality for the valuation of the sum turns into an equality when the leading terms in the series $\boldsymbol{f}$ and $\boldsymbol{g}$ do not cancel. In particular, this is the case when $\boldsymbol{f}$ and $\boldsymbol{g}$ belong to $\mathbb{K}_+$.

We point out that $\mathbb{K}$ actually agrees with the field of generalized Dirichlet series originally considered by Hardy and Riesz [HR15]. Classical Dirichlet series can be written as $\sum_k a_k k^s$, and these are obtained from (10) by substituting $t = \exp(s)$ and $\alpha_k = \log k$. It follows from results of van den Dries and Speissegger [vdDS98] that the field $\mathbb{K}$ is real closed. The interest in such fields comes from *Tarski's principle*, which says that every real closed field has the same first-order properties as the reals.

As a consequence of the previous fact, we can define polyhedra over Puiseux series as usual. In more detail, given $d \geqslant 1$, a *(Puiseux) polyhedron* is a set of the form

$$\boldsymbol{\mathcal{P}} = \{\boldsymbol{x} \in \mathbb{K}^d \colon \boldsymbol{A}\boldsymbol{x} \leqslant \boldsymbol{b}\}\,, \tag{12}$$

where $\boldsymbol{A} \in \mathbb{K}^{p\times d}$, $\boldsymbol{b} \in \mathbb{K}^p$ (with $p \geqslant 0$), and $\leqslant$ stands for the partial order over $\mathbb{K}^p$. By Tarski's principle, Puiseux polyhedra have the same (first-order) properties as their analogues over $\mathbb{R}$. In particular, the Minkowski–Weyl theorem applies, so that every Puiseux polyhedron admits an internal representation by means of a finite set of points and rays in $\mathbb{K}^d$.

Using a field of convergent series allows us to think of Puiseux polyhedra as parametric families of ordinary polyhedra. Indeed, to any Puiseux polyhedron $\boldsymbol{\mathcal{P}}$ of the form (12), we associate the family of polyhedra $\boldsymbol{\mathcal{P}}(t) \subset \mathbb{R}^d$, defined for t large enough,

$$\boldsymbol{\mathcal{P}}(t) := \{x \in \mathbb{R}^d \colon \boldsymbol{A}(t)x \leqslant \boldsymbol{b}(t)\}\,.$$

The next proposition implies in particular that the family of polyhedra $\boldsymbol{\mathcal{P}}(t)$ is independent of the choice of the external representation of $\boldsymbol{\mathcal{P}}$.

Proposition 3. *Suppose that the Puiseux polyhedron $\boldsymbol{\mathcal{P}}$ is the Minkowski sum of the convex hull of vectors $\boldsymbol{u}^1, \dots, \boldsymbol{u}^q \in \mathbb{K}^d$ and of the convex cone generated by vectors $\boldsymbol{v}^1, \dots, \boldsymbol{v}^r \in \mathbb{K}^d$ (here, the notions of convex hull and convex cone are understood over $\mathbb{K}$). Then, for t large enough, the polyhedron $\boldsymbol{\mathcal{P}}(t)$ is the Minkowski sum of the convex hull of vectors $\boldsymbol{u}^1(t), \dots, \boldsymbol{u}^q(t) \in \mathbb{R}^d$ and of the convex cone generated by vectors $\boldsymbol{v}^1(t), \dots, \boldsymbol{v}^r(t) \in \mathbb{R}^d$ (the notions of convex hull and convex cone are now understood over $\mathbb{R}$).*

Proof. Let $\boldsymbol{\mathcal{Q}}(t)$ denote the Minkowski sum of the convex hull of vectors $\boldsymbol{u}^1(t), \dots, \boldsymbol{u}^q(t)$ and of the convex cone generated by vectors $\boldsymbol{v}^1(t), \dots, \boldsymbol{v}^r(t)$. Since $\boldsymbol{\mathcal{P}}$ contains $\boldsymbol{u}^1, \dots, \boldsymbol{u}^q$, together with the rays $\mathbb{K}_+\boldsymbol{v}^1, \dots, \mathbb{K}_+\boldsymbol{v}^r$, we have $\boldsymbol{A}\boldsymbol{u}^i \leqslant \boldsymbol{b}$ and $\boldsymbol{A}\boldsymbol{v}^j \leqslant 0$ for all $i \in [q]$ and $j \in [r]$. It follows that $\boldsymbol{A}(t)\boldsymbol{u}^i(t) \leqslant \boldsymbol{b}(t)$ and $\boldsymbol{A}(t)\boldsymbol{v}^j(t) \leqslant 0$ hold for all $i \in [q]$ and $j \in [r]$ and for t large enough. Hence, $\boldsymbol{\mathcal{P}}(t) \supset \boldsymbol{\mathcal{Q}}(t)$ for t large enough.

Let us now consider an extreme point $\boldsymbol{u}$ of $\boldsymbol{\mathcal{P}}$. Then, a characterization of the extreme points of a polyhedron shows that the collection of gradients of the constraints $\boldsymbol{A}_k\boldsymbol{x} \leqslant \boldsymbol{b}_k$, $k \in [p]$, which are active at point $\boldsymbol{x} = \boldsymbol{u}$ constitutes a family of full rank. This property can be expressed in the first-order theory of $\mathbb{K}$. It follows that, for t large enough, the same property holds for the collection of the gradients of the constraints $\boldsymbol{A}_k(t)x \leqslant \boldsymbol{b}_k(t)$, $k \in [p]$, that are active at point $x = \boldsymbol{u}(t)$. Hence $\boldsymbol{u}(t)$ is an extreme point of $\boldsymbol{\mathcal{P}}(t)$, and so $\boldsymbol{u}(t)$ must belong to the set $\{\boldsymbol{u}^1(t), \dots, \boldsymbol{u}^q(t)\}$. A similar argument shows that if $\boldsymbol{v} \in \mathbb{K}^d$ generates an extreme ray of $\boldsymbol{\mathcal{P}}$, the ray

generated by $\boldsymbol{v}(t)$ is extreme in $\boldsymbol{\mathcal{P}}(t)$ for t large enough, and so $\boldsymbol{v}(t) \in \cup_{\ell\in[r]}\mathbb{R}_+\boldsymbol{v}^\ell(t)$. It follows that $\boldsymbol{\mathcal{P}}(t) \subset \boldsymbol{\mathcal{Q}}(t)$ holds for t large enough. ∎

Remark 4. One can show, by arguments of the same nature as in the latter proof, that for all $\boldsymbol{x} \in \mathbb{K}^d$,

$$\boldsymbol{x} \in \boldsymbol{\mathcal{P}} \iff (\boldsymbol{x}(t) \in \boldsymbol{\mathcal{P}}(t) \text{ for } t \text{ large enough}) . \tag{13}$$

Note, however, that the smallest value t_0 such that $\boldsymbol{x}(t) \in \boldsymbol{\mathcal{P}}(t)$ for all $t \geqslant t_0$ cannot be bounded uniformly in $\boldsymbol{x}$.

4.2. Tropical polyhedra. Tropical polyhedra may be informally thought of as the analogues of convex polyhedra over the tropical semifield $\mathbb{T}$. Note that in this semifield, the zero and unit elements are $-\infty$ and 0, respectively. Given $\lambda \in \mathbb{T}\backslash\{-\infty\}$, we shall also denote by $\lambda^{\odot(-1)}$ the inverse of λ for the tropical multiplication, i.e., $\lambda^{\odot(-1)} \coloneqq -\lambda$. The tropical addition and multiplication extend to vectors and matrices in the usual way. More precisely, $A\oplus B \coloneqq (A_{ij}\oplus B_{ij})_{ij}$ and $A\odot B \coloneqq (\bigoplus_k A_{ik}\odot B_{kj})_{ij}$, where A and B are two matrices of appropriate sizes with entries in $\mathbb{T}$. Further, the d-fold Cartesian product $\mathbb{T}^d$ is equipped with the structure of a semimodule, thanks to the tropical multiplication $\lambda \odot v \coloneqq (\lambda \odot v_i)_i$ of a vector v with a scalar λ.

A *tropical halfspace* of $\mathbb{T}^d$ is the set of points $x \in \mathbb{T}^d$ which satisfy one tropical linear (affine) inequality,

$$\max(\alpha_1 + x_1, \dots, \alpha_d + x_d, \beta) \leqslant \max(\alpha'_1 + x_1, \dots, \alpha'_d + x_d, \beta') ,$$

where $\alpha, \alpha' \in \mathbb{T}^d$ and $\beta, \beta' \in \mathbb{T}$. A *tropical polyhedron* is the intersection of finitely many tropical halfspaces. Equivalently, it can be written in the form

$$\{x \in \mathbb{T}^d \colon A \odot x \oplus b \leqslant A' \odot x \oplus b'\},$$

where $A, A' \in \mathbb{T}^{p\times d}$ and $b, b' \in \mathbb{T}^p$ for some $p \geqslant 0$. The tropical semifield $\mathbb{T} = \mathbb{R}\cup\{-\infty\}$ is equipped with the order topology, which gives rise to the product topology on $\mathbb{T}^d$. Tropical halfspaces, and thus tropical polyhedra, are closed in this topology. Note also that the subset topology on $\mathbb{R}^d \subset \mathbb{T}^d$ agrees with the usual Euclidean topology.

An analogue of the Minkowski–Weyl theorem allows for the "interior representation" of a tropical polyhedron $\mathcal{P}$ in terms of linear combinations of points and rays [GK11]. That is to say, there exist finite sets $U, V \subset \mathbb{T}^d$ such that $\mathcal{P}$ is the set of all points of the form

$$\left(\bigoplus_{u\in U} \alpha_u \odot u\right) \oplus \left(\bigoplus_{v\in V} \beta_v \odot v\right), \tag{14}$$

where $\alpha_u, \beta_v \in \mathbb{T}$ and $\bigoplus_{u\in U} \alpha_u$ is equal to the tropical unit, i.e., the real number 0. We shall say that the tropical polyhedron $\mathcal{P}$ is *generated* by the sets U and V. The term $\bigoplus_{u\in U} \alpha_u \odot u$ is a tropical convex combination of the points in U, while $\bigoplus_{v\in V} \beta_v \odot v$ is a tropical linear combination of the vectors in V. These are the tropical analogues of convex and conic hulls, respectively. Indeed, all scalars α_u, β_v are implicitly nonnegative in the tropical sense, i.e., they are greater than or equal to the tropical zero element $-\infty$. We point out that the "tropical polytopes" considered by Develin and Sturmfels [DS04] are obtained by omitting the term $\bigoplus_{u\in U} \alpha_u \odot u$ and requiring the vectors $v \in V$ to have finite coordinates in the representation (14).

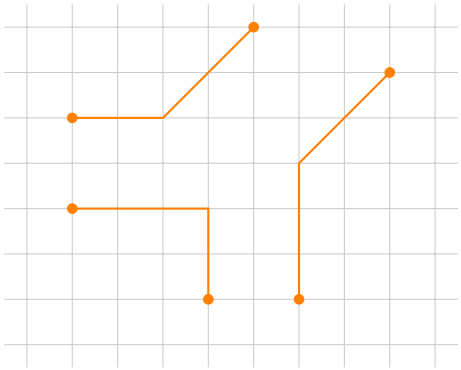


FIG. 3. *The three possible shapes of (generic) tropical segments in dimension* 2.

If $\mathcal{P}$ is a nonempty tropical polyhedron, the supremum $\sup(u,v) = u \oplus v$ with respect to the partial order $\leqslant$ of $\mathbb{T}^d$ of any two points $u, v \in \mathcal{P}$ also belongs to $\mathcal{P}$. If in addition $\mathcal{P}$ is compact, then the supremum of an arbitrary subset of $\mathcal{P}$ is well defined and belongs to $\mathcal{P}$. Consequently, there is a unique element in $\mathcal{P}$ which is the coordinatewise maximum of all elements in $\mathcal{P}$. We call it the *(tropical) barycenter* of $\mathcal{P}$, as it is the mean of $\mathcal{P}$ with respect to the uniform idempotent measure.

The *tropical segment* between the points $u, v \in \mathbb{T}^d$, denoted by $\mathsf{tsegm}(u,v)$, is defined as the set of points of the form $\lambda \odot u \oplus \mu \odot v$ such that $\lambda \oplus \mu = 0$. Equivalently, the set $\mathsf{tsegm}(u,v)$ is the tropical polyhedron generated by the sets $U = \{u, v\}$ and $V = \emptyset$. As illustrated in Figure 3, tropical segments are polygonal curves, and the direction vectors supporting every ordinary segment have their entries in $\{0, \pm 1\}$ [DS04, Proposition 3]. We shall slightly refine this statement in the case where $u \leqslant v$, and propose a characterization of tropical segments between any two such points. To this end, for $K \subset [d]$, we denote by e^K the vector whose kth entry is equal to 1 if $k \in K$, and 0 otherwise.

LEMMA 5. *Let $u, v \in \mathbb{T}^d$ such that $u \leqslant v$. The tropical segment $\mathsf{tsegm}(u,v)$ is a polygonal curve which, when oriented from u to v, consists of ordinary segments supported by direction vectors of the form $e^{K_1}, \dots, e^{K_\ell}$ where $K_1 \subsetneq \cdots \subsetneq K_\ell$ and $\ell \leqslant d$.*

Conversely, any polygonal curve from the point u to v in which the successive ordinary segments are supported by direction vectors of the form $e^{K_1}, \dots, e^{K_\ell}$ where $K_1 \subsetneq \cdots \subsetneq K_\ell$ coincides with the tropical segment $\mathsf{tsegm}(u,v)$.

Proof. We start by dealing with the first part of the statement. Since $u \leqslant v$, the set $\mathsf{tsegm}(u,v)$ is reduced to the set of the points of the form $u \oplus (\mu \odot v)$, where $\mu \leqslant 0$. Let $K(\mu)$ be the set of $i \in [d]$ such that $u_i < \mu + v_i$. When μ ranges from $-\infty$ to 0, $K(\mu)$ takes a finite number of values $K_0 = \emptyset \subsetneq K_1 \subsetneq \cdots \subsetneq K_\ell$, where $\ell \leqslant d$. It is immediate that the ordinary segments constituting the tropical segment $\mathsf{tsegm}(u,v)$ are supported by the vectors $e^{K_1}, \dots, e^{K_\ell}$.

We prove the second part of the statement by induction on ℓ. When $\ell = 1$, the polygonal curve reduces to a line segment $[u,v]$, where $v := u + \nu e^K$ and $\nu \geqslant 0$. If $\mu \leqslant 0$, the combination $u \oplus (\mu \odot v)$ is equal to u when $\mu < -\nu$, and to $u + (\mu + \nu)e^K$ when $\mu \geqslant -\nu$. We deduce that the ordinary segment $[u,v]$ coincides with the tropical segment $\mathsf{tsegm}(u,v)$.

Now, we assume that $\ell > 1$, and we split the polygonal curve, which we denote by

$\mathcal{C}$, into the subparts formed by the $(\ell-1)$ first ordinary segments and the last segment, respectively. Let w be the common vertex of the two subparts. We point out that $u \leqslant w \leqslant v$. Moreover, $v = w + \nu e^{K_\ell}$ for some $\nu \geqslant 0$. By the induction hypothesis, the two subparts respectively coincide with the tropical segments $\mathsf{tsegm}(u, w)$ and $\mathsf{tsegm}(w, v)$. Moreover, w can be written as $u \oplus ((-\nu) \odot v)$. Indeed, for all $i \in [d]$, we have $\big(u \oplus ((-\nu) \odot v)\big)_i = \max(u_i, -\nu + v_i)$. The latter quantity is equal to $\max(u_i, w_i)$ if $i \in K_\ell$, and to $\max_i(u_i, -\nu + w_i)$ otherwise. Besides, we know that $w_i \geqslant u_i$, and that the equality holds when $i \notin K_\ell$, as for all $j < \ell$ the set K_j is included in K_ℓ. We deduce that w belongs to the tropical segment $\mathsf{tsegm}(u, v)$, which implies that $\mathcal{C} = \mathsf{tsegm}(u, w) \cup \mathsf{tsegm}(w, v)$ is included in $\mathsf{tsegm}(u, v)$. Reciprocally, let $x = u \oplus (\mu \odot v)$ be an element of $\mathsf{tsegm}(u, v)$. If $\mu \leqslant -\nu$, note that we have $x = u \oplus \big((\mu + \nu) \odot w\big)$, as $w = u \oplus ((-\nu) \odot v)$. In this case, x belongs to $\mathsf{tsegm}(u, w)$. If $\mu > -\nu$, then $x \geqslant w$. Then, we can write $x = x \oplus w = w \oplus (\mu \odot v)$ as $w \geqslant u$. We deduce that $x \in \mathsf{tsegm}(w, v)$. This shows that $\mathcal{C}$ coincides with $\mathsf{tsegm}(u, v)$. ∎

We now relate tropical polyhedra with their classical analogues over Puiseux series via the valuation map. The fact that sums of nonnegative Puiseux series do not suffer from cancellation translates into the following.

LEMMA 6. *The valuation map is a monotone and surjective semifield homomorphism from $\mathbb{K}_+$ to $\mathbb{T}$.*

Proof. That val is a homomorphism is a consequence of (11) and the subsequent discussion. Monotonicity and surjectivity are straightforward. ∎

This carries over to Puiseux polyhedra in the nonnegative orthant.

PROPOSITION 7. *The image under the valuation map of any Puiseux polyhedron $\boldsymbol{\mathcal{P}} \subset \mathbb{K}_+^d$ is a tropical polyhedron in $\mathbb{T}^d$.*

Proof. Let $\boldsymbol{U}, \boldsymbol{V} \subset \mathbb{K}^d$ be two finite collections of vectors such that $\boldsymbol{\mathcal{P}}$ is the set of combinations of the form

$$\boldsymbol{x} = \sum_{\boldsymbol{u} \in \boldsymbol{U}} \boldsymbol{\alpha_u} \boldsymbol{u} + \sum_{\boldsymbol{v} \in \boldsymbol{V}} \boldsymbol{\beta_v} \boldsymbol{v}\,, \tag{15}$$

where $\boldsymbol{\alpha_u}, \boldsymbol{\beta_v} \in \mathbb{K}_+$ and $\sum_{\boldsymbol{u} \in \boldsymbol{U}} \boldsymbol{\alpha_u} = 1$. Observe that $\boldsymbol{U}$ and $\boldsymbol{V}$ both lie in $\mathbb{K}_+^d$ since $\boldsymbol{\mathcal{P}} \subset \mathbb{K}_+^d$. From Lemma 6, we deduce that $\mathrm{val}(\boldsymbol{\mathcal{P}})$ is contained in the tropical polyhedron $\mathcal{P}$ generated by the sets $U \coloneqq \mathrm{val}(\boldsymbol{U})$ and $V \coloneqq \mathrm{val}(\boldsymbol{V})$. Conversely, any point in $\mathcal{P}$ of the form (14) is the image under the valuation map of

$$\sum_{\boldsymbol{u} \in \boldsymbol{U}} \frac{1}{\boldsymbol{Z}} t^{\alpha_u} \boldsymbol{u} + \sum_{\boldsymbol{v} \in \boldsymbol{V}} t^{\beta_v} \boldsymbol{v}\,,$$

where $\boldsymbol{Z} = \sum_{u \in U} t^{\alpha_u}$ is such that $\mathrm{val}\, \boldsymbol{Z} = 0$. ∎

The special case of Proposition 7 concerning "tropical polytopes" in the sense of [DS04] was already proved by Develin and Yu [DY07, Proposition 2.1]. One can show that, conversely, each tropical polyhedron arises as the image under the valuation map of a polyhedron included in $\mathbb{K}_+^d$; see [ABGJ15, Proposition 2.6].

4.3. Metric properties. In this section, we establish various metric estimates which will be used in the analysis of the central path in section 7. These estimates involve several kinds of metrics. We start with the nonsymmetric metric δ_{F}, defined by

$$\delta_{\mathrm{F}}(x, y) \coloneqq \inf\big\{\rho \geqslant 0 \colon x + \rho e \geqslant y\big\}\,,$$

where $x, y \in \mathbb{T}^d$. Recall that e denotes the all-1-vector. Writing the inequality $x + \rho e \geqslant y$ as $\rho \odot x \geqslant y$ reveals that δ_{F} is the tropical analogue of the *Funk metric* which appears in Hilbert's geometry [PT14]. Equivalently, we can write

$$\delta_{\mathrm{F}}(x, y) = \max(0, \max_k(y_k - x_k)) \,,$$

with the convention $-\infty - (-\infty) = -\infty$. In this way, we observe that δ_{F} is derived from the one-sided ℓ_∞-norm $\|\cdot\|$ which we used to define the wide neighborhood of the central path in (8), i.e., $\delta_{\mathrm{F}}(x, y) = \|x - y\|$. We point out that $\delta_{\mathrm{F}}(x, y) < +\infty$ if and only if the support of x contains the support of y, i.e., $\{k \colon x_k \neq -\infty\} \supset \{k \colon y_k \neq -\infty\}$.

The metric d_∞ induced by the ordinary ℓ_∞-norm is obtained by symmetrizing δ_{F} as follows:

$$d_\infty(x, y) \coloneqq \max(\delta_{\mathrm{F}}(x, y), \delta_{\mathrm{F}}(y, x)) \,.$$

We shall consider another symmetrization of δ_{F}, leading to the affine version of *Hilbert's projective metric*:

$$d_{\mathrm{H}}(x, y) \coloneqq \delta_{\mathrm{F}}(x, y) + \delta_{\mathrm{F}}(y, x) \,.$$

The metric d_{H} was shown in [CGQ04] to be the canonical metric in tropical convexity. For instance, the projection onto a convex set is well defined and is a best approximation in this metric. The relevance of Hilbert's geometry to the study of the central path was already observed by Bayer and Lagarias [BL89]. Notice that $d_{\mathrm{H}}(x, y) < +\infty$ if and only if the supports of the two vectors $x, y \in \mathbb{T}^d$ are identical.

We extend our notation to sets as follows. Given $X, Y \subset \mathbb{T}^d$ we define

$$d_{\mathrm{H}}(X, Y) := \sup_{x \in X} \inf_{y \in Y} d_{\mathrm{H}}(x, y) \quad \text{and} \quad d_\infty(X, Y) := \sup_{x \in X} \inf_{y \in Y} d_\infty(x, y) \,.$$

These are the *directed Hausdorff distances* from X to Y induced by d_{H} and d_∞, respectively.

In order to establish the metric properties of this section, we repeatedly use the following elementary inequalities: if $t > 1$ and $\gamma_1, \dots, \gamma_p \in \mathbb{R}_+$,

(16)
$$\max(\log_t \gamma_1, \dots, \log_t \gamma_p) \leqslant \log_t(\gamma_1 + \cdots + \gamma_p) \leqslant \max(\log_t \gamma_1, \dots, \log_t \gamma_p) + \log_t p \,.$$

We start with a metric estimate over classical and tropical segments. Given $t > 1$ and a vector $x \in \mathbb{R}^d_+$, we denote by $\log_t x$ the vector of $\mathbb{T}^d$ with entries $\log_t x_i$. Moreover, if $X \subset \mathbb{R}^d_+$, $\log_t X$ is short for $\{\log_t x \colon x \in X\}$.

LEMMA 8. *Let $S = [u, v]$ be a segment in $\mathbb{R}^d_+$, and let S^{trop} be the tropical segment between the points $\log_t u$ and $\log_t v$, where $t > 1$. Then*

$$d_\infty(S^{\mathsf{trop}}, \log_t S) \leqslant \log_t 2 \,.$$

Proof. Let $x \coloneqq \lambda \odot (\log_t u) \oplus \mu \odot (\log_t v)$ be a point of the tropical segment S^{trop}, where $\lambda, \mu \in \mathbb{T}$ are such that $\lambda \oplus \mu = 0$. Now the point

$$x' \coloneqq \frac{t^\lambda u + t^\mu v}{t^\lambda + t^\mu}$$

belongs to S. Using (16) and $\lambda \oplus \mu = 0$, we get $0 \leqslant \log_t(t^\lambda + t^\mu) \leqslant \log_t 2$. Similarly, for all $i \in [d]$, we have $x_i \leqslant \log_t(t^\lambda u_i + t^\mu v_i) \leqslant x_i + \log_t 2$. We deduce that $x_i - \log_t 2 \leqslant \log_t x'_i \leqslant x_i + \log_t 2$. As a consequence, $d_\infty(x, \log_t x') \leqslant \log_t 2$. It follows that $d_\infty(S^{\mathsf{trop}}, \log_t S) \leqslant \log_t 2$. ∎

We now focus on estimating the distance between tropical polyhedra and related logarithmic deformations of convex polyhedra. To this end, we consider a Puiseux polyhedron $\boldsymbol{\mathcal{P}}$ included in the nonnegative orthant, as well as the associated parametric family of polyhedra $\boldsymbol{\mathcal{P}}(t)$ over $\mathbb{R}$. The following theorem shows that the tropical polyhedron $\operatorname{val}(\boldsymbol{\mathcal{P}})$ is the log-limit of the polyhedra $\boldsymbol{\mathcal{P}}(t)$, and that the convergence is uniform. This is related to a result of Briec and Horvath, who established in [BH04] a uniform convergence property for a parametric family of convex hulls.

THEOREM 9. *Let $\boldsymbol{\mathcal{P}} \subset \mathbb{K}_+^d$ be a Puiseux polyhedron. Then the sequence $(\log_t \boldsymbol{\mathcal{P}}(t))_t$ of logarithms of real polyhedra converges to the tropical polyhedron $\operatorname{val}(\boldsymbol{\mathcal{P}})$ with respect to the directed Hausdorff distance d_{H}.*

Proof. By Proposition 3, we can find finite sets $\boldsymbol{U}, \boldsymbol{V} \subset \mathbb{K}_+^d$ such that for sufficiently large t, the real polyhedron $\boldsymbol{\mathcal{P}}(t)$ is generated by the sets of points $\boldsymbol{U}(t) \coloneqq \{\boldsymbol{u}(t) \colon \boldsymbol{u} \in \boldsymbol{U}\} \subset \mathbb{R}_+^d$ and rays $\boldsymbol{V}(t) \coloneqq \{\boldsymbol{v}(t) \colon \boldsymbol{v} \in \boldsymbol{V}\} \subset \mathbb{R}_+^d$. Let $\boldsymbol{u} \in \boldsymbol{U}$. If t is large enough, then $\boldsymbol{u}_i(t) = 0$ is equivalent to $\operatorname{val}(\boldsymbol{u})_i = -\infty$ for all $i \in [d]$. Thus $\delta_{\mathrm{F}}(\log_t \boldsymbol{u}(t), \operatorname{val} \boldsymbol{u})$, as well as $\delta_{\mathrm{F}}(\operatorname{val} \boldsymbol{u}, \log_t \boldsymbol{u}(t))$, converge to 0 when $t \to +\infty$. The situation is similar for $\boldsymbol{v}(t)$ and $\operatorname{val} \boldsymbol{v}$ for any $\boldsymbol{v} \in \boldsymbol{V}$.

Moreover, the tropical polyhedron $\mathcal{P} \coloneqq \operatorname{val}(\boldsymbol{\mathcal{P}})$ is generated by the sets $\operatorname{val}(\boldsymbol{U})$ and $\operatorname{val}(\boldsymbol{V})$, as shown in the proof of Proposition 7.

Now consider $x \in \boldsymbol{\mathcal{P}}(t)$. From Carathéodory's theorem we know that there exist subsets $\{\boldsymbol{u}^k(t)\}_{k \in K} \subset \boldsymbol{U}(t)$ and $\{\boldsymbol{v}^\ell(t)\}_{\ell \in L} \subset \boldsymbol{V}(t)$ with $|K| + |L| \leqslant d+1$ such that the point x can be written as

$$x = \sum_{k \in K} \alpha_k \boldsymbol{u}^k(t) + \sum_{\ell \in L} \beta_\ell \boldsymbol{v}^\ell(t)\,,$$

where $\alpha_k, \beta_\ell > 0$ for $k \in K$, $\ell \in L$, and $\sum_{k \in K} \alpha_k = 1$. By using (16), we get that, for all $i \in [d]$,

$$\begin{aligned} (17)\quad & \bigoplus_{k \in K} \big((\log_t \alpha_k) \odot \log_t \boldsymbol{u}_i^k(t)\big) \oplus \bigoplus_{\ell \in L} \big((\log_t \beta_\ell) \odot \log_t \boldsymbol{v}_i^\ell(t)\big) \leqslant \log_t x_i \\ & \leqslant \left[\bigoplus_{k \in K} \big((\log_t \alpha_k) \odot \log_t \boldsymbol{u}_i^k(t)\big) \oplus \bigoplus_{\ell \in L} \big((\log_t \beta_\ell) \odot \log_t \boldsymbol{v}_i^\ell(t)\big)\right] + \log_t(|K| + |L|)\,. \end{aligned}$$

Setting $\gamma \coloneqq \max_{k \in K} \alpha_k$, we have $\frac{1}{|K|} \leqslant \gamma \leqslant 1$. Now we define

$$x' \coloneqq \left(\bigoplus_{k \in K} \alpha'_k \odot u^k\right) \oplus \left(\bigoplus_{\ell \in L} \beta'_\ell \odot v^\ell\right),$$

where $\alpha'_k \coloneqq \log_t(\alpha_k/\gamma)$, $\beta'_\ell \coloneqq \log_t \beta_\ell$, $u^k \coloneqq \operatorname{val}(\boldsymbol{u}^k)$, and $v^\ell \coloneqq \operatorname{val}(\boldsymbol{v}^\ell)$. By the choice of γ we have $\bigoplus_{k \in K} \alpha'_k = 0$, and thus $x' \in \mathcal{P}$. Further, $x_i > 0$ if and only if there exists $k \in K$ such that $\boldsymbol{u}_i^k(t) > 0$ or $\ell \in L$ such that $\boldsymbol{v}_i^\ell(t) > 0$. Provided that t is sufficiently large, this is equivalent to the fact that $u_i^k > -\infty$ for some $k \in K$, or $v_i^\ell > -\infty$ for some $\ell \in L$. This latter property amounts to $x'_i > -\infty$. Consequently, we have $d_{\mathrm{H}}(\log_t x, x') < +\infty$, and we can derive from (17) that

$$\begin{aligned} (18)\quad & x'_i - \max\left(\log_t |K| + \max_{k \in K} \delta_{\mathrm{F}}(\log_t \boldsymbol{u}^k(t), u^k), \max_{\ell \in L} \delta_{\mathrm{F}}(\log_t \boldsymbol{v}^\ell(t), v^\ell)\right) \leqslant \log_t x_i \\ & \leqslant x'_i + \log_t(|K| + |L|) + \max\left(\max_{k \in K} \delta_{\mathrm{F}}(u^k, \log_t \boldsymbol{u}^k(t)), \max_{\ell \in L} \delta_{\mathrm{F}}(v^\ell, \log_t \boldsymbol{v}^\ell(t))\right) \end{aligned}$$

for all $i \in [d]$. Finally, we deduce that

(19)
$$d_{\mathrm{H}}(\log_t \boldsymbol{\mathcal{P}}(t), \mathcal{P}) \leqslant \log_t(d+1) + \max\left(\max_{\boldsymbol{u}\in\boldsymbol{U}} \delta_{\mathrm{F}}(\operatorname{val}\boldsymbol{u}, \log_t \boldsymbol{u}(t)), \max_{\boldsymbol{v}\in\boldsymbol{V}} \delta_{\mathrm{F}}(\operatorname{val}\boldsymbol{v}, \log_t \boldsymbol{v}(t))\right) + \max\left(\log_t(d+1) + \max_{\boldsymbol{u}\in\boldsymbol{U}} \delta_{\mathrm{F}}(\log_t \boldsymbol{u}(t), \operatorname{val}\boldsymbol{u}), \max_{\boldsymbol{v}\in\boldsymbol{V}} \delta_{\mathrm{F}}(\log_t \boldsymbol{v}(t), \operatorname{val}\boldsymbol{v})\right),$$

which tends to 0 when $t \to +\infty$. ∎

Remark 10. For the sake of brevity, we only stated and proved here the one-sided metric estimates which we will use in the proof of our main results, leaving it to the interested reader to derive the symmetrical metric estimates. For instance, the inequality $d_\infty(\log_t S, S^{\mathsf{trop}}) \leqslant \log_t 2$ can be shown by a method similar to that of Lemma 8. Likewise, a variant of the proof of Theorem 9 shows that the reversed Hausdorff distance $d_{\mathrm{H}}(\operatorname{val}(\boldsymbol{\mathcal{P}}), \log_t \boldsymbol{\mathcal{P}}(t))$ tends to zero when $t \to \infty$.

Next we refine the convergence result just obtained by providing a metric estimate in the special case where $\boldsymbol{\mathcal{P}}$ is a polyhedron given by constraints with monomial coefficients. Here a Puiseux series of the form $\pm t^\alpha$ is called *monomial*, with the convention $t^{-\infty} = 0$. Further, a vector or a matrix is *monomial* if all its entries are. For a monomial matrix $\boldsymbol{M} = (\pm t^{\alpha_{ij}})$ of size $d \times d$ we introduce the quantity $\eta(\boldsymbol{M}) > 0$ by letting

$$\eta(\boldsymbol{M}) \coloneqq \min\left\{\eta \colon \sigma, \tau \in \mathrm{Sym}_d,\ \eta = \sum_{i=1}^d \alpha_{i\sigma(i)} - \sum_{i=1}^d \alpha_{i\tau(i)} > 0\right\},$$

where Sym_d stands for the symmetric group over $[d]$. We use the convention $\min \emptyset = +\infty$. Phrased differently, the determinant of $\boldsymbol{M}$ is a Puiseux series with finitely many terms with decreasing exponents, and $\eta(\boldsymbol{M})$ provides a lower bound on the gap between any two successive exponents (if any). This allows us to obtain explicit upper and lower bounds for $\log_t|\det \boldsymbol{M}(t)|$ in terms of $\operatorname{val}(\det \boldsymbol{M})$. Note that these bounds hold without any assumption on the genericity of the matrix $\boldsymbol{M}$.

Lemma 11. *Let $\boldsymbol{M} \in \mathbb{K}^{d\times d}$ be a monomial matrix. Then, for all $t > 1$, we have*

$$\log_t|\det \boldsymbol{M}(t)| \leqslant \operatorname{val}(\det \boldsymbol{M}) + \log_t d!\,,$$

and, if additionally $t \geqslant (d!)^{1/\eta(\boldsymbol{M})}$, then we get

$$\operatorname{val}(\det \boldsymbol{M}) \leqslant \log_t|\det \boldsymbol{M}(t)| + \log_t d!\,.$$

Proof. First, note that the statement is trivial when $\det \boldsymbol{M} = 0$, since in this case, the determinant of $\boldsymbol{M}(t)$ vanishes for all $t > 1$. Now suppose that $\det \boldsymbol{M} \neq 0$. Since $\boldsymbol{M}$ is monomial, every entry of $\boldsymbol{M}$ is of the form $\epsilon_{ij} t^{\alpha_{ij}}$ where $\epsilon_{ij} \in \{\pm 1\}$. Therefore, we obtain

$$\det \boldsymbol{M} = \sum_{k=1}^p c_k t^{\beta_k}\,,$$

where the following conditions are met: (i) each c_k is a nonnull integer, and $\sum_k |c_k| \leqslant d!$; (ii) every β_k is of the form $\sum_{i=1}^d \alpha_{i\sigma(i)}$ for a certain permutation $\sigma \in \mathrm{Sym}_d$; and

(iii) $\beta_1 > \cdots > \beta_p > -\infty$. With this notation, we have $\operatorname{val}(\det \boldsymbol{M}) = \beta_1$, and $\beta_i - \beta_{i+1} \geqslant \eta(\boldsymbol{M})$. Similarly, for all $t > 1$, we have $\det \boldsymbol{M}(t) = \sum_{k=1}^{p} c_k t^{\beta_k}$. This leads to

$$\log_t |\det \boldsymbol{M}(t)| \leqslant \beta_1 + \log_t \left(\sum_{k=1}^{p} |c_k| t^{\beta_k - \beta_1} \right) \leqslant \operatorname{val}(\det \boldsymbol{M}) + \log_t d! \,.$$

Further, provided that $t \geqslant (d!)^{1/\eta(\boldsymbol{M})}$, we have

$$\sum_{k=2}^{p} |c_k| t^{\beta_k - \beta_1} \leqslant (d! - 1) t^{\beta_2 - \beta_1} \leqslant (d! - 1) t^{-\eta(\boldsymbol{M})} \leqslant 1 - 1/d! \,,$$

and so

$$\log_t |\det \boldsymbol{M}(t)| \geqslant \beta_1 + \log_t \left(1 - \sum_{k=2}^{p} |c_k| t^{\beta_k - \beta_1} \right) \geqslant \operatorname{val}(\det \boldsymbol{M}) - \log_t d! \,. \qquad \square$$

Recall that we write e for the all-1-vector of an appropriate size.

THEOREM 12. *Let $\boldsymbol{\mathcal{P}} \subset \mathbb{K}^d_+$ be a polyhedron of the form $\{\boldsymbol{x} \in \mathbb{K}^d \colon \boldsymbol{A}\boldsymbol{x} \leqslant \boldsymbol{b}\}$ where $\boldsymbol{A}$ and $\boldsymbol{b}$ are monomial. Let η_0 be the minimum of the quantities $\eta(\boldsymbol{M})$ where $\boldsymbol{M}$ is a square submatrix of $\left(\begin{smallmatrix} \boldsymbol{A} & \boldsymbol{b} & 0 \\ e^\top & 0 & 1 \end{smallmatrix}\right)$ of order d. Then, for all $t \geqslant (d!)^{1/\eta_0}$, we have*

$$d_{\mathrm{H}}(\log_t \boldsymbol{\mathcal{P}}(t), \operatorname{val}(\boldsymbol{\mathcal{P}})) \leqslant \log_t \big((d+1)^2 (d!)^4 \big) \,.$$

Proof. We employ the notation introduced in the proof of Theorem 9. Note that the inequality given in (19) holds for any sets $\boldsymbol{U}, \boldsymbol{V}$ generating the polyhedron $\boldsymbol{\mathcal{P}}$. In particular, we can set $\boldsymbol{U}$ to the set of vertices of $\boldsymbol{\mathcal{P}}$, and $\boldsymbol{V}$ to a set consisting of precisely one representative of every extreme ray of the recession cone of $\boldsymbol{\mathcal{P}}$.

Let $\boldsymbol{u} \in \boldsymbol{U}$. Since $\boldsymbol{u}$ is a vertex, there exists a subset I of cardinality d such that $\boldsymbol{A}_I \boldsymbol{u} = \boldsymbol{b}_I$, where $\boldsymbol{A}_I$ and $\boldsymbol{b}_I$ consist of the rows of $\boldsymbol{A}$ and $\boldsymbol{b}$, respectively, which are indexed by $i \in I$, and $\boldsymbol{A}_I$ is invertible. Therefore, by Cramer's rule, every coordinate $\boldsymbol{u}_i$ can be expressed as a fraction of the form $\pm \det \boldsymbol{M} / \det \boldsymbol{A}_I$ where $\boldsymbol{M}$ is a submatrix of $\begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \end{pmatrix}$ of size $d \times d$. Recall that $\boldsymbol{u}_i$ is nonnegative, and hence $\boldsymbol{u}_i = |\det \boldsymbol{M}| / |\det \boldsymbol{A}_I|$. By definition, $\eta(\boldsymbol{M})$ and $\eta(\boldsymbol{A}_I)$ are greater than or equal to η_0. From Lemma 11 we derive that

$$\operatorname{val}(\boldsymbol{u}_i) - 2 \log_t d! \leqslant \log_t \boldsymbol{u}_i(t) \leqslant \operatorname{val}(\boldsymbol{u}_i) + 2 \log_t d!$$

for all $t \geqslant (d!)^{1/\eta_0}$. Since these inequalities hold for all $i \in [d]$, we deduce that the two quantities $\delta_{\mathrm{F}}(\operatorname{val} \boldsymbol{u}, \log_t \boldsymbol{u}(t))$ and $\delta_{\mathrm{F}}(\log_t \boldsymbol{u}(t), \operatorname{val} \boldsymbol{u})$ are bounded by $2 \log_t d!$.

The recession cone of $\boldsymbol{\mathcal{P}}$ is the set $\{\boldsymbol{z} \in \mathbb{K}^d \colon \boldsymbol{A}\boldsymbol{z} \leqslant 0\}$. Since $\boldsymbol{\mathcal{P}}$ is contained in the positive orthant, so is the recession cone. Therefore, without loss of generality, we can assume that every $\boldsymbol{v} \in \boldsymbol{V}$ satisfies $\sum_i \boldsymbol{v}_i = 1$. In this way, the elements of $\boldsymbol{V}$ precisely correspond to the vertices of the polyhedron $\{\boldsymbol{z} \in \mathbb{K}^d \colon \boldsymbol{A}\boldsymbol{z} \leqslant 0 \,, e^\top \boldsymbol{z} = 1\}$. Using the same arguments as above, we infer that, for all $\boldsymbol{v} \in \boldsymbol{V}$, the two values $\delta_{\mathrm{F}}(\operatorname{val} \boldsymbol{v}, \log_t \boldsymbol{v}(t))$ and $\delta_{\mathrm{F}}(\log_t \boldsymbol{v}(t), \operatorname{val} \boldsymbol{v})$ are bounded by $2 \log_t d!$, as soon as $t \geqslant (d!)^{1/\eta_0}$. Now the claim follows from (19). $\square$

5. The tropical central path. Our core idea is to introduce the *tropical central path* of a linear program over Puiseux series. This is defined as the image of the classical primal-dual central path under the valuation map. By (6) the classical central

path is a segment of a real algebraic curve, and so its tropicalization is a segment of a (real) tropical curve and thus piecewise linear. It will turn out that for certain Puiseux linear programs the tropical central path carries a substantial amount of metric information. In sections 6 and 7 we will see that this applies to the linear programs $\mathbf{LW}_r(t)$. As its key advantage the tropical central path turns out to be much easier to analyze than its classical counterpart.

5.1. A geometric characterization of the tropical central path. As in section 3, we consider a dual pair of linear programs, except that now the coefficients lie in the field $\mathbb{K}$ of absolutely convergent real Puiseux series from section 4.1:

$$\begin{array}{lllll} \mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c}) & \text{minimize} & \langle \boldsymbol{c},\boldsymbol{x}\rangle & \text{subject to} & \boldsymbol{A}\boldsymbol{x}+\boldsymbol{w}=\boldsymbol{b}\,,\ (\boldsymbol{x},\boldsymbol{w})\in\mathbb{K}^N_+\,, \\ \mathbf{DualLP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c}) & \text{maximize} & \langle -\boldsymbol{b},\boldsymbol{y}\rangle & \text{subject to} & \boldsymbol{s}-\boldsymbol{A}^\top\boldsymbol{y}=\boldsymbol{c}\,,\ (\boldsymbol{s},\boldsymbol{y})\in\mathbb{K}^N_+\,, \end{array}$$

where $\boldsymbol{A}\in\mathbb{K}^{m\times n}$, $\boldsymbol{b}\in\mathbb{K}^m$, $\boldsymbol{c}\in\mathbb{K}^n$, and $N:=n+m$. Here, the Euclidean scalar product $\langle\cdot,\cdot\rangle$ is extended by setting $\langle \boldsymbol{u},\boldsymbol{v}\rangle:=\sum_i \boldsymbol{u}_i\boldsymbol{v}_i$ for any vectors $\boldsymbol{u},\boldsymbol{v}$ with entries over $\mathbb{K}$. Further, we define

$$\boldsymbol{\mathcal{P}}:=\{(\boldsymbol{x},\boldsymbol{w})\in\mathbb{K}^N_+\colon \boldsymbol{A}\boldsymbol{x}+\boldsymbol{w}=\boldsymbol{b}\}\quad\text{and}\quad \boldsymbol{\mathcal{Q}}:=\{(\boldsymbol{s},\boldsymbol{y})\in\mathbb{K}^N_+\colon \boldsymbol{s}-\boldsymbol{A}^\top\boldsymbol{y}=\boldsymbol{c}\}\,,$$

which correspond to the feasible regions of the two Puiseux linear programs. Moreover, we let $\boldsymbol{\mathcal{F}}:=\boldsymbol{\mathcal{P}}\times\boldsymbol{\mathcal{Q}}$ be the set of primal-dual feasible points, and $\boldsymbol{\mathcal{F}}^\circ:=\{\boldsymbol{z}\in\boldsymbol{\mathcal{F}}\colon \boldsymbol{z}>0\}$ is the strictly feasible subset. Throughout we will make the following assumption.

Assumption 13. The set $\boldsymbol{\mathcal{F}}^\circ$ is nonempty.

This allows us to define the central path of the Puiseux linear programs $\mathbf{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ and $\mathbf{DualLP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$. Indeed, applying Tarski's principle to the real closed field $\mathbb{K}$ shows that, for all $\boldsymbol{\mu}\in\mathbb{K}$ such that $\boldsymbol{\mu}>0$, the system

$$\begin{aligned} \boldsymbol{A}\boldsymbol{x}+\boldsymbol{w}&=\boldsymbol{b},\\ \boldsymbol{s}-\boldsymbol{A}^\top\boldsymbol{y}&=\boldsymbol{c},\\ \begin{pmatrix}\boldsymbol{x}\boldsymbol{s}\\ \boldsymbol{w}\boldsymbol{y}\end{pmatrix}&=\boldsymbol{\mu}e,\\ \boldsymbol{x},\boldsymbol{w},\boldsymbol{y},\boldsymbol{s}&>0 \end{aligned} \tag{20}$$

has a unique solution in $\mathbb{K}^{2N}$. We denote this solution as $\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})=(\boldsymbol{x}^{\boldsymbol{\mu}},\boldsymbol{w}^{\boldsymbol{\mu}},\boldsymbol{s}^{\boldsymbol{\mu}},\boldsymbol{y}^{\boldsymbol{\mu}})$ and refer to it as the *point of the central path with parameter* $\boldsymbol{\mu}$. Similarly, by Tarski's principle, $\boldsymbol{\mathcal{F}}\neq\emptyset$ ensures that the two linear programs have the same optimal value $\boldsymbol{\nu}\in\mathbb{K}$. We let $(\boldsymbol{x}^*,\boldsymbol{w}^*)\in\boldsymbol{\mathcal{P}}$ and $(\boldsymbol{s}^*,\boldsymbol{y}^*)\in\boldsymbol{\mathcal{Q}}$ be a pair of primal and dual optimal solutions. Equivalently, we have $\boldsymbol{\nu}=\langle\boldsymbol{c},\boldsymbol{x}^*\rangle=\langle-\boldsymbol{b},\boldsymbol{y}^*\rangle$.

Let $\mathcal{P}\subset\mathbb{T}^N$, $\mathcal{Q}\subset\mathbb{T}^N$, and $\mathcal{F}\subset\mathbb{T}^{2N}$ be the images under the valuation map of the primal and dual feasible polyhedra $\boldsymbol{\mathcal{P}}$, $\boldsymbol{\mathcal{Q}}$, and $\boldsymbol{\mathcal{F}}$, respectively. Similarly we write $(x^*,w^*)\in\mathbb{T}^N$ and $(s^*,y^*)\in\mathbb{T}^N$ for the coordinatewise valuations of the optimal solutions $(\boldsymbol{x}^*,\boldsymbol{w}^*)$ and $(\boldsymbol{s}^*,\boldsymbol{y}^*)$.

Given a primal-dual feasible point $\boldsymbol{z}=(\boldsymbol{x},\boldsymbol{w},\boldsymbol{s},\boldsymbol{y})\in\boldsymbol{\mathcal{F}}$, the *duality gap*, denoted by $\mathsf{gap}(\boldsymbol{z})$, is defined as the difference between the values of the primal and dual objective functions, i.e., $\langle\boldsymbol{c},\boldsymbol{x}\rangle+\langle\boldsymbol{b},\boldsymbol{y}\rangle$. We recall that $\mathsf{gap}(\boldsymbol{z})$ is equivalently given by the *complementarity gap* defined as the sum of the pairwise product of primal/dual variables:

$$\mathsf{gap}(\boldsymbol{z})=\langle\boldsymbol{x},\boldsymbol{s}\rangle+\langle\boldsymbol{w},\boldsymbol{y}\rangle\,.$$

Observe that the right-hand side of the latter equality consists of sums of nonnegative terms (of the form $\boldsymbol{x}_j \boldsymbol{s}_j$ and $\boldsymbol{w}_i \boldsymbol{y}_i$). Consequently, we can apply the valuation map termwise and define, for all $z = (x, w, s, y) \in \mathcal{F}$, the *tropical duality gap* as

$$\mathsf{tgap}(z) \coloneqq \langle x, s \rangle_{\mathbb{T}} \oplus \langle w, y \rangle_{\mathbb{T}} \,,$$

where $\langle \cdot, \cdot \rangle_{\mathbb{T}}$ stands for the tropical analogue of the scalar product, i.e., $\langle u, v \rangle_{\mathbb{T}} \coloneqq \bigoplus_i (u_i \odot v_i)$. Then the quantity $\mathsf{tgap}(z)$ equals the valuation of the duality gap of any primal-dual feasible $\boldsymbol{z}$ with $\mathrm{val}(\boldsymbol{z}) = z$. The study of the tropical central path will require the following *tropical sublevel set* induced by the tropical duality gap:

$$\mathcal{F}^\lambda \coloneqq \{ z \in \mathcal{F} \colon \mathsf{tgap}(z) \leqslant \lambda \} \,,$$

which is defined for any $\lambda \in \mathbb{R}$. Now we present a few basic facts about this collection of sublevel sets.

PROPOSITION 14. *Let $\boldsymbol{\mu} \in \mathbb{K}$ such that $\boldsymbol{\mu} > 0$, and let $\lambda = \mathrm{val}(\boldsymbol{\mu})$. Then*

(i) *the set $\mathcal{F}^\lambda$ is a bounded tropical polyhedron given by $\mathcal{P}^\lambda \times \mathcal{Q}^\lambda$ where*

$$\mathcal{P}^\lambda \coloneqq \{ (x, w) \in \mathcal{P} \colon \langle s^*, x \rangle_{\mathbb{T}} \oplus \langle y^*, w \rangle_{\mathbb{T}} \leqslant \lambda \} \,,$$
$$\mathcal{Q}^\lambda \coloneqq \{ (s, y) \in \mathcal{Q} \colon \langle x^*, s \rangle_{\mathbb{T}} \oplus \langle w^*, y \rangle_{\mathbb{T}} \leqslant \lambda \} \,;$$

(ii) *the image under* val *of the point $\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})$ lies in $\mathcal{F}^\lambda$.*

Proof. We start with the proof of (ii). By definition, $\boldsymbol{\mathcal{C}}(\boldsymbol{\mu}) \in \boldsymbol{\mathcal{F}}$ so that $\mathrm{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})) \in \mathcal{F}$. Moreover, we have $\mathsf{tgap}\big(\mathrm{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu}))\big) = \mathrm{val}(\mathsf{gap}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu}))$. Since $\mathsf{gap}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})) = N\boldsymbol{\mu}$ by (5), we deduce that the previous quantity is equal to the valuation of $N\boldsymbol{\mu}$, which is λ. This forces $\mathrm{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu}))$ to lie in $\mathcal{F}^\lambda$.

We need to show that $\mathcal{F}^\lambda$ is a tropical polyhedron which is bounded. To this end we consider $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{s}, \boldsymbol{y}) \in \boldsymbol{\mathcal{F}}$. Recall that $\boldsymbol{\nu}$ is the common optimal value of the primal and the dual Puiseux linear programs. Then we obtain

$$\mathsf{gap}(\boldsymbol{z}) = \big(\langle \boldsymbol{c}, \boldsymbol{x} \rangle - \boldsymbol{\nu}\big) + \big(\langle \boldsymbol{b}, \boldsymbol{y} \rangle + \boldsymbol{\nu}\big) = \mathsf{gap}(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{s}^*, \boldsymbol{y}^*) + \mathsf{gap}(\boldsymbol{x}^*, \boldsymbol{w}^*, \boldsymbol{s}, \boldsymbol{y}) \,.$$

Since the right-hand side of this identity is a sum of two nonnegative terms, applying the valuation map yields $\mathsf{tgap}(z) = \mathsf{tgap}(x, w, s^*, y^*) \oplus \mathsf{tgap}(x^*, w^*, s, y)$ for all $z = (x, w, s, y) \in \mathcal{F}$. Thus we can express the tropical sublevel set $\mathcal{F}^\lambda$ as

$$\mathcal{F}^\lambda = \{ (x, w, s, y) \in \mathcal{F} \colon \mathsf{tgap}(x, w, s^*, y^*) \oplus \mathsf{tgap}(x^*, w^*, s, y) \leqslant \lambda \} = P^\lambda \times Q^\lambda \,.$$

In particular, the set $\mathcal{F}^\lambda$ is a tropical polyhedron.

To complete our proof we still have to show that $\mathcal{F}^\lambda$ is bounded. As the entries of every $z \in \mathcal{F}^\lambda$ are greater than or equal to $-\infty \in \mathbb{T}$, we just need to prove that $\mathcal{F}^\lambda$ is bounded from above. To this end, pick some point $z^\circ = (x^\circ, w^\circ, s^\circ, y^\circ)$ in $\mathcal{F}^\lambda$ which has finite coordinates. For instance, we can take $z^\circ = \mathrm{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu}))$ by making use of (ii). Now consider an arbitrary point $(x, w, s, y) \in \mathcal{F}^\lambda$. As $\mathcal{F}^\lambda = \mathcal{P}^\lambda \times \mathcal{Q}^\lambda$, we know that $(x, w, s^\circ, y^\circ) \in \mathcal{F}^\lambda$. In particular, $\mathsf{tgap}(x, w, s^\circ, y^\circ) \leqslant \lambda$ or, equivalently,

$$\begin{cases} x_j \odot s_j^\circ \leqslant \lambda & \text{for all } j \in [n] \,, \\ w_i \odot y_i^\circ \leqslant \lambda & \text{for all } i \in [m] \,. \end{cases}$$

Since $s_j^\circ, y_i^\circ > -\infty$, this entails that $x_j \leqslant \lambda \odot (s_j^\circ)^{\odot(-1)}$ and $w_i \leqslant \lambda \odot (y_i^\circ)^{\odot(-1)}$ for all $i \in [m]$ and $j \in [n]$. Similarly, the entries of (s, y) are bounded, too. This completes the proof of (i). ∎

As a bounded tropical polyhedron the set $\mathcal{F}^\lambda$ admits a (tropical) barycenter. Recall that the latter was defined as the coordinatewise maximum of that set. The following theorem relates this barycenter with the valuation of the central path, and gives rise to the definition of the tropical central path.

THEOREM 15. *The image under the valuation map of the central path of the pair of primal-dual linear programs* $\boldsymbol{LP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *and* $\boldsymbol{DualLP}(\boldsymbol{A},\boldsymbol{b},\boldsymbol{c})$ *can be described by*

$$\operatorname{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})) = \textit{barycenter of } \{z \in \operatorname{val}(\boldsymbol{\mathcal{F}})\colon \mathtt{tgap}(z) \leqslant \operatorname{val}(\boldsymbol{\mu})\} \tag{21}$$

for any $\boldsymbol{\mu} \in \mathbb{K}$ *such that* $\boldsymbol{\mu} > 0$.

Proof. Let $\lambda \coloneqq \operatorname{val}(\boldsymbol{\mu})$, and denote by $\bar{z} = (\bar{x}, \bar{w}, \bar{s}, \bar{y})$ the barycenter of the tropical polyhedron $\mathcal{F}^\lambda$. By construction we have $\operatorname{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})) \leqslant \bar{z}$. Moreover, since $\mathtt{tgap}(\bar{z}) \leqslant \lambda$, we also have $\bar{x}_j \odot \bar{s}_j \leqslant \lambda$ and $\bar{w}_i \odot \bar{y}_i \leqslant \lambda$ for all $i \in [m]$ and $j \in [n]$. It follows that

$$\begin{cases} \lambda = \operatorname{val}(\boldsymbol{x}_j^{\boldsymbol{\mu}} \boldsymbol{s}_j^{\boldsymbol{\mu}}) = \operatorname{val}(\boldsymbol{x}_j^{\boldsymbol{\mu}}) \odot \operatorname{val}(\boldsymbol{s}_j^{\boldsymbol{\mu}}) \leqslant \bar{x}_j \odot \bar{s}_j \leqslant \lambda\,, \\ \lambda = \operatorname{val}(\boldsymbol{w}_j^{\boldsymbol{\mu}} \boldsymbol{y}_j^{\boldsymbol{\mu}}) = \operatorname{val}(\boldsymbol{w}_i^{\boldsymbol{\mu}}) \odot \operatorname{val}(\boldsymbol{y}_i^{\boldsymbol{\mu}}) \leqslant \bar{w}_i \odot \bar{y}_i \leqslant \lambda\,. \end{cases} \tag{22}$$

As a consequence, the inequality $\operatorname{val}(\boldsymbol{\mathcal{C}}(\boldsymbol{\mu})) \leqslant \bar{z}$ is necessarily an equality. ∎

The quantity (21), which depends only on the valuation of $\boldsymbol{\mu}$, is called *the tropical central path* at $\lambda = \operatorname{val}(\boldsymbol{\mu})$ and is denoted by

$$\mathcal{C}^{\mathrm{trop}}(\lambda) = (x^\lambda, w^\lambda, s^\lambda, y^\lambda)\,.$$

Analogously, the *primal* and the *dual tropical central paths* are defined by projecting to (x^λ, w^λ) and (s^λ, y^λ), respectively. As shown in (22), the primal and dual components of the tropical central path are characterized by

$$x_j^\lambda \odot s_j^\lambda = \lambda = w_i^\lambda \odot y_i^\lambda \tag{23}$$

for all $i \in [m]$ and $j \in [n]$. The next statement shows that the tropical central path is a polygonal curve with a particularly simple structure.

PROPOSITION 16. *The tropical central path* $\lambda \mapsto \mathcal{C}^{\mathrm{trop}}(\lambda)$ *is a monotone piecewise linear function, whose derivative at each smooth point is a vector of the form* $(e^K, e^{[N]\setminus K})$ *for some* $K \subset [N]$.

Proof. Let us denote by $g\colon \mathbb{T}^N \to \mathbb{T}$ the function which sends (x, w) to $\langle s^*, x\rangle_{\mathbb{T}} \oplus \langle y^*, w\rangle_{\mathbb{T}}$. Pick finite generating sets $U \subset \mathbb{T}^N$ and $V \subset \mathbb{T}^N \setminus \{(-\infty, \dots, -\infty)\}$ for the tropical polyhedron $\mathcal{P}$. Since (x^λ, w^λ) lies in $\mathcal{P}$, it can be expressed as

$$(x^\lambda, w^\lambda) = \left(\bigoplus_{u\in U} \alpha_u \odot u\right) \oplus \left(\bigoplus_{v\in V} \beta_v \odot v\right),$$

where $\bigoplus_{u\in U} \alpha_u = 0$. The inequality $g(x^\lambda, w^\lambda) \leqslant \lambda$ now amounts to $\alpha_u \odot g(u) \leqslant \lambda$ and $\beta_v \odot g(v) \leqslant \lambda$ for all $u \in U$ and $v \in V$. As (x^λ, w^λ) is the barycenter of $\mathcal{P}^\lambda$, the coefficients α_u and β_v can be chosen to be maximal. This enforces $\alpha_u = \min\bigl(0, \lambda \odot (g(u))^{\odot(-1)}\bigr)$ and $\beta_v = \lambda \odot (g(v))^{\odot(-1)}$, using the convention $(-\infty)^{\odot(-1)} = +\infty$. Note that $g(v) \neq -\infty$ for all $v \in V$. Indeed, if there were a ray $v \in V$ with $g(v) = -\infty$, then any point of the form $(x^*, w^*) \oplus (\beta \odot v)$ would belong to $\mathcal{P}^\lambda$. The latter conclusion would contradict the boundedness of $\mathcal{P}^\lambda$. Therefore, all α_u and β_v

belong to $\mathbb{R}$. Observe also that α_u and β_v, thought of as functions of λ, are monotone and piecewise linear, and that their derivatives at any smooth point take value in $\{0,1\}$. It follows that $\lambda \mapsto (x^\lambda, s^\lambda)$ is piecewise linear and monotone, and that its derivative at any smooth point is of the form e^K.

A similar argument reveals that the map $\lambda \mapsto (s^\lambda, y^\lambda)$ is also piecewise linear and monotone, with a derivative at any smooth point of the form $e^{K'}$ for some set $K' \subset [N]$. As a consequence, the derivative of the map $\mathcal{C}^{\mathrm{trop}}$ at any smooth point is of the form $(e^K, e^{K'})$, and, from (23), we get that $K' = [N] \setminus K$. This proves our claim. ∎

We end this section with the following direct consequence of Proposition 16.

COROLLARY 17. *If $\lambda \leqslant \lambda'$, then $\mathcal{C}^{\mathrm{trop}}(\lambda) \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda') \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda) + (\lambda' - \lambda)e$.*

5.2. A uniform metric estimate on the convergence of the central path. So far, we have related the tropical central path with the central path of the linear programs $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ and $\mathbf{DualLP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ over Puiseux series. These give rise to a parametric family of dual linear programs $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $\mathrm{DualLP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ over the reals. The purpose of this section is to relate the resulting family of central paths with the tropical central path. More precisely, we will show that the logarithmic deformation of these central paths uniformly converges to the tropical curve. In fact, we will even show that the logarithmic deformation of the wide neighborhoods $\mathcal{N}_\theta^{-\infty}$ of these central paths collapses onto the tropical central path.

Let $\boldsymbol{\mathcal{P}}(t), \boldsymbol{\mathcal{Q}}(t) \subset \mathbb{R}^N$ be the feasible sets of the linear programs $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $\mathrm{DualLP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$, respectively. (Note that this notation is compatible with that introduced in section 4.1 thanks to Proposition 3.) Then $\boldsymbol{\mathcal{F}}(t) \coloneqq \boldsymbol{\mathcal{P}}(t) \times \boldsymbol{\mathcal{Q}}(t)$ is the primal-dual feasible set, while $\boldsymbol{\mathcal{F}}^\circ(t) \coloneqq \{z \in \boldsymbol{\mathcal{F}}(t) \colon z > 0\}$ comprises only those primal-dual points which are strictly feasible. The following lemma relates the optimal solutions of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ and $\mathbf{DualLP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ with those of $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $\mathrm{DualLP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$.

LEMMA 18. *There exists a real number $t_0 > 1$ such that for all $t > t_0$ the following three properties hold:*

(i) *The set $\boldsymbol{\mathcal{F}}^\circ(t)$ is nonempty;*
(ii) *the number $\boldsymbol{\nu}(t)$ is the optimal value of $LP(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $DualLP(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$;*
(iii) *$(\boldsymbol{x}^*(t), \boldsymbol{w}^*(t))$ and $(\boldsymbol{s}^*(t), \boldsymbol{y}^*(t))$ constitute optimal solutions.*

Proof. For two series $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{K}$ the equality $\boldsymbol{u} = \boldsymbol{v}$ forces $\boldsymbol{u}(t) = \boldsymbol{v}(t)$ for all t large enough. A similar statement holds for inequalities like $\boldsymbol{u} \leqslant \boldsymbol{v}$ and $\boldsymbol{u} < \boldsymbol{v}$. We infer that the set $\boldsymbol{\mathcal{F}}^\circ(t)$ is not empty if $t \gg 1$. Moreover, $(\boldsymbol{x}^*(t), \boldsymbol{w}^*(t)) \in \mathcal{P}(t)$, $(\boldsymbol{s}^*(t), \boldsymbol{y}^*(t)) \in \mathcal{Q}(t)$, and $\langle \boldsymbol{c}(t), \boldsymbol{x}^*(t) \rangle = \langle -\boldsymbol{b}(t), \boldsymbol{y}^*(t) \rangle = \boldsymbol{\nu}(t)$. Since $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $\mathrm{DualLP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ are dual to one another, we conclude that $(\boldsymbol{x}^*(t), \boldsymbol{w}^*(t))$ and $(\boldsymbol{s}^*(t), \boldsymbol{y}^*(t))$ form a pair of optimal solutions. ∎

Throughout the following we will keep the value t_0 from Lemma 18. When $t > t_0$, we know from Lemma 18(i) that the primal-dual central path of the linear programs $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and $\mathrm{DualLP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ is well defined. In this case we denote by $\mathcal{C}_t(\mu)$ the point of this central path with parameter μ, where $\mu \in \mathbb{R}$ and $\mu > 0$. Let us fix the real precision parameter θ in the open interval from 0 to 1. Then the set

$$\mathcal{N}_{\theta,t}^{-\infty}(\mu) \coloneqq \left\{ z = (x, w, s, y) \in \boldsymbol{\mathcal{F}}^\circ(t) \colon \bar{\mu}(z) = \mu \text{ and } \begin{pmatrix} xs \\ wy \end{pmatrix} \geqslant (1-\theta)\mu e \right\}$$

is a neighborhood of the point $\mathcal{C}_t(\mu)$. A direct inspection shows that the union of the sets $\mathcal{N}_{\theta,t}^{-\infty}(\mu)$ for $\mu > 0$ agrees with the wide neighborhood $\mathcal{N}_\theta^{-\infty}$ of the entire central

path of the linear program $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ over $\mathbb{R}$; see (8). In order to stress the dependence on t, we denote this neighborhood by $\mathcal{N}_{\theta,t}^{-\infty}$. With this notation, we have

$$\mathcal{N}_{\theta,t}^{-\infty} = \bigcup_{\mu>0} \mathcal{N}_{\theta,t}^{-\infty}(\mu)\,.$$

Further, let

$$\delta(t) \coloneqq 2d_{\mathrm{H}}(\log_t \boldsymbol{\mathcal{F}}(t), \mathcal{F})\,, \tag{24}$$

which, by Theorem 9, tends to 0 when t goes to $+\infty$. The following result states that we can uniformly bound the distance from the image of $\mathcal{N}_{\theta,t}^{-\infty}(\mu)$ under $\log_t$ to the point $\mathcal{C}^{\mathrm{trop}}(\log_t \mu)$ of the tropical central path, independently of μ.

THEOREM 19. *For all $t > t_0$ and $\mu > 0$ we have*

$$d_\infty\big(\log_t \mathcal{N}_{\theta,t}^{-\infty}(\mu), \mathcal{C}^{\mathrm{trop}}(\log_t \mu)\big) \leqslant \log_t\Big(\frac{N}{1-\theta}\Big) + \delta(t)\,.$$

Proof. Choose $t > t_0$, $\mu > 0$, and $z = (x, w, s, y) \in \mathcal{N}_{\theta,t}^{-\infty}(\mu)$. Letting $\lambda \coloneqq \log_t \mu$, we claim that it suffices to prove that

$$\log_t z \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda) + (\log_t N + \delta(t))e\,. \tag{25}$$

Indeed, by definition of the wide neighborhood $\mathcal{N}_{\theta,t}^{-\infty}(\mu)$, we have $\log_t(x, w) \geqslant -\log_t(s, y) + \big(\lambda + \log_t(1-\theta)\big)e$. Using (25), we obtain

$$\log_t(x, w) \geqslant -(s^\lambda, y^\lambda) + \Big(\lambda - \log_t\Big(\frac{N}{1-\theta}\Big) - \delta(t)\Big)e = (x^\lambda, w^\lambda) - \Big(\log_t\Big(\frac{N}{1-\theta}\Big) + \delta(t)\Big)e\,,$$

where the last equality is due to (23). Analogously, we can prove that $\log_t(s, y) \geqslant (s^\lambda, y^\lambda) - \Big(\log_t\big(\frac{N}{1-\theta}\big) + \delta(t)\Big)e$. Finally, the inequality in (25) implies that $\log_t z \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda) + \Big(\log_t\big(\frac{N}{1-\theta}\big) + \delta(t)\Big)e$, since $\theta > 0$.

Now let us show that (25) holds. By definition of the duality measure $\bar{\mu}(z)$, we have $\mathsf{gap}(z) = N\bar{\mu}(z) = N\mu$. Applying the map $\log_t$ yields

$$\mathsf{tgap}(\log_t z) \leqslant \log_t \mathsf{gap}(z) = \lambda + \log_t N\,, \tag{26}$$

where the inequality is a consequence of the first inequality in (16).

Let $z' \in \mathcal{F}$ such that $d_{\mathrm{H}}(\log_t z, z') < +\infty$. Recall that

$$z' - \delta_{\mathrm{F}}(\log_t z, z')e \leqslant \log_t z \leqslant z' + \delta_{\mathrm{F}}(z', \log_t z)e\,. \tag{27}$$

The first inequality in (27) gives $\mathsf{tgap}(z') \leqslant \mathsf{tgap}(\log_t z) + 2\delta_{\mathrm{F}}(\log_t z, z')$. In combination with (26) this shows that z' lies in $\mathcal{F}^{\lambda'}$ for $\lambda' \coloneqq \lambda + \log_t N + 2\delta_F(\log_t z, z')$. The second inequality in (27) now yields

$$\log_t z \leqslant z' + \delta_{\mathrm{F}}(z', \log_t z)e \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda') + \delta_{\mathrm{F}}(z', \log_t z)e \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda) + \big(\log_t N + 2d_{\mathrm{H}}(\log_t z, z')\big)e\,,$$

where the second inequality follows from $\mathcal{C}^{\mathrm{trop}}(\lambda')$ being the barycenter of $\mathcal{F}^{\lambda'}$, and the last inequality is a consequence of Corollary 17. As this argument is valid for all $z' \in \mathcal{F}$ within a finite distance from $\log_t z$, we obtain that $\log_t z \leqslant \mathcal{C}^{\mathrm{trop}}(\lambda) + \big(\log_t N + \delta(t)\big)e$. □

5.3. Main example. The family $\mathbf{LW}_r(t)$ of linear programs over the reals from the introduction may also be read as a linear program over the field $\mathbb{K}$, when t is thought of as a formal parameter. We denote this linear program by $\mathbf{LW}_r$. The goal of this section is to obtain a complete description of the corresponding tropical central path.

Introducing slack variables $\boldsymbol{w}_1, \dots, \boldsymbol{w}_{3r-1}$ in the first $3r-1$ inequalities of $\mathbf{LW}_r$, and adding the redundant inequalities $\boldsymbol{x}_i \geqslant 0$ for $1 \leqslant i < 2r-1$, gives rise to a linear program $\mathbf{LW}_r^=$, which is of the form $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ in dimension $N = 5r-1$. (Note that the last two inequalities of $\mathbf{LW}_r$ are nonnegativity constraints, which is why we do not need slack variables for them.) The dual Puiseux linear program (with slacks) is referred to as $\textsc{DualLW}_r^=$. We retain the notation introduced in section 5.1; for instance, we denote by $\boldsymbol{\mathcal{P}}$ and $\boldsymbol{\mathcal{Q}}$ the primal and dual feasible sets, respectively.

To begin with, we verify that Assumption 13 is satisfied. Due to the lower triangular nature of the system of inequalities in $\mathbf{LW}_r$, we can easily find a vector $\boldsymbol{x}$ satisfying every inequality of this system in a strict manner. In other words, we can find $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}$ such that $\boldsymbol{x} > 0$ and $\boldsymbol{w} > 0$. Moreover, since the feasible set of $\mathbf{LW}_r$ is bounded, the set $\boldsymbol{\mathcal{P}}$ is bounded as well. This implies that the dual feasible set $\boldsymbol{\mathcal{Q}}$ contains a point $(\boldsymbol{s}, \boldsymbol{y})$ satisfying $\boldsymbol{s} > 0$, $\boldsymbol{y} > 0$. As a result, the set $\boldsymbol{\mathcal{F}}^\circ$ is nonempty.

We focus on the description of the primal part $\lambda \mapsto (x^\lambda, w^\lambda)$ of the tropical central path, since the dual part can be readily obtained by using the relations (23). It can be checked that the optimal value of $\mathbf{LW}_r^=$, and subsequently of $\textsc{DualLW}_r^=$, is equal to 0. Since in our case, the primal objective vector $\boldsymbol{c}$ is given by the nonnegative vector $(1, 0, \dots, 0) \in \mathbb{K}^n$, we deduce that we can choose the dual optimal solution $(\boldsymbol{s}^*, \boldsymbol{y}^*)$ as $(1, 0, \dots, 0) \in \mathbb{K}^{n+m}$. As a consequence of Theorem 15 and Proposition 14(i), the point (x^λ, w^λ) on the primal tropical central path agrees with the barycenter of the tropical sublevel set

$$\mathcal{P}^\lambda = \{(x, w) \in \mathcal{P} \colon x_1 \leqslant \lambda\}\,. \tag{28}$$

Recall that $\mathcal{P}$ stands for $\mathrm{val}(\boldsymbol{\mathcal{P}})$.

We first restrict our attention to the x-component of the tropical central path. To this end, let $\boldsymbol{\mathcal{P}'}$ be the projection of the primal feasible set $\boldsymbol{\mathcal{P}}$ onto the coordinates $\boldsymbol{x}_1, \dots, \boldsymbol{x}_{2r}$. This is precisely the feasible set of the Puiseux linear program $\mathbf{LW}_r$. Further, let $\mathcal{P}'$ be the image under val of $\boldsymbol{\mathcal{P}'}$. Equivalently, this is the projection of $\mathcal{P}$ onto $x_1, \dots, x_{2r}$. We claim that $\mathcal{P}'$ is given by the $3r+1$ tropical linear inequalities

$$\left.\begin{array}{l} x_1 \leqslant 2\,,\ x_2 \leqslant 1, \\ x_{2j+1} \leqslant 1 + x_{2j-1}\,,\ x_{2j+1} \leqslant 1 + x_{2j} \\ x_{2j+2} \leqslant (1 - 1/2^j) + \max(x_{2j-1}, x_{2j}) \end{array}\right] \ 1 \leqslant j < r, \tag{29}$$

which are obtained by applying the valuation map to the inequalities in $\mathbf{LW}_r$ coefficientwise. While this can be checked by hand, we can also apply [AGS20, Corollary 14], as $\mathcal{P}'$ is a regular set in $\mathbb{T}^{2r}$, i.e., it coincides with the closure of its interior.

By (28) we deduce that the point x^λ is the barycenter of the tropical polyhedron $\{x \in \mathcal{P}' \colon x_1 \leqslant \lambda\}$. We arrive at the following explicit description of x^λ.

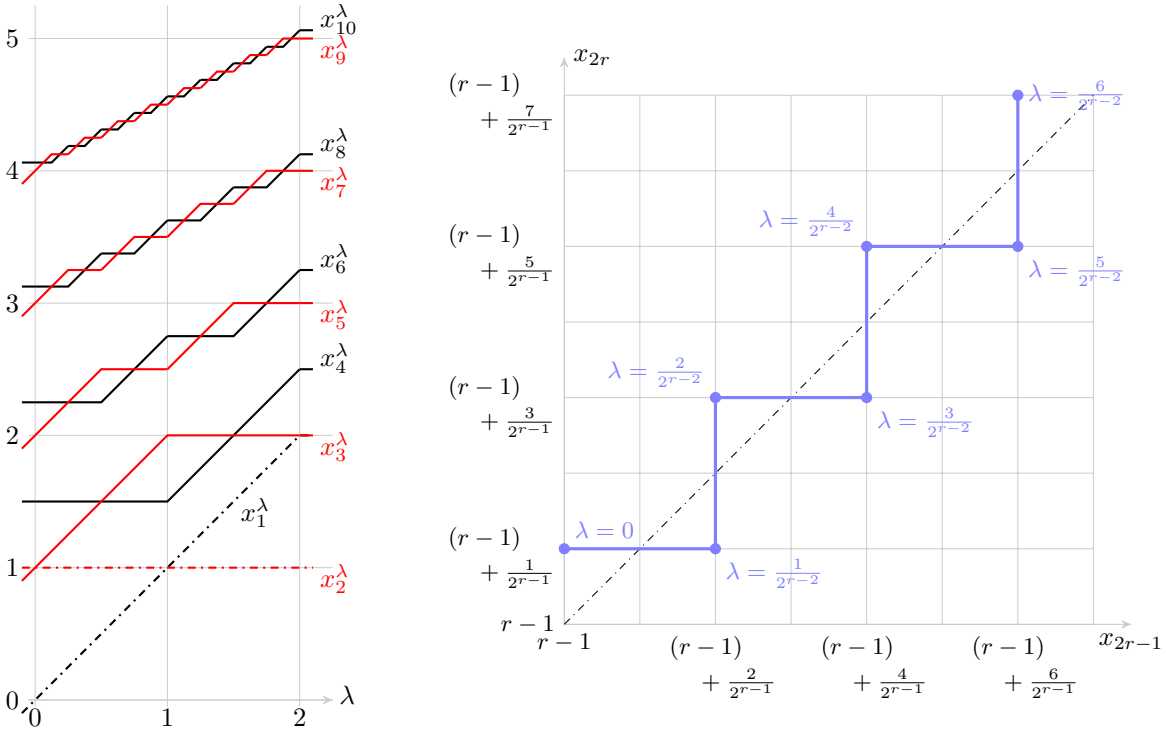


FIG. 4. *Left: the x-components of the primal tropical central path of $\mathbf{LW}_r$ for $r \geqslant 5$ and $0 \leqslant \lambda \leqslant 2$. Right: the projection of the tropical central path of $\mathbf{LW}_r$ onto the (x_{2r-1}, x_{2r})-plane.*

PROPOSITION 20. *For all $\lambda \in \mathbb{R}$, the point x^λ is given by the recursion*

$$\left.\begin{aligned} x_1^\lambda &= \min(\lambda, 2), \\ x_2^\lambda &= 1, \\ x_{2j+1}^\lambda &= 1 + \min(x_{2j-1}^\lambda, x_{2j}^\lambda) \\ x_{2j+2}^\lambda &= (1 - 1/2^j) + \max(x_{2j-1}^\lambda, x_{2j}^\lambda) \end{aligned}\right] \quad 1 \leqslant j < r.$$

Proof. We introduce the family of maps $F_j \colon (a, b) \mapsto (1 + \min(a, b), 1 - 1/2^j + \max(a, b))$ where $1 \leqslant j < r$. With this notation, the point x lies in $\mathcal{P}'$ and also satisfies $x_1 \leqslant \lambda$ if and only if

$$x_1 \leqslant \min(\lambda, 2)\,, \quad x_2 \leqslant 1\,, \quad (x_{2j+1}^\lambda, x_{2j+2}^\lambda) \leqslant F_j(x_{2j-1}^\lambda, x_{2j}^\lambda) \tag{30}$$

for every $1 \leqslant j < r$. Since the maps F_j are order preserving, the barycenter of the tropical polyhedron defined by (30) is the point which attains equality in (30). ∎

Observe that the map $\lambda \mapsto x^\lambda$ is constant on the interval $[2, \infty[$, while it is linear on $]-\infty, 0]$. In contrast, on the remaining interval $[0, 2]$, the shape of this map is much more complicated. This is illustrated in Figure 4.

We now incorporate the slack variables $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{3r-1}$ into our analysis. The

point $(\boldsymbol{x}, \boldsymbol{w})$ in the primal feasible set $\boldsymbol{\mathcal{P}}$ is defined by the following constraints:

$$\left.\begin{aligned} \boldsymbol{x}_1 + \boldsymbol{w}_1 &= t^2, \\ \boldsymbol{x}_2 + \boldsymbol{w}_2 &= t, \\ \boldsymbol{x}_{2j+1} + \boldsymbol{w}_{3j} &= t\,\boldsymbol{x}_{2j-1} \\ \boldsymbol{x}_{2j+1} + \boldsymbol{w}_{3j+1} &= t\,\boldsymbol{x}_{2j} \\ \boldsymbol{x}_{2j+2} + \boldsymbol{w}_{3j+2} &= t^{1-1/2^j}(\boldsymbol{x}_{2j-1} + \boldsymbol{x}_{2j}) \\ (\boldsymbol{x}, \boldsymbol{w}) &\in \mathbb{K}_+^N. \end{aligned}\right\} \quad 1 \leqslant j < r. \tag{31}$$

This entails that the points (x, w) in $\mathcal{P} = \mathrm{val}(\boldsymbol{\mathcal{P}})$ satisfy the inequalities

$$\left.\begin{aligned} & w_1 \leqslant 2\,,\ w_2 \leqslant 1, \\ & w_{3j} \leqslant 1 + x_{2j-1}\,,\ w_{3j+1} \leqslant 1 + x_{2j} \\ & w_{3j+2} \leqslant (1 - 1/2^j) + \max(x_{2j-1}, x_{2j}) \end{aligned}\right\} \quad 1 \leqslant j < r. \tag{32}$$

The following result states that all these inequalities are tight for all points (x^λ, w^λ) on the primal tropical central path.

Proposition 21. *For all $\lambda \in \mathbb{R}$, the point w^λ is described by the following relations:*

$$\left.\begin{aligned} w_1^\lambda &= 2\,,\ w_2^\lambda = 1, \\ w_{3j}^\lambda &= 1 + x_{2j-1}^\lambda \\ w_{3j+1}^\lambda &= 1 + x_{2j}^\lambda \\ w_{3j+2}^\lambda &= (1 - 1/2^j) + \max(x_{2j-1}^\lambda, x_{2j}^\lambda) = x_{2j+2}^\lambda \end{aligned}\right\} \quad 1 \leqslant j < r. \tag{33}$$

Proof. Let $\bar{w}$ be the element defined by the relations in (33). We want to prove that $w^\lambda = \bar{w}$. By Theorem 15 and Proposition 14(i), it suffices to show that the point $(x^\lambda, \bar{w})$ is the barycenter of the tropical polyhedron $\mathcal{P}^\lambda$. Given $(x, w) \in \mathcal{P}^\lambda$, we have $x \leqslant x^\lambda$, as x belongs to $\{x \in \mathcal{P}' \colon x_1 \leqslant \lambda\}$ and x^λ is the barycenter of the latter set. Moreover, w satisfies the inequalities given in (32). We deduce that $w \leqslant \bar{w}$.

It now remains to show that $(x^\lambda, \bar{w})$ belongs to $\mathcal{P}$, since this immediately leads to $(x^\lambda, \bar{w}) \in \mathcal{P}^\lambda$. In other words, we want to find a point $(\boldsymbol{x}, \boldsymbol{w}) \in \boldsymbol{\mathcal{P}}$ such that $\mathrm{val}(\boldsymbol{x}, \boldsymbol{w}) = (x^\lambda, \bar{w})$. Let us fix a sequence of positive numbers $\alpha_0 = \frac{1}{2} > \alpha_1 > \cdots > \alpha_{r-1} > 0$. We claim that letting

$$\boldsymbol{x}_{2j+1} \coloneqq \alpha_j t^{x_{2j+1}^\lambda}\,, \qquad \boldsymbol{x}_{2j+2} \coloneqq \alpha_j t^{x_{2j+2}^\lambda} \qquad (0 \leqslant j < r) \tag{34}$$

and defining $\boldsymbol{w}$ in terms of the equalities in (31) yields such an admissible lift.

First, observe that $\boldsymbol{x} \geqslant 0$ and $\mathrm{val}(\boldsymbol{x}) = x^\lambda$. Second, we have $\boldsymbol{w}_1 = t^2 - \alpha_1 t^{x_1^\lambda}$ and $\boldsymbol{w}_2 = \alpha_0 t$. Recall that $x_1^\lambda \leqslant 2$. Thus, $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ are nonnegative, and they satisfy $\mathrm{val}\,\boldsymbol{w}_1 = 2$ and $\mathrm{val}\,\boldsymbol{w}_2 = 1$. Now, let us consider j for $1 \leqslant j < r$. We have

$$\boldsymbol{w}_{3j} = t\,\boldsymbol{x}_{2j-1} - \boldsymbol{x}_{2j+1} = \alpha_{j-1} t^{1+x_{2j-1}^\lambda} - \alpha_j t^{x_{2j+1}^\lambda} = (\alpha_{j-1} - \alpha_j) t^{\bar{w}_{3j}} + \alpha_j (t^{\bar{w}_{3j}} - t^{x_{2j+1}^\lambda})\,.$$

As $\bar{w}_{3j} = 1 + x_{2j-1}^\lambda \geqslant x_{2j+1}^\lambda$, we have $0 \leqslant t^{\bar{w}_{3j}} - t^{x_{2j+1}^\lambda} \leqslant t^{\bar{w}_{3j}}$, and this gives us $\boldsymbol{w}_{3j} \geqslant 0$ and $\mathrm{val}\,\boldsymbol{w}_{3j} = \bar{w}_{3j}$. A similar argument shows that $\boldsymbol{w}_{3j+1} \geqslant 0$ and $\mathrm{val}\,\boldsymbol{w}_{3j+1} = \bar{w}_{3j+1}$.

TABLE 1

Coordinates of points on the primal tropical central path of $\mathbf{LW}_r$ *for some specific values of* λ*, where* $1 \leqslant j < r$ *and* $k = 0, 2, \ldots, 2^{j-1} - 2$.

λ	$\frac{4k}{2^j}$	$\frac{4k+2}{2^j}$	$\frac{4k+4}{2^j}$	$\frac{4k+6}{2^j}$	$\frac{4k+8}{2^j}$
x_{2j+1}	$j + \frac{2k}{2^j}$	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+4}{2^j}$	$j + \frac{2k+4}{2^j}$
x_{2j+2}	$j + \frac{2k+1}{2^j}$	$j + \frac{2k+1}{2^j}$	$j + \frac{2k+3}{2^j}$	$j + \frac{2k+3}{2^j}$	$j + \frac{2k+5}{2^j}$
w_{3j}	$j + \frac{2k}{2^j}$	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+4}{2^j}$	$j + \frac{2k+4}{2^j}$	$j + \frac{2k+4}{2^j}$
w_{3j+1}	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+2}{2^j}$	$j + \frac{2k+4}{2^j}$	$j + \frac{2k+6}{2^j}$
w_{3j+2}	$j + \frac{2k+1}{2^j}$	$j + \frac{2k+1}{2^j}$	$j + \frac{2k+3}{2^j}$	$j + \frac{2k+3}{2^j}$	$j + \frac{2k+5}{2^j}$

Finally, we can write

$$\begin{aligned}\boldsymbol{w}_{3j+2} &= t^{1-1/2^j}(\boldsymbol{x}_{2j-1} + \boldsymbol{x}_{2j}) - \boldsymbol{x}_{2j+2} \\ &= \begin{cases} (2\alpha_{j-1} - \alpha_j)t^{x^\lambda_{2j+2}} & \text{if } x^\lambda_{2j-1} = x^\lambda_{2j}\,, \\ (\alpha_{j-1} - \alpha_j)t^{x^\lambda_{2j+2}} + o(t^{x^\lambda_{2j+2}}) & \text{otherwise.} \end{cases}\end{aligned}$$

Since $2\alpha_{j-1} > \alpha_{j-1} > \alpha_j$, we obtain that $\boldsymbol{w}_{3j+2} \geqslant 0$, and $\operatorname{val} \boldsymbol{w}_{3j+2} = x^\lambda_{2j+2} = \bar{w}_{3j+2}$. ∎

Table 1 gives a summary of the values of the coordinates of the primal tropical central path for specific values of λ which we shall use below.

6. Curvature analysis. The purpose of this section is to show how the combinatorial analysis of the tropical central path translates into lower bounds on the total curvature of the central path of a parametric family of linear programs over the reals. Our main application will be a detailed version of Theorem B from the introduction, and a proof of this result.

Let us recall some basic facts concerning total curvature. Let σ be a twice continuously differentiable curve parameterized over an interval $[a, b]$. Without loss of generality, assume that σ is parameterized by its arc length, so that $\|\sigma'(s)\| = 1$ for all $s \in [a, b]$. Then, the *total curvature* $\kappa(\sigma, [a, b])$ is defined as $\int_a^b \|\sigma''(s)\| ds$. In other words, this is the integral of the norm of the acceleration if we move on the curve at unit speed.

The definition of the total curvature can be generalized to nonsmooth curves as follows. For two nonnull vectors $x, y \in \mathbb{R}^d$ we denote by $\angle xy$ the measure $\alpha \in [0, \pi]$ of the angle of the vectors x and y, so that

$$\cos \alpha = \frac{\langle x, y \rangle}{\|x\| \|y\|}\,,$$

where $\|\cdot\|$ refers to the Euclidean norm. Given three points $U, V, W \in \mathbb{R}^d$ such that $U \neq V$ and $V \neq W$, we extend this notation to write $\angle UVW$ for the angle formed by the vectors UV and VW. If τ is a polygonal curve in $\mathbb{R}^d$ parameterized over an interval $[a, b]$, the total curvature $\kappa(\tau, [a, b])$ is defined as the sum of angles between the consecutive segments of the curve. Then, if σ is an arbitrary curve, its total curvature $\kappa(\sigma, [a, b])$ is defined as the supremum of $\kappa(\tau, [a, b])$ over all polygonal curves τ inscribed in σ. We point out that, for smooth curves, this coincides with the

definition of the total curvature that we have previously given; see [AR89, Chapter V] for more background.

Our approach is based on estimating the curvature of the central path using approximations by polygonal curves. Our first observation is concerned with limits of angles between families of vectors arising from vectors over $\mathbb{K}$.

LEMMA 22. *Let $\boldsymbol{x}, \boldsymbol{y}$ be two nonnull vectors in $\mathbb{K}^d$, and let $x \coloneqq \mathrm{val}(\boldsymbol{x})$ and $y \coloneqq \mathrm{val}(\boldsymbol{y})$. Then the limit of $\angle \boldsymbol{x}(t)\boldsymbol{y}(t)$ for $t \to +\infty$ exists. Moreover, if the sets $\arg\max_{i\in[d]} x_i$ and $\arg\max_{i\in[d]} y_i$ are disjoint, then we have*

$$\lim_{t\to+\infty} \angle \boldsymbol{x}(t)\boldsymbol{y}(t) = \frac{\pi}{2}\,.$$

Proof. Since the field $\mathbb{K}$ is real closed, the Euclidean norm $\|\cdot\|$ can be extended to a function from $\mathbb{K}^d$ to $\mathbb{K}$ by $\|\boldsymbol{u}\| \coloneqq \sqrt{\sum_i \boldsymbol{u}_i^2}$ for all $\boldsymbol{u} \in \mathbb{K}^d$. As a consequence, the quotient $\langle \boldsymbol{x}, \boldsymbol{y}\rangle / \big(\|\boldsymbol{x}\|\|\boldsymbol{y}\|\big)$ is an element of $\mathbb{K}$. Let α be its valuation. We obtain

$$\alpha \leqslant \max_{i\in[d]}(x_i + y_i) - \left(\max_{i\in[d]} x_i + \max_{i\in[d]} y_i\right) \leqslant 0\,.$$

Suppose without loss of generality that $\langle \boldsymbol{x}, \boldsymbol{y}\rangle \neq 0$. Then, there exists a nonzero number $c \in \mathbb{R}$ such that

$$\langle \boldsymbol{x}(t), \boldsymbol{y}(t)\rangle / \big(\|\boldsymbol{x}(t)\|\|\boldsymbol{y}(t)\|\big) = ct^\alpha + o(t^\alpha)$$

when $t \to +\infty$. If $\arg\max_{i\in[d]} x_i \cap \arg\max_{i\in[d]} y_i = \emptyset$, then $\alpha < 0$, implying that the limit of $\cos \angle \boldsymbol{x}(t)\boldsymbol{y}(t) = \langle \boldsymbol{x}(t), \boldsymbol{y}(t)\rangle / \big(\|\boldsymbol{x}(t)\|\|\boldsymbol{y}(t)\|\big)$ as $t \to +\infty$ is equal to 0. ∎

We will use Lemma 22 in order to estimate the limit when $t \to +\infty$ of the angle between segments formed by triplets of successive points of the tropical central path. One remarkable property is that the tropical central path of any Puiseux linear program is monotone; see Proposition 16. We refine Lemma 22 to fit this setting.

LEMMA 23. *Let $\boldsymbol{U}, \boldsymbol{V}, \boldsymbol{W} \in \mathbb{K}^d$, $U \coloneqq \mathrm{val}(\boldsymbol{U})$, $V \coloneqq \mathrm{val}(\boldsymbol{V})$, and $W \coloneqq \mathrm{val}(\boldsymbol{W})$. If $\max_{i\in[d]} U_i < \max_{i\in[d]} V_i < \max_{i\in[d]} W_i$, and the sets $\arg\max_{i\in[d]} V_i$ and $\arg\max_{i\in[d]} W_i$ are disjoint, we have*

$$\lim_{t\to+\infty} \angle \boldsymbol{U}(t)\boldsymbol{V}(t)\boldsymbol{W}(t) = \frac{\pi}{2}\,.$$

Proof. Let us remark that for all $i \in [d]$, we have $\mathrm{val}(\boldsymbol{V}_i - \boldsymbol{U}_i) \leqslant \max(U_i, V_i)$, and this inequality is an equality if $U_i \neq V_i$. Since $\max_{i\in[d]} U_i < \max_{i\in[d]} V_i$, we deduce that $\max_{i\in[d]} \mathrm{val}(\boldsymbol{V}_i - \boldsymbol{U}_i) = \max_{i\in[d]} V_i$, and that the arguments of the two maxima are equal. The same applies to the coordinates of the vector $\mathrm{val}(\boldsymbol{W} - \boldsymbol{V})$. We infer from Lemma 22 that $\angle \boldsymbol{U}(t)\boldsymbol{V}(t)\boldsymbol{W}(t)$ tends to $\pi/2$ whenever $\arg\max_{i\in[d]} V_i \cap \arg\max_{i\in[d]} W_i = \emptyset$. ∎

This motivates us to introduce a *(weak) tropical angle*

$$\angle^* UVW \coloneqq \begin{cases} \frac{\pi}{2} & \text{if } U, V, W \text{ satisfy the conditions of Lemma 23}\,, \\ 0 & \text{otherwise} \end{cases}$$

for any three points $U, V, W \in \mathbb{T}^d$. We now consider a Puiseux linear program of the form $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, together with the associated family of linear programs $\mathrm{LP}(\boldsymbol{A}(t), \boldsymbol{b}(t), \boldsymbol{c}(t))$ and their primal-dual central path $\mathcal{C}_t$. We obtain the following.

PROPOSITION 24. *Let $\underline{\lambda}, \overline{\lambda} \in \mathbb{R}$ and $\lambda_0 = \underline{\lambda} < \lambda_1 < \cdots < \lambda_{p-1} < \lambda_p = \overline{\lambda}$. Then*

$$\liminf_{t\to\infty} \kappa\Big(\mathcal{C}_t, \big[t^{\underline{\lambda}}, t^{\overline{\lambda}}\big]\Big) \geqslant \sum_{k=1}^{p-1} \angle^* \mathcal{C}^{\mathrm{trop}}(\lambda_{k-1}) \mathcal{C}^{\mathrm{trop}}(\lambda_k) \mathcal{C}^{\mathrm{trop}}(\lambda_{k+1}) \,.$$

Proof. Let $\lambda \in \mathbb{R}$, and let $\boldsymbol{U}$ be the point of the central path of $\mathbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ with parameter $\boldsymbol{\mu}$ equal to the Puiseux series t^λ. If t is substituted by a sufficiently large real number, the points $\mathcal{C}_t(t^\lambda)$ and $\boldsymbol{U}(t)$ are identical, since both satisfy the constraints given in (4) for $A = \boldsymbol{A}(t)$, $b = \boldsymbol{b}(t)$, $c = \boldsymbol{c}(t)$, and $\mu = t^\lambda$.

The monotonicity of the tropical central path shown in Proposition 16 allows us to apply Lemma 23. From the previous discussion, we get that

$$\lim_{t\to\infty} \angle \mathcal{C}_t(t^{\lambda_{k-1}}) \mathcal{C}_t(t^{\lambda_k}) \mathcal{C}_t(t^{\lambda_{k+1}}) \geqslant \angle^* \mathcal{C}^{\mathrm{trop}}(\lambda_{k-1}) \mathcal{C}^{\mathrm{trop}}(\lambda_k) \mathcal{C}^{\mathrm{trop}}(\lambda_{k+1})$$

for all $k \in [p-1]$. Since the total curvature can be approximated from below by measuring angles of polygonal paths, we obtain $\kappa\big(\mathcal{C}_t, [t^{\underline{\lambda}}, t^{\overline{\lambda}}]\big) \geqslant \sum_{k=1}^{p-1} \angle \mathcal{C}_t(t^{\lambda_{k-1}}) \mathcal{C}_t(t^{\lambda_k}) \mathcal{C}_t(t^{\lambda_{k+1}})$. □

We denote by $\mathbf{LW}_r^=(t)$ and $\mathbf{DUALLW}_r^=(t)$ the linear programs over $\mathbb{R}$ obtained by substituting the parameter t with a real value in the Puiseux linear programs $\mathbf{LW}_r^=$ and $\mathbf{DUALLW}_r^=$, respectively. We are now ready to state and prove the following detailed version of Theorem B.

THEOREM 25. *Assume that $r \geqslant 2$. For all $\epsilon > 0$, the total curvature of the primal central path of the linear program $\mathbf{LW}_r^=(t)$ is greater than $(2^{r-2} - 1)\frac{\pi}{2} - \epsilon$, provided that $t > 1$ is sufficiently large. Moreover, the same holds for the primal-dual central path.*

Proof. We first focus on the primal-dual central path $\mathcal{C}_t$ of $\mathbf{LW}_r^=(t)$. We will use Proposition 24 to provide a lower bound on $\liminf_{t\to\infty} \kappa(\mathcal{C}_t, [t^0, t^2])$ by considering the subdivision of the closed interval $[0, 2]$ by the scalars $\lambda_k = \frac{4k}{2^{r-1}}$ for $k = 0, 1, \ldots, 2^{r-2}$.

Given $\lambda \in [0, 2]$, all the dual components of the point $\mathcal{C}^{\mathrm{trop}}(\lambda)$ of the tropical central path are less than or equal to $\max(0, \lambda - 1)$. This is a consequence of the identity (23) and the fact that all the primal components are greater than or equal to $\min(1, \lambda)$; see Propositions 20 and 21. Using Table 1, we deduce that the maximal component of the vector $\mathcal{C}^{\mathrm{trop}}(\lambda_k)$ is equal to $r - 1 + \frac{2k+2}{2^{r-1}}$, and that it is uniquely attained by the coordinate $w^{\lambda_k}_{3(r-1)}$ when k is odd, and by $w^{\lambda_k}_{3(r-1)+1}$ when k is even. This implies $\angle^* \mathcal{C}^{\mathrm{trop}}(\lambda_{k-1}) \mathcal{C}^{\mathrm{trop}}(\lambda_k) \mathcal{C}^{\mathrm{trop}}(\lambda_{k+1}) = \frac{\pi}{2}$. We deduce that the limit inferior of $\kappa(\mathcal{C}_t, [t^0, t^2])$ is greater than or equal to $(2^{r-2} - 1)\frac{\pi}{2}$.

We now deal with the primal central path. Similarly to Proposition 24, we can show that the limit inferior of the total curvature of the primal central path over $[t^0, t^2]$ is bounded from below by the sum of the tropical angles $\angle^*(x^{\lambda_{k-1}}, w^{\lambda_{k-1}})(x^{\lambda_k}, w^{\lambda_k})(x^{\lambda_{k+1}}, w^{\lambda_{k+1}})$. Using the arguments of the previous paragraph, we know that every angle is equal to $\pi/2$. We obtain the expected lower bound. □

Remark 26. One can refine Theorem 25 to additionally obtain a lower bound on the curvature of the dual central path at the same time. This requires considering a slightly modified version of $\mathbf{LW}_r^=(t)$. More precisely, it can be shown that it suffices to add the constraints $x_{2r+1} + w_{3r} = \frac{1}{t^r} x_{2r-1}$ and $x_{2r+2} + w_{3r+1} = \frac{1}{t^r} x_{2r}$ involving the two extra variables x_{2r+1}, x_{2r+2} and the slack variables w_{3r} and w_{3r+1}.

Remark 27. Let us compare the lower bound of Theorem 25 with the upper bound of Dedieu, Malajovich, and Shub [DMS05] obtained from averaging. Given a real $m \times n$

matrix A, vectors $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$, and an $m \times m$ diagonal matrix E with diagonal entries ± 1, we consider the linear program

$$P_E \qquad \min c^\top x,\ Ax + w = b,\ Ew \geqslant 0\,.$$

It is shown there that the sum of the total curvatures of the dual central paths of the 2^m linear programs P_E arising from the various choices of sign matrices E does not exceed

$$2\pi n \binom{m-1}{n}\,. \tag{35}$$

It can be verified that the dual linear program $\textbf{DUALLW}_r^=(t)$ is of the form P_E for $E = -I$, where I is the identity matrix, $n = 3r+1$, and $m = 5r-1$. By applying Stirling's formula to (35) we see that the sum of the total curvatures of the dual central paths of the 2^m linear programs P_E, arising from varying E, is bounded by

$$2\pi(3r+1)\binom{5r-2}{3r+1} = O\left(\sqrt{r}\left(\frac{3125}{108}\right)^r\right).$$

The lower bound of order $\Omega(2^r)$ from Theorem 25 shows that the total curvature of the dual central path of at least one of these 2^m linear programs is exponential in r.

Remark 28. The weak tropical angle $\angle^* UVW$ used in Proposition 24 yields a bound on the total curvature of nondecreasing paths. A similar approach allows one, more generally, to define a notion of tropical curvature for arbitrary paths. This should also be compared with the notion of curvature for tropical hypersurfaces introduced in [BdMR13]. We leave this for future work.

7. Tropical lower bound on the complexity of interior point methods. In this section, we derive a general lower bound on the number of iterations of interior point methods with a log-barrier. That lower bound is given by the smallest number of tropical segments needed to describe the tropical central path; see Theorem 30. Applying this result to the parametric family of linear programs $\textbf{LW}_r^=(t)$ provides a proof of Theorem A.

7.1. Approximating the tropical central path by tropical segments. We return to the general situation from section 5.1 and consider a dual pair of linear programs $\textbf{LP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ and $\textbf{DualLP}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ over Puiseux series.

Let $\epsilon > 0$. For $z \in \mathbb{R}^{2N}$ we denote by $\mathcal{B}_\infty(z;\epsilon)$ the closed d_∞-ball centered at z and with radius ϵ. Further, we fix $\underline{\lambda}, \overline{\lambda} \in \mathbb{R}$ such that $\underline{\lambda} \leqslant \overline{\lambda}$. The set

$$\mathcal{T}([\underline{\lambda}, \overline{\lambda}]; \epsilon) := \bigcup_{\underline{\lambda} \leqslant \lambda \leqslant \overline{\lambda}} \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda); \epsilon)$$

is the *tubular neighborhood* of the section $\mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])$ of the tropical central path; see Figure 5. Lemma 5 and Proposition 16 yield that the tropical central path can be described as a concatenation of finitely many tropical segments. By $\gamma\big([\underline{\lambda}, \overline{\lambda}]\big)$ we denote the smallest number of tropical segments needed to describe the section $\mathcal{C}^{\mathrm{trop}}\big([\underline{\lambda}, \overline{\lambda}]\big)$.

Let us consider the union $\mathcal{S}$ of a finite sequence of consecutive tropical segments

$$\mathcal{S} := \mathsf{tsegm}(z^0, z^1) \cup \mathsf{tsegm}(z^1, z^2) \cup \cdots \cup \mathsf{tsegm}(z^{p-1}, z^p) \qquad \text{for } z^0, \ldots, z^p \in \mathbb{T}^{2N}\,,$$

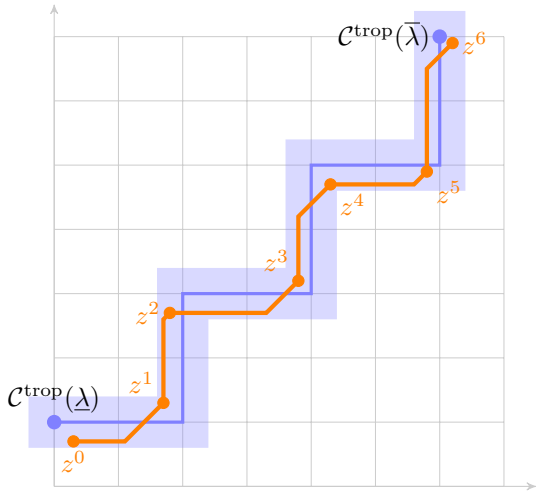


FIG. 5. *Tubular neighborhood $\mathcal{T}([\underline{\lambda},\overline{\lambda}];\epsilon)$ of the tropical central path (in light blue), containing an approximation by six tropical segments (in orange).*

which is contained in the tubular neighborhood $\mathcal{T} \coloneqq \mathcal{T}([\underline{\lambda},\overline{\lambda}];\epsilon)$, and which further satisfies $z^0 \in \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\underline{\lambda});\epsilon)$ and $z^p \in \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\overline{\lambda});\epsilon)$. That is, $\mathcal{S}$ approximates the tropical central path by p tropical segments, starting and ending in small neighborhoods of z^0 and z^p, respectively; see Figure 5 for an illustration. Next we will show that, in this situation, the number of tropical segments in $\mathcal{S}$ is bounded from below by $\gamma\big([\underline{\lambda},\overline{\lambda}]\big)$, provided that the tubular neighborhood $\mathcal{T}$ is tight enough. To this end, we set $\epsilon_0\big([\underline{\lambda},\overline{\lambda}]\big) > 0$ to one-sixteenth of the minimal d_∞-distance between any two distinct vertices in the polygonal curve $\mathcal{C}^{\mathrm{trop}}\big([\underline{\lambda},\overline{\lambda}]\big)$.

PROPOSITION 29. *If $\epsilon < \epsilon_0\big([\underline{\lambda},\overline{\lambda}]\big)$, then $p \geqslant \gamma\big([\underline{\lambda},\overline{\lambda}]\big)$.*

Proof. We use the abbreviations $\gamma \coloneqq \gamma\big([\underline{\lambda},\overline{\lambda}]\big)$ and $\epsilon_0 \coloneqq \epsilon_0\big([\underline{\lambda},\overline{\lambda}]\big)$. Let us consider a sequence $\lambda_0 = \underline{\lambda} < \lambda_1 < \cdots < \lambda_\gamma = \overline{\lambda}$ such that $\mathcal{C}^{\mathrm{trop}}([\underline{\lambda},\overline{\lambda}])$ can be decomposed as the union of γ successive tropical segments, i.e.,

$$\mathcal{C}^{\mathrm{trop}}([\underline{\lambda},\overline{\lambda}]) = \mathsf{tsegm}(\mathcal{C}^{\mathrm{trop}}(\lambda_0),\mathcal{C}^{\mathrm{trop}}(\lambda_1)) \cup \cdots \cup \mathsf{tsegm}(\mathcal{C}^{\mathrm{trop}}(\lambda_{\gamma-1}),\mathcal{C}^{\mathrm{trop}}(\lambda_\gamma))\,.$$

We claim that it suffices to show that, for all $0 < i < \gamma$, at least one point z^j belongs to the neighborhood $\mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i);8\epsilon)$ of the intermediate point $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$. Indeed, this property also trivially holds for $i = 0$ or $i = p$, by the choices of z^0 and z^p. As $\epsilon < \epsilon_0$, all these $(p+1)$ balls $\mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i);8\epsilon)$ are pairwise disjoint. We infer that $p \geqslant \gamma$ in these cases.

Consequently, we consider i such that $0 < i < \gamma$, and study the shape of the tropical central path in the neighborhood of the intermediate point $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$. Without loss of generality, we assume that $\mathcal{C}^{\mathrm{trop}}(\lambda_i) = 0$. Let $K, L \subset [2N]$ such that e^K and e^L are the left and right derivatives of $\mathcal{C}^{\mathrm{trop}}$ at λ_i, respectively. If $K \subset L$, Lemma 5 ensures that the union of $\mathsf{tsegm}(\mathcal{C}^{\mathrm{trop}}(\lambda_{i-1}),\mathcal{C}^{\mathrm{trop}}(\lambda_i))$ and $\mathsf{tsegm}(\mathcal{C}^{\mathrm{trop}}(\lambda_i),\mathcal{C}^{\mathrm{trop}}(\lambda_{i+1}))$ is a tropical segment. This contradicts the minimality of γ. Therefore, we have $K \not\subset L$. In particular, $K \neq L$, which implies that $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$ is a vertex of the polygonal curve $\mathcal{C}^{\mathrm{trop}}([\underline{\lambda},\overline{\lambda}])$. Moreover, since $\epsilon < \epsilon_0$, the point $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$ is the only vertex

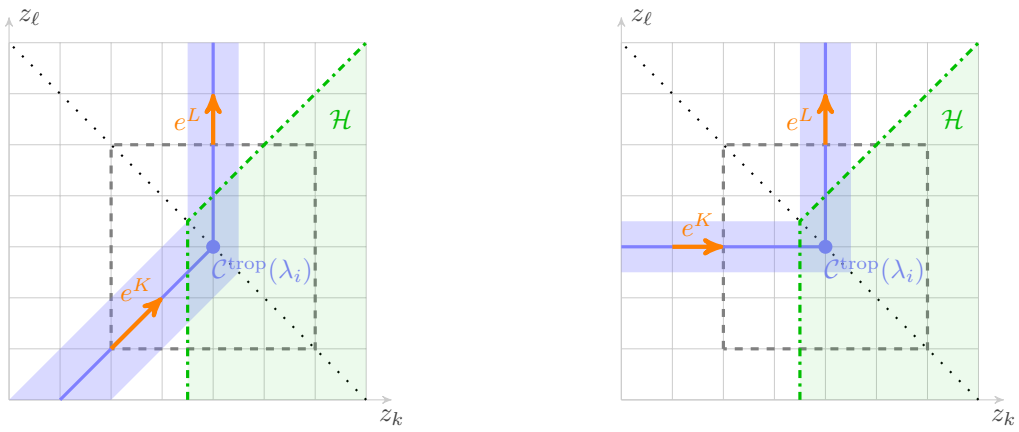


FIG. 6. *The tropical central path (in dark blue) and its tubular neighborhood $\mathcal{T}$ (in light blue) near a breakpoint $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$, projected onto the plane (z_k, z_ℓ). The tropical halfspace $\mathcal{H}$ defined in (39) is shown in green. The cases $\ell \in K$ and $\ell \notin K$ are depicted on the left and right, respectively. The dashed square delimits the ball $\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$, while the dotted line corresponds to the line $z_k + z_\ell = 0$.*

of $\mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])$ contained in the ball $\mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i); 8\epsilon)$. We derive

$$\mathcal{C}^{\mathrm{trop}}(\lambda_i + \lambda) = \begin{cases} \lambda e^K & \text{if } -8\epsilon \leqslant \lambda \leqslant 0\,, \\ \lambda e^L & \text{if } 0 \leqslant \lambda \leqslant 8\epsilon\,. \end{cases} \tag{36}$$

By Proposition 16, the set L is nonempty. We fix $k \in K \setminus L$ and $\ell \in L$, and denote by π the projection mapping $z \in \mathbb{R}^{2N}$ onto its coordinates $(z_k, z_\ell) \in \mathbb{R}^2$. We refer the reader to Figure 6 for an illustration of the projection of the central path near the breakpoint $\mathcal{C}^{\mathrm{trop}}(\lambda_i)$.

First, we will prove the following useful inclusion of sets:

$$\mathcal{T} \cap \pi^{-1}\big(\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)\big) \subset \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i); 8\epsilon)\,. \tag{37}$$

Indeed, let $z \in \mathcal{T}$ such that $\pi(z) \in \mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$. Then $z = \mathcal{C}^{\mathrm{trop}}(\lambda) + z'$ for some $\lambda \in [\underline{\lambda}, \overline{\lambda}]$ and $z' \in \mathcal{B}_\infty(0; \epsilon)$. If $\lambda \leqslant \lambda_i - 6\epsilon$, then by (36) and Corollary 17,

$$z_k = (\mathcal{C}^{\mathrm{trop}}(\lambda))_k + z'_k \leqslant (\mathcal{C}^{\mathrm{trop}}(\lambda_i - 6\epsilon))_k + \epsilon \leqslant -5\epsilon\,.$$

Similarly, for $\lambda \geqslant \lambda_i + 6\epsilon$ we can show that $z_\ell \geqslant 5\epsilon$. Summing up we have $\lambda_i - 6\epsilon \leqslant \lambda \leqslant \lambda_i + 6\epsilon$, and thus

$$d_\infty(z, \mathcal{C}^{\mathrm{trop}}(\lambda_i)) \leqslant d_\infty(z, \mathcal{C}^{\mathrm{trop}}(\lambda)) + d_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda), \mathcal{C}^{\mathrm{trop}}(\lambda_i)) \leqslant \epsilon + 6\epsilon = 7\epsilon\,.$$

Second, observe that at least one segment $\mathsf{tsegm}(z^j, z^{j+1})$ intersects the line $\{z \in \mathbb{R}^{2N} : z_k + z_\ell = 0\}$. Indeed, we have $z^0_k + z^0_\ell \leqslant (\mathcal{C}^{\mathrm{trop}}(\underline{\lambda}))_k + (\mathcal{C}^{\mathrm{trop}}(\underline{\lambda}))_\ell + 2\epsilon$. Moreover, we have $\underline{\lambda} < \lambda_i - 8\epsilon$, since if $\lambda_i - 8\epsilon \leqslant \underline{\lambda} < \lambda$, then the point $\mathcal{C}^{\mathrm{trop}}(\underline{\lambda})$ would belong to the ball $\mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i; 8\epsilon))$ by Corollary 17. We deduce that

$$z^0_k + z^0_\ell \leqslant (\mathcal{C}^{\mathrm{trop}}(\underline{\lambda}))_k + (\mathcal{C}^{\mathrm{trop}}(\underline{\lambda}))_\ell + 2\epsilon \leqslant (\mathcal{C}^{\mathrm{trop}}(\lambda_i - 8\epsilon))_k + (\mathcal{C}^{\mathrm{trop}}(\lambda_i - 8\epsilon))_\ell \leqslant -6\epsilon\,.$$

We can similarly prove that $z^p_k + z^p_\ell > 0$.

Our aim is to show that at least one extremity of the tropical segment $\mathsf{tsegm}(z^j, z^{j+1})$ belongs to $\mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i); 8\epsilon)$. By contradiction, suppose that $z^j, z^{j+1} \notin \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\lambda_i); 8\epsilon)$. Our argument is based on the projection of the segment onto the plane (z_k, z_ℓ). By (37), we know that $\pi(z^j), \pi(z^{j+1}) \notin \mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$.

Note that the intersection of $\mathsf{tsegm}(\pi(z^j), \pi(z^{j+1}))$ with $\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$ is a tropical segment $\mathsf{tsegm}(u, v)$ where $u, v \in \partial\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$. By assumption on $\mathsf{tsegm}(z^j, z^{j+1})$, the intersection of $\mathsf{tsegm}(\pi(z^j), \pi(z^{j+1}))$ with the line $\{(z_k, z_\ell) \in \mathbb{R}^2 \colon z_k + z_\ell = 0\}$ is nonempty. As the intersection of the latter line with $\pi(\mathcal{T})$ is contained in $\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$ (see Figure 6), we deduce that

(38)
$$\begin{aligned}
&\mathsf{tsegm}(u, v) \cap \{(z_k, z_\ell) \in \pi(\mathcal{T}) \colon z_k + z_\ell = 0\} \\
&\quad = \mathsf{tsegm}(\pi(z^j), \pi(z^{j+1})) \cap \mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon) \cap \{(z_k, z_\ell) \in \pi(\mathcal{T}) \colon z_k + z_\ell = 0\} \\
&\quad = \mathsf{tsegm}(\pi(z^j), \pi(z^{j+1})) \cap \{(z_k, z_\ell) \in \mathbb{R}^2 \colon z_k + z_\ell = 0\} \neq \emptyset\,.
\end{aligned}$$

Now, we introduce the tropical halfspace

$$\mathcal{H} = \{(z_k, z_\ell) \in \mathbb{T}^2 \colon \max(0, z_\ell - \epsilon) \leqslant z_k + \epsilon\}\,. \tag{39}$$

As shown in Figure 6, we have

$$\mathcal{H} \cap \pi(\mathcal{T}) \cap \mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon) \subset \mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 3\epsilon)\,.$$

Since the points u and v are in $\pi(\mathcal{T}) \cap \partial\mathcal{B}_\infty(\pi(\mathcal{C}^{\mathrm{trop}}(\lambda_i)); 4\epsilon)$, we deduce that they do not belong to $\mathcal{H}$. The complement of $\mathcal{H}$ is a tropically convex set, i.e., if $z, z' \notin \mathcal{H}$, then the tropical segment between z and z' is contained in the complement of $\mathcal{H}$. As a result, we have $\mathsf{tsegm}(u, v) \cap \mathcal{H} = \emptyset$. However, as shown in Figure 6, we have $\{(z_k, z_\ell) \in \pi(\mathcal{T}) \colon z_k + z_\ell = 0\} \subset \mathcal{H}$. We deduce that there is a contradiction with (38). ∎

We are now ready to establish a general lower bound on the number of iterations performed by the class of log-barrier interior point methods described in section 3. The result is stated in terms of polygonal curves contained in the wide neighborhood $\mathcal{N}_{\theta,t}^{-\infty}$ of the central path, since such curves are the trajectories followed by the log-barrier interior point methods. Our proof combines Theorem 19 with Proposition 29.

THEOREM 30. *For $0 < \theta < 1$ suppose that $t > t_0$ satisfies*

$$\log_t\Big(\frac{2N}{1-\theta}\Big) + \delta(t) < \epsilon_0\big([\underline{\lambda}, \overline{\lambda}]\big)\,. \tag{40}$$

Then, every polygonal curve $[z^0, z^1] \cup [z^1, z^2] \cup \cdots \cup [z^{p-1}, z^p]$ contained in the neighborhood $\mathcal{N}_{\theta,t}^{-\infty}$, with $\bar{\mu}(z^0) \leqslant t^{\underline{\lambda}}$ and $\bar{\mu}(z^p) \geqslant t^{\overline{\lambda}}$, contains at least $\gamma\big([\overline{\lambda}, \underline{\lambda}]\big)$ segments.

Proof. We first assume that $\bar{\mu}(z^0) = t^{\underline{\lambda}}$, $\bar{\mu}(z^p) = t^{\overline{\lambda}}$, and $\bar{\mu}(z^0) \leqslant \bar{\mu}(z^i) \leqslant \bar{\mu}(z^p)$ for all i. Consider $z \in [z^i, z^{i+1}]$ for some $0 \leqslant i < p$. By Proposition 1 we have $\bar{\mu}(z^0) \leqslant \bar{\mu}(z) \leqslant \bar{\mu}(z^p)$. Since, $z \in \mathcal{N}_{\theta,t}^{-\infty}$ by assumption, we deduce that $z \in \mathcal{N}_{\theta,t}^{-\infty}(\mu)$ for some μ with $\bar{\mu}(z^0) \leqslant \mu \leqslant \bar{\mu}(z^p)$.

We define $\mathcal{S}$ as the union of the tropical segments $\mathsf{tsegm}(\log_t z^i, \log_t z^{i+1})$ for

$0 \leqslant i < p$. We have

$$\begin{aligned}
d_\infty\big(\mathcal{S}, \mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])\big) &\leqslant \max_{0 \leqslant i < p} d_\infty\big(\mathsf{tsegm}(\log_t z^i, \log_t z^{i+1}), \mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])\big) \\
&\leqslant \log_t 2 + \max_{0 \leqslant i < p} \sup_{z \in [z^i, z^{i+1}]} d_\infty\big(\log_t z, \mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])\big) \qquad \text{by Lemma 8} \\
&\leqslant \log_t 2 + \sup_{t^{\underline{\lambda}} \leqslant \mu \leqslant t^{\overline{\lambda}}} d_\infty\big(\log_t \mathcal{N}_{\theta,t}^{-\infty}(\mu), \mathcal{C}^{\mathrm{trop}}([\underline{\lambda}, \overline{\lambda}])\big) \\
&\leqslant \log_t 2 + \sup_{t^{\underline{\lambda}} \leqslant \mu \leqslant t^{\overline{\lambda}}} d_\infty\big(\log_t \mathcal{N}_{\theta,t}^{-\infty}(\mu), \mathcal{C}^{\mathrm{trop}}(\log_t \mu)\big) \\
&\leqslant \log_t\Big(\frac{2N}{1-\theta}\Big) + \delta(t) \qquad \text{by Theorem 19.}
\end{aligned}$$

By choosing $\epsilon \coloneqq \log_t\big(\frac{2N}{1-\theta}\big) + \delta(t)$, this implies that $\mathcal{S}$ is contained in the tubular neighborhood $\mathcal{T}([\underline{\lambda}, \overline{\lambda}]; \epsilon)$. Moreover, $z^0 \in \mathcal{N}_{\theta,t}^{-\infty}(t^{\underline{\lambda}})$, and hence $d_\infty(z^0, \mathcal{C}^{\mathrm{trop}}(\underline{\lambda})) \leqslant \log_t\big(\frac{N}{1-\theta}\big) + \delta(t)$ by Theorem 19. We deduce that $z^0 \in \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\underline{\lambda}); \epsilon)$. An analogous argument shows that $z^p \in \mathcal{B}_\infty(\mathcal{C}^{\mathrm{trop}}(\overline{\lambda}); \epsilon)$. Since $\epsilon < \epsilon_0\big([\underline{\lambda}, \overline{\lambda}]\big)$ (see (40)), we can apply Proposition 29, which yields the claim in this special case.

We need to deal with the general case. We assume that $\underline{\lambda} < \overline{\lambda}$, since the case $\underline{\lambda} = \overline{\lambda}$ is trivial; note that $\gamma\big([\overline{\lambda}, \underline{\lambda}]\big) = 0$. Let $j \in \{0, \dots, p-1\}$ be the largest integer such that $\bar{\mu}(z^j) \leqslant t^{\underline{\lambda}}$, and let $k \in \{j+1, \dots, p\}$ be the smallest integer such that $\bar{\mu}(z^k) \geqslant t^{\overline{\lambda}}$. With this notation, we have $\bar{\mu}(z^j) \leqslant t^{\underline{\lambda}} < \bar{\mu}(z^i) < t^{\overline{\lambda}} \leqslant \bar{\mu}(z^k)$ for all $i \in \{j+1, \dots, k-1\}$. By Proposition 1, we can find $z' \in [z^j, z^{j+1}[$ such that $\bar{\mu}(z) = t^{\underline{\lambda}}$. Similarly, let $z'' \in]z^{k-1}, z^k]$ such that $\bar{\mu}(z'') = t^{\overline{\lambda}}$. Then, the polygonal curve $[z', z^{j+1}] \cup \dots \cup [z^{k-1}, z'']$ satisfies the assumptions made in the first part of the proof, in addition to being included in the neighborhood $\mathcal{N}_{\theta,t}^{-\infty}$. We deduce that $p \geqslant k - j \geqslant \gamma\big([\overline{\lambda}, \underline{\lambda}]\big)$. ∎

7.2. An exponential lower bound on the number of iterations for our main example. Let us consider the linear programs $\mathbf{LW}_r^{=}(t)$ and $\mathbf{DualLW}_r^{=}(t)$. A direct inspection reveals that choosing any $t_0 > 1$ is sufficient to meet the requirements of Lemma 18.

We focus on the section $\mathcal{C}^{\mathrm{trop}}([0,2])$ of the associated tropical central path, i.e., we consider $\underline{\lambda} = 0$ and $\overline{\lambda} = 2$. As explained in section 5.3 and illustrated in Figure 4, the projection of $\mathcal{C}^{\mathrm{trop}}\big([0,2]\big)$ onto the plane (x_{2r-1}, x_{2r}) consists of 2^{r-1} ordinary segments, which alternate their directions. This projection to two dimensions cannot be expressed as a concatenation of less than 2^{r-1} tropical segments in the plane (see Figure 3). Therefore, also the tropical central path cannot be written as the union of fewer tropical segments in any higher dimensional space. With our notation from section 7.1 this means that

$$\gamma\big([0,2]\big) \geqslant 2^{r-1}\,.$$

It remains to find t such that the condition (40) of Theorem 30 holds. It can be verified that the minimal d_∞-distance between any two vertices in $\mathcal{C}^{\mathrm{trop}}\big([0,2]\big)$ equals $1/2^{r-2}$. Therefore, we have $\epsilon_0\big([0,2]\big) = 1/(2^4 \cdot 2^{r-2})$. Every nonnull coefficient in the constraint matrix of $\mathbf{LW}_r$ is a monomial of degree in $\frac{1}{2^{r-1}}\mathbb{Z}$, and thus we may apply Theorem 12, where $\eta_0 \geqslant 1/2^{r-1}$. As a consequence, if $t \geqslant ((2N)!)^{2^{r-1}}$, we have

$$d_{\mathrm{H}}(\log_t \boldsymbol{\mathcal{F}}(t), \mathrm{val}(\boldsymbol{\mathcal{F}})) \leqslant \log_t\big((2N+1)^2((2N)!)^4\big)\,.$$

Recall that $N = 5r - 1$ is the total number of variables (including slacks), and that $\delta(t) = 2d_{\mathrm{H}}(\log_t \boldsymbol{\mathcal{F}}(t), \mathrm{val}(\boldsymbol{\mathcal{F}}))$. Now Theorem 30 specializes to the following result.

THEOREM 31. *Let $0 < \theta < 1$, and suppose that*

$$t > \left(\frac{\big((10r-1)!\big)^8}{1-\theta} \right)^{2^{r+2}} . \tag{41}$$

Then, every polygonal curve $[z^0, z^1] \cup [z^1, z^2] \cup \cdots \cup [z^{p-1}, z^p]$ contained in the neighborhood $\mathcal{N}^{-\infty}_{\theta,t}$ of the primal-dual central path of $\mathbf{LW}^{=}_r(t)$, with $\bar{\mu}(z^0) \leqslant 1$ and $\bar{\mu}(z^p) \geqslant t^2$, contains at least 2^{r-1} segments.

Taking into account the discussion at the end of section 3, we may restate Theorem 31 in terms of the complexity of interior point methods and prove Theorem A.

COROLLARY 32. *Let $0 < \theta < 1$, and suppose that t satisfies (41). Then, any log-barrier interior point method which describes a trajectory contained in the neighborhood $\mathcal{N}^{-\infty}_{\theta,t}$ of the primal-dual central path of $\mathbf{LW}^{=}_r(t)$, needs to perform at least 2^{r-1} iterations to reduce the duality measure from t^2 to 1.*

Remark 33. Corollary 32 requires the size θ of the neighborhood to be fixed independently of the parameter t. While this requirement can be relaxed slightly (the lower bound holds as soon as $\log(1-\theta) = o(\log t)$), we point out that it is met by the interior point methods discussed in Section 3. For instance, for predictor-corrector methods, the radius of the outer neighborhood is usually set to $\theta = 1/2$. Although this setting can be refined, one can show that the proof of the convergence requires θ to be chosen less than $4/5$ (see [Wri97, Exercise 5.6]).

Remark 34. Smale's 9th problem in its original version [Sma00] is about the existence of a strongly polynomial decision algorithm for the feasibility of a linear program. In the setting of interior point methods, the feasibility problem is usually tackled by solving some auxiliary linear program or linear complementarity problem involving additional constraints or variables; see, e.g., [Wri97]. It would be interesting to check that the nonstrongly polynomial behavior exhibited in this paper still appears in such auxiliary problems.

8. Combinatorial experiments. We now want to give some hints to the combinatorial properties of the feasible region of the Puiseux linear program $\mathbf{LW}_r$, which we denote as $\boldsymbol{\mathcal{R}}_r$. These are based on experiments for the first few values of r, which have been performed with `polymake` [GJ00]. Notice that, since version 3.0, `polymake` offers linear programming and convex hull computations over the field of Puiseux fractions with rational coefficients [JLLS16], and the coefficients of $\mathbf{LW}_r$ lie in this subfield. Notice that this is entirely independent of the metric analysis which was necessary for our main results. Throughout we assume that $r \geqslant 1$.

By construction $\boldsymbol{\mathcal{R}}_r$ is a convex polyhedron in the nonnegative orthant of $\mathbb{K}^{2r}$ which contains interior points, which means that it is full-dimensional. Moreover, it is easy to check that the exterior normal vectors of the defining inequalities positively span the entire space, and hence $\boldsymbol{\mathcal{R}}_r$ is bounded, i.e., a polytope over Puiseux series. None of these $3r+1$ inequalities is redundant, i.e., each inequality defines a facet. For instance, $\mathcal{R}_1$ is a quadrangle. However, the polytope $\boldsymbol{\mathcal{R}}_r$ is not *simple* for $r \geqslant 2$; i.e., there are vertices which are contained in more than $2r$ facets. All these nonsimple vertices lie in the optimal face, which is given by $x_1 = 0$. By modifying the inequality $x_{2r} \geqslant 0$ to $x_{2r} \geqslant \epsilon$ for sufficiently small $\epsilon > 0$ we obtain a simple polytope $\boldsymbol{\mathcal{R}}^{\epsilon}_r$ as the

feasible region of the *perturbed linear program*

$$\mathbf{LW}_r^\epsilon \qquad \begin{array}{rl} \text{minimize} & x_1, \\ \text{subject to} & x_1 \leqslant t^2, \\ & x_2 \leqslant t, \\ & \left.\begin{array}{l} x_{2j+1} \leqslant t\, x_{2j-1}\,,\ x_{2j+1} \leqslant t\, x_{2j} \\ x_{2j+2} \leqslant t^{1-1/2^j}(x_{2j-1} + x_{2j}) \end{array}\right] \ 1 \leqslant j < r, \\ & x_{2r-1} \geqslant 0\,,\ x_{2r} \geqslant \epsilon. \end{array}$$

For the remainder of this section we refer to our original construction $\mathbf{LW}_r$ and its feasible region as the *unperturbed case*. The unperturbed Puiseux polytope $\mathbf{LW}_r$ can be seen as the limit of the perturbed Puisuex polytopes $\mathbf{LW}_r^\epsilon$ when ϵ goes to zero. See Figure 7 for a visualization of $\boldsymbol{\mathcal{R}}_2^\epsilon$.

The two facets $x_1 = t^2$ and $x_3 = tx_2$ play a special role. It can be verified that they do not share any vertices, neither in the perturbed nor in the unperturbed case. In the unperturbed $r = 2$ case these two facets cover all the vertices except for one, while four and 16 are uncovered for $r = 3$ and $r = 4$, respectively.

Since the perturbation is only very slight, it follows that the dual graphs of the perturbed and the unperturbed Puiseux polytopes are the same. For $2 \leqslant r \leqslant 6$ we found that this is a complete graph minus one edge, which corresponds to the special pair of disjoint facets mentioned above. This should be compared with the following constructions. Let D be any dual to 2-neighborly polytope; any two facets of D share a common *ridge*, i.e., a face of codimension 2. Now pick any ridge and truncate it to produce a new polytope D'. The dual graph of D' is a complete graph minus one edge. Note that the polytope D' may have many vertices, as it is still very close to a polytope which is 2-neighborly.

In our experiments, for all $r \leqslant 6$, the primal graph of $\boldsymbol{\mathcal{R}}_r^\epsilon$ has diameter $r + 1 = (3r + 1) - 2r$, which is precisely the Hirsch bound. This is in stark contrast with the unperturbed case in which the diameter equals 3 for $2 \leqslant r \leqslant 6$.

An interesting question is, Which values for ϵ are small enough? Our experiments suggest that $\epsilon = t^{-1}$ works for all $r \geqslant 2$. Employing generalized Puiseux series with valuations of higher rank offers an alternative approach, which will always work because we may introduce a second large infinitesimal $s \gg t$ and set $\epsilon = s^{-1}$.

9. Concluding remarks. In the present work, we obtained a family of counterexamples showing that standard polynomial time interior point methods exhibit a nonstrongly polynomial time behavior. To do so, we considered non-Archimedean instances with a degenerate tropical limit that we characterized by combinatorial means. This strategy is likely to be applicable to other problems in computational complexity: tropicalization generally permits us to test the sensitivity of classical algorithms to the bit length of the input.

Interior point methods work more generally when replacing the logarithmic barrier with any self-concordant barrier functions (see [NN94] for background). There has been lately an intensive research on finding barrier functions that lead to improved complexity estimates [LS14, BE15, LY18, LS19]. A natural question is whether the present approach can be extended to these barriers. A first result in this direction is given in [AAGH20], where it is shown that the tropicalization of the central path associated with the entropic barrier is the same as the one of the logarithmic barrier. We leave it as an open problem to establish a "universality result" showing that a large class of self-concordant barrier functions are subject to the same pathology. We

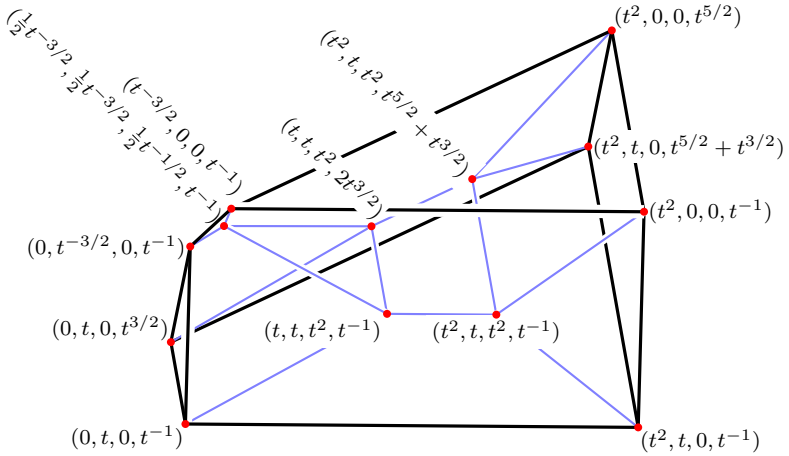


FIG. 7. *Schlegel diagram of perturbed polytope* $\boldsymbol{\mathcal{R}}_2^\epsilon$ *(for* $\epsilon = t^{-1}$*) projected onto the facet* $x_3 = 0$.

point out that the tools from model theory like the ones used in [Ale13, ABGJ14b] may help to analyze such functions.

Acknowledgments. We thank the colleagues who commented on [ABGJ14b], in particular Antoine Deza, Tamás Terlaky, and Yuriy Zinchenko. Moreover, we are indebted to Frank Sottile and two anonymous reviewers of the paper [ABGJ18] for their very useful and detailed comments.

REFERENCES

[AAGH20] X. Allamigeon, A. Aznag, S. Gaubert, and Y. Hamdi. The tropicalization of the entropic barrier. arXiv:2010.10205, October 2020.

[ABGJ14a] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Combinatorial simplex algorithms can solve mean payoff games. *SIAM J. Opt.*, 24(4):2096–2117, 2014. E-print arXiv:1309.5925.

[ABGJ14b] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Long and winding central paths. E-print arXiv:1405.4161, 2014.

[ABGJ15] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Tropicalizing the simplex algorithm. *SIAM J. Discrete Math.*, 29(2):751–795, 2015. E-print arXiv:1308.0454.

[ABGJ18] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Log-barrier interior point methods are not strongly polynomial. *SIAM Journal on Applied Algebra and Geometry*, 2(1):140–178, 2018.

[AGS18] X. Allamigeon, S. Gaubert, and M. Skomra. Solving generic nonarchimedean semidefinite programs using stochastic game algorithms. *Journal of Symbolic Computation*, 85:25 – 54, 2018. 41th International Symposium on Symbolic and Alge-braic Computation (ISSAC'16).

[AGS20] X. Allamigeon, S. Gaubert, and M. Skomra. Tropical spectrahedra. *Discrete & Computational Geometry*, 63:507–548, 2020. E-print arXiv:1610.06746.

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

[AR89] A. D. Alexandrov and Yu. G. Reshetnyak. *General theory of irregular curves*. Kluwer, 1989.

[BdMR13] B. Bertrand, L. López de Medrano, and J.-J. Risler. On the total curvature of tropical hypersurfaces. In *Algebraic and combinatorial aspects of tropical geometry*, number 589 in Contemp. Math., pages 21–43. Amer. Math. Soc., Providence, RI, 2013.

[BE15] S. Bubeck and R. Eldan. The entropic barrier: a simple and optimal universal self-concordant barrier. volume 40 of *Proceedings of Machine Learning Research*,

pages 279–279, Paris, France, 03–06 Jul 2015. PMLR.

[BEM18] M. A. Bennouna and Y. El Maazouz. Implementation of a predictor-corrector interior point method in arbitrary precision, 2018. https://github.com/yassineELMAAZOUZ/Projet-Barrieres-Exotiques.

[Ber71] G.M. Bergman. The logarithmic limit-set of an algebraic variety. *Transactions of the American Mathematical Society*, 157:459–469, 1971.

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

[BL89] D. A. Bayer and J. C. Lagarias. The nonlinear geometry of linear programming. I. Affine and projective scaling trajectories. *Trans. Amer. Math. Soc.*, 314(2):499–526, 1989.

[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.

[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, March 2006.

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

[DH20] D. Dadush and S. Huiberts. Smoothed analysis of the simplex method. In T. Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2020.

[DHNV20] D. Dadush, S. Huiberts, B. Natura, and L. A. Végh. A scaling-invariant algorithm for linear programming whose running time depends only on the constraint matrix. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 761–774, New York, NY, USA, 2020. Association for Computing Machinery.

[DL19] J. De Loera. Algebraic and topological tools in linear optimization. *Notices of the American Mathematical Society*, 66:1, 08 2019.

[DLSV12] J.A. De Loera, B. Sturmfels, and C. Vinzant. The central curve in linear programming. *Foundations of Computational Mathematics*, 12(4):509–540, 2012.

[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.

[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.

[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.

[GJ00] E. Gawrilow and M. Joswig. `polymake`: a framework for analyzing convex polytopes. In *Polytopes—combinatorics and computation (Oberwolfach, 1997)*, volume 29 of *DMV Sem.*, pages 43–73. Birkhäuser, Basel, 2000.

[GK11] S. Gaubert and R.D. Katz. Minimal half-spaces and external representation of tropical polyhedra. *Journal of Algebraic Combinatorics*, 33(3):325–348, 2011.

[HR15] G.H. Hardy and M. Riesz. *The general theory of Dirichlet's series*. Cambridge University Press, 1915.

[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. *Math. Intelligencer*, 18(4):19–28, 1996.

[Jer73] R.G. Jeroslow. Asymptotic linear programming. *Operations Research*, 21(5):1128–1141, 1973.

[JL20] M. Joswig and G. Loho. Monomial tropical cones for multicriteria optimization. *SIAM J. Discrete Math.*, 34(2):1172–1191, 2020.

[JLLS16] M. Joswig, G. Loho, B. Lorenz, and B. Schröter. Linear programs and convex hulls over fields of Puiseux fractions. In *Proceedings of MACIS 2015, Berlin, November 11–13, 2015. LNCS 9582*, pages 429–445. Springer, 2016.

[Kar84] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, 1984.

[KMY89a] M. Kojima, S. Mizuno, and A. Yoshise. A polynomial-time algorithm for a class of linear complementarity problems. *Math. Programming*, 44(1, (Ser. A)):1–26, 1989.

[KMY89b] M. Kojima, S. Mizuno, and A. Yoshise. A primal-dual interior point algorithm for linear programming. In *Progress in mathematical programming (Pacific Grove, CA, 1987)*, pages 29–47. Springer, New York, 1989.

[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.

[KS10] E.D. Kim and F. Santos. An update on the Hirsch conjecture. *Jahresber. Dtsch. Math.-Ver.*, 112(2):73–98, 2010.

[KT13] T. Kitahara and T. Tsuchiya. A simple variant of the mizuno–todd–ye predictor-corrector algorithm and its objective-function-free complexity. *SIAM Journal on Optimization*, 23(3):1890–1903, 2013.

[Lit07] G.L. Litvinov. Maslov dequantization, idempotent and tropical mathematics: a brief introduction. *Journal of Mathematical Sciences*, 140(3):426–444, 2007.

[LS14] Y. T. Lee and A. Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 424–433, 2014.

[LS19] Y. T. Lee and A. Sidford. Solving linear programs with $\tilde{O}(\sqrt{\text{rank}})$ linear system solves. arXiv:1910.08033, 2019.

[LY18] Y. T. Lee and M.-C. Yue. Universal barrier is n-self-concordant. arXiv:1809.03011, 2018.

[MA89] R.D.C. Monteiro and I. Adler. Interior path following primal-dual algorithms. part I: Linear programming. *Mathematical Programming*, 44:27–41, 1989.

[Mar10] T. Markwig. A field of generalised Puiseux series for tropical geometry. *Rend. Semin. Mat., Univ. Politec. Torino*, 68(1):79–92, 2010.

[Mik05] G. Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *J. Amer. Math. Soc.*, 18:313–377, 2005.

[MMT98] N. Megiddo, S. Mizuno, and T. Tsuchiya. A modified layered-step interior-point algorithm for linear programming. *Mathematical Programming*, 82(3):339–355, Aug 1998.

[MS15] D. Maclagan and B. Sturmfels. *Introduction to Tropical Geometry*, volume 161 of *Graduate Texts in Math.* American Math. Soc., 2015.

[MT03] R. D. C. Monteiro and T. Tsuchiya. A variant of the vavasis–ye layered-step interior-point algorithm for linear programming. *SIAM Journal on Optimization*, 13(4):1054–1079, 2003.

[MTY93] S. Mizuno, M. J. Todd, and Y. Ye. On adaptive-step primal-dual interior-point algorithms for linear programming. *Mathematics of Operations Research*, 18(4):964–981, 1993.

[NN94] Y. Nesterov and A. Nemirovskii. *Interior Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics, 1994.

[PT14] A. Papadopoulos and M. Troyanov. Weak Minkowski spaces. In *Handbook of Hilbert geometry*, volume 22 of *IRMA Lect. Math. Theor. Phys.*, pages 11–32. Eur. Math. Soc., Zürich, 2014.

[Ren88] J. Renegar. A polynomial-time algorithm, based on Newton's method, for linear programming. *Math. Programming*, 40(1, (Ser. A)):59–93, 1988.

[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.

[Sma00] S. Smale. Mathematical problems for the next century. In *Mathematics: frontiers and perspectives*, pages 271–294. Amer. Math. Soc., Providence, RI, 2000.

[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.

[Tar86] E. Tardos. A Strongly Polynomial Algorithm to Solve Combinatorial Linear Programs. *Operations Research*, 34(2):250–256, March 1986.

[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.

[vdW93] B. L. van der Waerden. *Algebra II.* Springer, Berlin, 1993. Unter Benutzung von Vorlesungen von E. Artin und E. Noether. Mit einem Geleitwort von Jürgen Neukirch, 6th edn.

[Vir01] O. Viro. Dequantization of real algebraic geometry on logarithmic paper. In *European Congress of Mathematics*, pages 135–146. Springer, 2001.

[VY96] S. A. 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, Jul 1996.

[Wri97] S. Wright. *Primal-Dual Interior-Point Methods.* Society for Industrial and Applied Mathematics, 1997.

[Ye11] Y. Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Math. Oper. Res.*, 36(4):593–603, 2011.

[Zie12] G. M. Ziegler. Who solved the Hirsch conjecture? *Documenta Mathematica*, pages 75—85, 2012. Extra Volume "Optimization Stories".

[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.