

THE CENTRAL CURVE IN LINEAR PROGRAMMING

JESÚS A. DE LOERA, BERND STURMFELS, AND CYNTHIA VINZANT

ABSTRACT. The central curve of a linear program is an algebraic curve specified by linear and quadratic constraints arising from complementary slackness. It is the union of the various central paths for minimizing or maximizing the cost function over any region in the associated hyperplane arrangement. We determine the degree, arithmetic genus and defining prime ideal of the central curve, thereby answering a question of Bayer and Lagarias. These invariants, along with the degree of the Gauss image of the curve, are expressed in terms of the matroid of the input matrix. Extending work of Dedieu, Malajovich and Shub, this yields an instance-specific bound on the total curvature of the central path, a quantity relevant for interior point methods. The global geometry of central curves is studied in detail.

1. INTRODUCTION

We consider the standard linear programming problem in its *primal* and *dual* formulation:

$$(1) \quad \text{Maximize } \mathbf{c}^T \mathbf{x} \text{ subject to } A\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0;$$

$$(2) \quad \text{Minimize } \mathbf{b}^T \mathbf{y} \text{ subject to } A^T \mathbf{y} - \mathbf{s} = \mathbf{c} \text{ and } \mathbf{s} \geq 0.$$

Here A is a fixed matrix of rank d having n columns. The vectors $\mathbf{c} \in \mathbb{R}^n$ and $\mathbf{b} \in \text{image}(A)$ may vary. In most of our results we assume that \mathbf{b} and \mathbf{c} are generic. This implies that both the primal optimal solution and the dual optimal solution are unique. We also assume that the problem (1) is bounded and both problems (1) and (2) are strictly feasible.

Before describing our contributions, we review some basics from the theory of linear programming [26, 33]. The (primal) *logarithmic barrier function* for (1) is defined as

$$f_\lambda(\mathbf{x}) := \mathbf{c}^T \mathbf{x} + \lambda \sum_{i=1}^n \log x_i,$$

where $\lambda > 0$ is a real parameter. This specifies a family of optimization problems:

$$(3) \quad \text{Maximize } f_\lambda(\mathbf{x}) \text{ subject to } A\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0.$$

Since the function f_λ is strictly concave, it attains a unique maximum $\mathbf{x}^*(\lambda)$ in the interior of the feasible polytope $P = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^n : A\mathbf{x} = \mathbf{b}\}$. Note that $f_\lambda(\mathbf{x})$ tends to $-\infty$ when \mathbf{x} approaches the boundary of P . The *primal central path* is the curve $\{\mathbf{x}^*(\lambda) \mid \lambda > 0\}$ inside the polytope P . There is an analogous logarithmic barrier function for the dual problem (2) and a corresponding *dual central path*. The central path connects the optimal solution of the linear program in question with its *analytic center*. This is the optimal point of f_∞ or equivalently $\text{argmax}_P(\sum_{i=1}^n \log x_i)$. The central path is homeomorphic to a line segment.

1991 *Mathematics Subject Classification*. Primary: 90C05; Secondary: 05B35, 13P25, 14H45, 52C35.

Key words and phrases. Linear programming, interior point methods, matroid, Tutte polynomial, hyperbolic polynomial, Gauss map, degree, curvature, projective variety, Gröbner basis, hyperplane arrangement.




FIGURE 1. The central curve of six lines for two choices of the cost function

The *complementary slackness* condition says that the pair of optimal solutions, to the primal linear program (1) and to the dual linear program (2), are characterized by

$$(4) \quad \mathbf{Ax} = \mathbf{b}, \quad \mathbf{A}^T \mathbf{y} - \mathbf{s} = \mathbf{c}, \quad \mathbf{x} \geq 0, \quad \mathbf{s} \geq 0, \quad \text{and } x_i s_i = 0 \text{ for } i = 1, 2, \dots, n.$$

The central path converges to the solution of this system of equations and inequalities:


Theorem 1 (cf. [33]). *If A has d rows, then for all $\lambda > 0$, the system of polynomial equations*

$$(5) \quad \mathbf{Ax} = \mathbf{b}, \quad \mathbf{A}^T \mathbf{y} - \mathbf{s} = \mathbf{c}, \quad \text{and } x_i s_i = \lambda \text{ for } i = 1, 2, \dots, n,$$

has a unique real solution $(\mathbf{x}^(\lambda), \mathbf{y}^*(\lambda), \mathbf{s}^*(\lambda))$ with the properties $\mathbf{x}^*(\lambda) > 0$ and $\mathbf{s}^*(\lambda) > 0$. The point $\mathbf{x}^*(\lambda)$ is the optimal solution of (3). The limit point $(\mathbf{x}^*(0), \mathbf{y}^*(0), \mathbf{s}^*(0))$ of these solutions for $\lambda \rightarrow 0$ is the unique solution of the complementary slackness constraints (4).*

Our object of study in this paper is the set of *all* solutions of the equations (5), not just those whose coordinates are real and positive. For general \mathbf{b} and \mathbf{c} , this set is the following irreducible algebraic curve. The *central curve* is the Zariski closure of the central path in $(\mathbf{x}, \mathbf{y}, \mathbf{s})$ -space, that is, it is the smallest algebraic variety in \mathbb{R}^{2n+d} that contains the central path. The *primal central curve* in \mathbb{R}^n is obtained by projecting the central curve into \mathbf{x} -space. We can similarly define the *dual central curve* by projecting into \mathbf{y} -space or into \mathbf{s} -space.

Example 2. Figure 1 depicts the primal central curve for a small *transportation problem*. Here A is the 5×6 node-edge matrix of the complete bipartite graph $K_{2,3}$, as shown below:



Here $n = 6$ and $d = 4$ because A has rank 4. We return to this example in Section 4. \diamond

As seen in Figure 1, and proved in Theorem 33, the primal central curve contains the central paths of every polytope in the arrangement in $\{A\mathbf{x} = \mathbf{b}\}$ defined by the coordinate hyperplanes $\{x_i = 0\}$ for the cost functions \mathbf{c} and $-\mathbf{c}$. The union over all central curves, as the right hand side \mathbf{b} varies, is an algebraic variety of dimension $d+1$, called the *central sheet*, which will play an important role. Our analysis will be based on results of Terao [31] and Proudfoot–Speyer [24] on algebras generated by reciprocals of linear forms; see also Berget [6]. Matroid theory will be our language for working with these algebras and their ideals.

The algebro-geometric study of central paths was pioneered by Bayer and Lagarias [3, 4]. Their 1989 articles are part of the early history of interior point methods. They observed (on pages 569-571 of [4]) that the central path defines an irreducible algebraic curve in \mathbf{x} -space or \mathbf{y} -space, and they identified a complete intersection that has the central curve as an irreducible component. The last sentence of [4, §11] states the open problem of identifying polynomials that cut out the central curve, without any extraneous components. It is worth stressing that one easily finds polynomials that vanish on the central curve from the gradient optimality conditions on the barrier function. Those polynomials vanish on high-dimensional components, other than the central curve. These extra components are contained in the coordinate hyperplanes, and the challenge is to remove them in our algebraic description.

In numerical optimization, the optimal solution to (1) is found by following a piecewise-linear approximation to the central path. Different strategies for generating the step-by-step moves correspond to different interior point methods. One way to estimate the number of Newton steps needed to reach the optimal solution is to bound the *total curvature* of the central path. This has been investigated by many authors (see e.g. [11, 22, 28, 34, 37]), the idea being that curves with small curvature are easier to approximate with line segments. The algebraic results in this paper contribute to the understanding of the total curvature.

Here is a list of our results. Precise statements are given in each section.

- Section 2 analyzes central curves in the plane, with emphasis on the dual formulation ($d = 2$). We show that our curves are *Vinnikov curves* [35] of degree $\leq n - 1$, obtained from an arrangement of n lines by taking a *Renegar derivative* [25]. The total curvature of a plane curve can be bounded in terms of its number of real inflection points. We derive a new bound from a classical formula due to Felix Klein [18].
- All our formulas and bounds in Sections 4, 5, 6, and 7 are expressed in the language of matroid theory. A particularly important role is played by matroid invariants, such as the *Tutte polynomial*, that are associated with the matrix A . In Section 3 we review the required background from matroid theory and geometric combinatorics.
- In Section 4 we present a complete solution to the Bayer-Lagarias problem. Under the assumption that \mathbf{b} and \mathbf{c} are general, while A is fixed and possibly special, we determine the prime ideal of all polynomials that vanish on the primal central curve. We express the degree of this curve as a matroid invariant. This yields the tight upper bound $\binom{n-1}{d}$ for the degree. For instance, the curves in Figure 1 have degree five. We also determine the Hilbert series and arithmetic genus of our curve in \mathbb{P}^n .
- Section 5 develops our approach to estimating the total curvature of the central curve. Dedieu, Malajovich and Shub [11] noted that the total curvature of any curve \mathcal{C} coincides with the arc length of the image of \mathcal{C} under the *Gauss map*. Hence any bound on the degree of the *Gauss curve* translates into a bound on the total




FIGURE 2. The DTZ snake with 6 constraints. On the left, a global view of the polygon and its central curve with the line $y_2 = 1$ appearing as part of the curve. On the right a close-up of the central path and its inflection points.

2. PLANE CURVES

When the central curve lives in a plane, the curve is cut out by a single polynomial equation. This occurs for the dual curve when $d = 2$ and the primal curve when $n = d - 2$. We now focus on the dual curve ($d = 2$). This serves as a warm-up to the full derivation of all equations in Section 4. In this section we derive the equations of the central curve from first principles, we show that these curves are hyperbolic and Renegar derivatives of products of lines, and we use this structure to bound the average total curvature of the curve.

Let $A = (a_{ij})$ be a fixed $2 \times n$ matrix of rank 2, and consider arbitrary vectors $\mathbf{b} = (b_1, b_2)^T \in \mathbb{R}^2$ and $\mathbf{c} = (c_1, \dots, c_n)^T \in \mathbb{R}^n$. Here the \mathbf{y} -space is the plane \mathbb{R}^2 with coordinates $\mathbf{y} = (y_1, y_2)$. The central curve is the Zariski closure in this plane of the parametrized path

$$\mathbf{y}^*(\lambda) = \underset{\{\mathbf{y} : A^T \mathbf{y} \geq \mathbf{c}\}}{\operatorname{argmin}} \quad b_1 y_1 + b_2 y_2 - \lambda \sum_{i=1}^n \log(a_{1i} y_1 + a_{2i} y_2 - c_i).$$

The conditions for optimality are obtained by setting the first partial derivatives to zero:

$$0 = b_1 - \lambda \sum_{i=1}^n \frac{a_{1i}}{a_{1i} y_1 + a_{2i} y_2 - c_i} \quad \text{and} \quad 0 = b_2 - \lambda \sum_{i=1}^n \frac{a_{2i}}{a_{1i} y_1 + a_{2i} y_2 - c_i}.$$

Multiplying these equations by b_2/λ or b_1/λ gives

$$(6) \quad \frac{b_1 b_2}{\lambda} = \sum_{i=1}^n \frac{b_2 a_{1i}}{a_{1i} y_1 + a_{2i} y_2 - c_i} = \sum_{i=1}^n \frac{b_1 a_{2i}}{a_{1i} y_1 + a_{2i} y_2 - c_i}.$$

This eliminates the parameter λ and we are left with the equation on the right. By clearing denominators, we get a single polynomial C that vanishes on the central curve in \mathbf{y} -space:

$$(7) \quad C(\mathbf{y}) = \sum_{i \in \mathcal{I}} (b_1 a_{2i} - b_2 a_{1i}) \prod_{j \in \mathcal{I} \setminus \{i\}} (a_{1j} y_1 + a_{2j} y_2 - c_j),$$




FIGURE 3. The degree-6 central path of a planar 7-gon in the affine charts $\{y_0 = 1\}$ and $\{y_2 = 1\}$. Every line passing through $[0 : -b_2 : b_1]$ intersects the curve in 6 real points, showing the real curve to be 3 completely-nested ovals.

where $\mathcal{I} = \{i : b_1 a_{2i} - b_2 a_{1i} \neq 0\}$. We see that the degree of $C(\mathbf{y})$ is $|\mathcal{I}| - 1$. This equals $n - 1$ for generic \mathbf{b} . In our derivation we assumed that λ is non-zero but the resulting equation is valid on the Zariski closure, which includes the important points with parameter $\lambda = 0$.

We consider the closure \mathcal{C} of the central curve in the complex projective plane \mathbb{P}^2 with coordinates $[y_0 : y_1 : y_2]$. Thus \mathcal{C} is the complex projective curve defined by $y_0^{|\mathcal{I}|-1} C(\frac{y_1}{y_0}, \frac{y_2}{y_0})$.

Proposition 4. *The curve \mathcal{C} is hyperbolic with respect to the point $[0 : -b_2 : b_1]$. This means that every line in $\mathbb{P}^2(\mathbb{R})$ passing through this special point meets \mathcal{C} only in real points.*

Proof. Any line passing through the point $[0 : -b_2 : b_1]$ (except the line $\{y_0 = 0\}$) has the form $\{b_1 y_1 + b_2 y_2 = b_0 y_0\}$ for some $b_0 \in \mathbb{R}$. See the left picture in Figure 3. We shall see in Remark 36 that, for any $b_0 \in \mathbb{R}$, the line meets \mathcal{C} in $\deg(\mathcal{C})$ real points. \square

Hyperbolic curves are also known as *Vinnikov curves*, in light of Vinnikov's seminal work [20, 35] relating them to semidefinite programming [27]. Semidefinite programming has been generalized to hyperbolic programming, in the work of Renegar [25] and others. A key construction in hyperbolic programming is the Renegar derivative which creates a (hyperbolic) polynomial of degree $D - 1$ from any (hyperbolic) polynomial of degree D . To be precise, the *Renegar derivative* of a homogeneous polynomial f with respect to a point \mathbf{e} is

$$R_{\mathbf{e}}f(\mathbf{y}) = \left(\frac{\partial}{\partial t} f(\mathbf{y} + t\mathbf{e}) \right) \Big|_{t=0}.$$

Renegar derivatives correspond to the *polar curves* of classical algebraic geometry [14, §1.1].

The Renegar derivative of $f = \prod_{i \in \mathcal{I}} (a_{1i} y_1 + a_{2i} y_2 - c_i y_0)$ with $\mathbf{e} = (0, -b_2, b_1)$ is seen to be

$$(8) \quad R_{\mathbf{e}}f(\mathbf{y}) = \sum_{i \in \mathcal{I}} (b_1 a_{2i} - b_2 a_{1i}) \prod_{j \in \mathcal{I} \setminus \{i\}} (a_{1j} y_1 + a_{2j} y_2 - c_j y_0) = C(\mathbf{y}).$$

In words: the central curve \mathcal{C} is the Renegar derivative, taken with respect to the cost function, of the product of the linear forms that define the convex polygon of feasible points.

The product of linear forms $f = \prod_i (a_{1i}y_1 + a_{2i}y_2 - c_i y_0)$ is a hyperbolic polynomial with respect to \mathbf{e} . Renegar [25] shows that if f is hyperbolic with respect to \mathbf{e} then so is $R_{\mathbf{e}}f$. This yields a second proof for Proposition 4.

Proposition 4 is visualized in Figure 3. The picture on the right is obtained from the picture on the left by a projective transformation. The point at infinity which represents the cost function is now in the center of the diagram. In this rendition, the central curve consists of three nested ovals around that point, highlighting the salient features of a Vinnikov curve. This beautiful geometry is found not just in the dual picture but also in the primal picture:

Remark 5. If $d = n - 2$ then the primal central curve lies in the plane $\{\mathbf{A}\mathbf{x} = \mathbf{b}\}$. The conditions for optimality of (1) state that the vector $\nabla(\sum_i \log x_i) = (x_1^{-1}, \dots, x_n^{-1})$ is in the span of \mathbf{c} and the rows of A . The Zariski closure of such vectors is the *central sheet*, to be seen in Section 4. Here, the central sheet is the hypersurface in \mathbb{R}^n with defining polynomial

$$(9) \quad \det \begin{pmatrix} A_1 & A_2 & \cdots & A_n \\ c_1 & c_2 & \cdots & c_n \\ x_1^{-1} & x_2^{-1} & \cdots & x_n^{-1} \end{pmatrix} \cdot \prod_{i \in \mathcal{I}} x_i,$$

where A_i is the i th column of A and $\mathcal{I} = \{i : \{(A_j)_{c_j}\}_{j \in [n] \setminus i} \text{ are linearly independent}\}$. We see that the degree of this hypersurface is $|\mathcal{I}| - 1$, so it is $n - 1$ for generic A . Intersecting this surface with the plane $\{\mathbf{A}\mathbf{x} = \mathbf{b}\}$ gives the primal central curve, which is hence a curve of degree $|\mathcal{I}| - 1$. The corresponding complex projective curve in $\mathbb{P}^2 = \{[x_0 : \mathbf{x}] \mid \mathbf{A}\mathbf{x} = x_0 \mathbf{b}\} \subset \mathbb{P}^n$ is hyperbolic with respect to the point $[0 : \mathbf{v}]$ in \mathbb{P}^n , where \mathbf{v} spans the kernel of $\begin{pmatrix} A \\ \mathbf{c} \end{pmatrix}$. \diamond

It is of importance for interior point algorithms to know the exact total curvature, formally introduced in equation (24), of the central path of a linear program (see [11, 22, 28, 34, 37]). Deza *et al.* [13] proved that even for $d = 2$ the total curvature grows linearly in n , and they conjectured that the total curvature is no more than $2\pi n$. They named this conjecture the *continuous Hirsch conjecture* because of its similarity with the discrete simplex method analogue (see [12]). In Section 5 we derive general bounds for total curvature, but for plane curves we can exploit an additional geometric feature, namely, *inflection points*.

Benedetti and Dedò [5] derived a general bound for the total curvature of a real plane curve in terms of its number of inflection points and its degree. We can make this very explicit for our central path $\{\mathbf{y}^*(\lambda) : \lambda \in \mathbb{R}_{\geq 0}\}$. Its total curvature is bounded above by

$$(10) \quad \text{total curvature of the central path} \leq \pi \cdot (\text{its number of inflection points} + 1).$$

To see this, consider the Gauss map γ that takes the curve into the circle S^1 by mapping any point on the curve to its unit tangent vector. See Section 5 for the general definition. The total curvature is the arc length of the image of the Gauss map. As λ decreases from ∞ to 0, the cost function $\mathbf{b}^T \mathbf{y}^*(\lambda)$ strictly decreases. This implies that, for any point $\mathbf{y}^*(\lambda)$ on the curve, its image under the Gauss map has positive inner product with \mathbf{b} , that is, $\mathbf{b}^T \gamma(\mathbf{y}^*(\lambda)) \geq 0$. Thus the image of the Gauss map is restricted to a half circle of S^1 , and it cannot wrap around S^1 . This shows that the Gauss map can achieve a length of at most π before it must “change direction”, which happens only at inflection points of the curve.

It is known that the total number of (complex) inflection points of a plane curve of degree D is at most $3D(D - 2)$. For real inflection points, there is an even better bound:

Proposition 6 (A classical result of Felix Klein [18]).

The number of real inflection points of a plane curve of degree D is at most $D(D - 2)$.

This provides only a quadratic bound for the total curvature of the central path in terms of its degree, but it does allow us to improve known bounds for the average total curvature. The *average total curvature* of the central curve of a hyperplane arrangement is the average, over all bounded regions of the arrangement, of the total curvature of the central curve in that region. Dedieu *et al.* [11] proved that the average total curvature in a simple arrangement (*i.e.* for a generic matrix A) defined by n hyperplanes in dimension d is not greater than $2\pi d$. When $d = 2$, we can use Proposition 6 to improve this bound by a factor of two. See Theorem 27 for the extension to general d .

Theorem 7. *The average total curvature of a central path of a generic line arrangement in the plane is at most 2π .*

Proof. The central curve for n general lines in \mathbb{R}^2 has degree $n - 1$ and consists of $n - 1$ (real affine) connected components. The argument above and Klein's theorem then show that

$$\begin{aligned} \sum_{i=1}^{n-1} (\text{curvature of the } i\text{th component}) &\leq \sum_{i=1}^{n-1} \pi(\#\text{inflection points on the } i\text{th component} + 1) \\ &\leq \pi(n-1)(n-2). \end{aligned}$$

Our arrangement of n general lines has $\binom{n-1}{2}$ bounded regions. The average total curvature over each of these regions is therefore at most $\pi(n-1)(n-2)/\binom{n-1}{2} = 2\pi$. \square

To bound the curvature of just the central path, we need to bound the number of inflection points appearing on that piece of the central curve. To address this issue, we posed the following problem in the manuscript version of this article: *What is the largest number of inflection points on a single oval of a hyperbolic curve of degree D in the real plane? In particular, is this number linear in the degree D ?* These questions have since been answered by Erwan Brugallé and Lucía López de Medrano, using an extension of their techniques in [9]. They constructed a Vinnikov curve of even degree D which has the maximal number $D(D-2)$ of inflection points and all of these inflection points lie on the outermost oval. It would be very interesting to see whether their approach can be applied to improve the DTZ snakes of Example 3 and lead to new lower bounds for the total curvature of planar central paths.

3. CONCEPTS FROM MATROID THEORY

We have seen in the previous section that the geometry of a central curve in the plane is intimately connected to that of the underlying arrangement of constraint lines. For instance, the degree of the central curve, $|\mathcal{I}| - 1$, is one less than the number of constraints not parallel to the cost function. The systematic study of this kind of combinatorial information, encoded in a geometric configuration of vectors or hyperplanes, is the subject of *matroid theory*.

Matroid theory will be crucial for stating and proving our results in the rest of this paper. This section offers an exposition of the relevant concepts. Of course, there is already a well-established connection between matroid theory and linear optimization (e.g., as outlined in [19] or in oriented matroid programming [2]). Our paper sets up yet another connection. The material that follows is well-known in algebraic combinatorics, but less so in optimization, so we aim to cover the basic facts. The missing details can be found in [7, 10].

We consider an r -dimensional linear subspace \mathcal{L} of the vector space K^n with its fixed standard basis. Here K is any field. Typically, \mathcal{L} will be given to us as the row space of an

$r \times n$ -matrix. The kernel of that matrix is denoted by \mathcal{L}^\perp . This is a subspace of dimension $n - r$ in K^n . We write x_1, \dots, x_n for the restriction of the standard coordinates on K^n to \mathcal{L} .

The two subspaces \mathcal{L} and \mathcal{L}^\perp specify a dual pair of matroids, denoted $M(\mathcal{L})$ and $M(\mathcal{L}^\perp)$, on the set $[n] = \{1, \dots, n\}$. The matroid $M(\mathcal{L})$ has rank r and its dual $M(\mathcal{L}^\perp) = M(\mathcal{L})^*$ has rank $n - r$. We now define the first matroid $M = M(\mathcal{L})$ by way of its *independent sets*. A subset I of $[n]$ is *independent* in M if the linear forms in $\{x_i : i \in I\}$ are linearly independent on \mathcal{L} . Maximal independent sets are called *bases*. These all have cardinality r . A subset I is *dependent* if it is not independent. It is a *circuit* if it is minimally dependent.

Example 8. Consider the linear space \mathcal{L} spanned by the rows of the rank 4 matrix A in Example 2. Because the first four columns of A are linearly independent, the linear forms $\{x_1, x_2, x_3, x_4\}$ are linearly independent on \mathcal{L} and $\{1, 2, 3, 4\}$ is an independent set of $M(\mathcal{L})$. As \mathcal{L} has dimension four, $\{1, 2, 3, 4\}$ is a basis of the matroid $M(\mathcal{L})$. On the other hand, the set of columns $\{1, 2, 4, 5\}$ is linearly dependent but every proper subset is linearly independent. Hence $\{1, 2, 4, 5\}$ is a circuit of $M(\mathcal{L})$. From similar considerations, we find that $M(\mathcal{L})$ has nine circuits, namely $\{i, j, k, l\}$ where $i, j \in \{1, 2, 3\}$ and $k, l \in \{4, 5, 6\}$. \diamond

One matroid application of importance for our study of central curves is the following formula for number of bounded components of a hyperplane arrangement. Let \mathbf{u} be a generic vector in \mathbb{R}^n and consider the $(n - r)$ -dimensional affine space $\mathcal{L}^\perp + \mathbf{u}$ of \mathbb{R}^n . The equations $x_i = 0$ define n hyperplanes in this affine space. The arrangement $\{x_i = 0\}_{i \in [n]}$ in $\mathcal{L}^\perp + \mathbf{u}$ is *simple*, which means that no point lies on more than $n - r$ of the n hyperplanes. The vertices of this hyperplane arrangement are in bijection with the bases of the matroid M . The complements of the hyperplanes are convex polyhedra; they are the *regions* of the arrangement. Each region is either bounded or unbounded, and we are interested in the bounded regions. These bounded regions are the feasibility regions for the linear programs with various sign restrictions on the variables x_i . Proposition 6.6.2 in [10], which is based on results of Zaslavsky [36], equates the number of such regions with the absolute value of the *Möbius invariant* $\mu(M)$ of the matroid of M :

$$(11) \quad |\mu(M)| = \# \text{ bounded regions of the hyperplane arrangement } \{x_i = 0\}_{i \in [n]} \text{ in } \mathcal{L}^\perp + \mathbf{u}.$$

Further below, in Equation (13), this invariant will be expressed in terms of the matroid M . We refer to $|\mu(M)|$ as the *Möbius number* of the matroid M .

To obtain the Möbius number and more refined invariants that we will need in Section 5 we introduce a simplicial complex associated to the matroid M called the *broken circuit complex*. We fix the standard ordering $1 < 2 < \dots < n$ of $[n]$. A *broken circuit* of M is any subset of $[n]$ of the form $C \setminus \{\min(C)\}$ where C is a circuit. The *broken circuit complex* of M is the simplicial complex $\text{Br}(M)$ whose minimal non-faces are the broken circuits. Hence, a subset of $[n]$ is a face of $\text{Br}(M)$ if it does not contain any broken circuit. It is known that $\text{Br}(M)$ is a shellable simplicial complex of dimension $r - 1$ (see Theorem 7.4.3 in [7]). We can recover the Möbius number of M as follows. Let $f_i = f_i(\text{Br}(M))$ denote the number of i -dimensional faces of the broken circuit complex $\text{Br}(M)$. The corresponding h-vector $(h_0, h_1, \dots, h_{r-1})$ can be read off from any shelling (cf. [7, §7.2] and [29, §2]). It satisfies

$$(12) \quad \sum_{i=0}^{r-1} \frac{f_{i-1} z^i}{(1-z)^i} = \frac{h_0 + h_1 z + h_2 z^2 + \dots + h_{r-1} z^{r-1}}{(1-z)^r}.$$

The relation between the f-vector and the h-vector holds for any simplicial complex [29]. The rational function (12) is the *Hilbert series* (see [29]) of the *Stanley-Reisner ring* of the

broken circuit complex $\text{Br}(M)$. The defining ideal of the Stanley-Reisner ring is generated by the monomials $\prod_{i \in C \setminus \{\min(C)\}} x_i$ representing broken circuits. Proudfoot and Speyer [24] constructed a *broken circuit ring*, which is the quotient of $K[x_1, \dots, x_n]$ modulo a prime ideal whose initial ideal is precisely this monomial ideal. Hence (12) is also the Hilbert series of the ring in [24]. In particular, the Möbius number is the common degree of both rings:

$$(13) \quad |\mu(M)| = h_0 + h_1 + h_2 + \dots + h_{r-1}.$$

Example 9 (Uniform matroids). If \mathcal{L} is a general r -dimensional subspace of K^n then $M = M(\mathcal{L})$ is the *uniform matroid* $M = U_{r,n}$, whose bases are all r -subsets in $[n]$ and whose circuits are all $(r+1)$ -subsets of $[n]$. The broken circuits of M are then all the r -subsets of $\{2, \dots, n\}$. The broken circuit complex $\text{Br}(M)$ is the $(r-1)$ -dimensional simplicial complex on $[n]$ whose maximal simplices are $\{1, j_1, \dots, j_{r-1}\}$ where $2 \leq j_1 < \dots < j_{r-1} \leq n$. This shows that f_i equals $\binom{n-1}{i}$ for $1 \leq i \leq r-1$. Using (12) one finds that $h_i = \binom{n-r+i-1}{i}$. We can then use (13) to compute the Möbius number of the uniform matroid $M = U_{r,n}$:

$$(14) \quad |\mu(U_{r,n})| = \sum_{i=0}^{r-1} \binom{n-r+i-1}{i} = \binom{n-1}{r-1}.$$

This binomial coefficient is an upper bound on $|\mu(M)|$ for any rank r matroid M on $[n]$.

To understand the geometric interpretation of $|\mu(M)|$, let us identify \mathcal{L}^\perp with \mathbb{R}^{n-r} . Here we are given n general hyperplanes through the origin in \mathbb{R}^{n-r} , and we replace each of them by a random parallel translate. The resulting arrangement of n affine hyperplanes in \mathbb{R}^{n-r} creates precisely $\binom{n-1}{r-1}$ bounded regions, as promised by the conjunction of (14) and (11). \diamond

The Möbius number is important to us because it computes the degree of the central curve of the primal linear program (1). See Theorem 13 below. We will take $r = d + 1$ and $\mathcal{L} = \mathcal{L}_{A,\mathbf{c}}$ to be the linear space spanned by the rows of A and the vector \mathbf{c} . The matroid $M(\mathcal{L}_{A,\mathbf{c}})$ we need there has rank $r = d + 1$ and it is denoted $M_{A,\mathbf{c}}$. We use the notation

$$(15) \quad |\mu(A, \mathbf{c})| := |\mu(M_{A,\mathbf{c}})| = |\mu(M(\mathcal{L}_{A,\mathbf{c}}))|.$$

Consider Example 2, with A the displayed 5×6 -matrix of rank $d = 4$, or the instance in Figure 3. Here, $n = 6$, $r = d + 1 = 5$, and $M_{A,\mathbf{c}} = U_{5,6}$ is the uniform matroid. Its Möbius number equals $|\mu(A, \mathbf{c})| = |\mu(U_{5,6})| = 5$. This number 5 counts the bounded segments on the vertical line on the left in Figure 3. Note that the relevant matroid for Example 2 is not, as one might expect, the graphic matroid of $K_{2,3}$. For higher-dimensional problems the matroids $M_{A,\mathbf{c}}$ we encounter are typically non-uniform.

There are many other interpretations of the Möbius invariant and the h-vector. For example, a useful identity for computations is the h-vector as an evaluation of the Tutte polynomial $T_M(x, y)$ of the matroid M (see [7, Eq. (7.15)] and the discussion in [7, §7.2]):

$$(16) \quad h_0 + h_1 z + h_2 z^2 + \dots + h_{r-1} z^{r-1} = z^r \cdot T_M(1/z, 0).$$

4. EQUATIONS DEFINING THE CENTRAL CURVE

In this section we determine the prime ideal of the central curve of the primal linear program (1). As a consequence we obtain explicit formulas for the degree, arithmetic genus and Hilbert function of the projective closure of the primal central curve. These results resolve the problem stated by Bayer and Lagarias at the end of [4, §11].

Our ground field is K will be $\mathbb{Q}(A)(\mathbf{b}, \mathbf{c})$. Here $\mathbb{Q}(A)$ denotes the subfield of \mathbb{R} generated by the entries of A and $\mathbb{Q}(A)(\mathbf{b}, \mathbf{c})$ is the rational function field generated by the coordinates b_i and c_j of the right hand side \mathbf{b} and the cost vector \mathbf{c} . We assume that these coordinates are algebraically independent over $\mathbb{Q}(A)$. This is a formal way of ensuring that our algebraic results remain valid for generic values of b_i and c_j in \mathbb{R} .

Let $\mathcal{L}_{A,\mathbf{c}}$ be the subspace of K^n spanned by the rows of A and the vector \mathbf{c} . We define the *central sheet* to be the coordinate-wise reciprocal $\mathcal{L}_{A,\mathbf{c}}^{-1}$ of that linear subspace. In precise terms, we define $\mathcal{L}_{A,\mathbf{c}}^{-1}$ to be the Zariski closure in the affine space \mathbb{C}^n of the set

$$(17) \left\{ \left(\frac{1}{u_1}, \frac{1}{u_2}, \dots, \frac{1}{u_n} \right) \in \mathbb{C}^n : (u_1, u_2, \dots, u_n) \in \mathcal{L}_{A,\mathbf{c}} \text{ and } u_i \neq 0 \text{ for } i = 1, \dots, n \right\}.$$

Lemma 10. *The Zariski closure of the primal central path $\{\mathbf{x}^*(\lambda) : \lambda \in \mathbb{R}_{\geq 0}\}$ is equal to the intersection of the central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$ with the affine-linear subspace defined by $A\mathbf{x} = \mathbf{b}$.*

Proof. We eliminate \mathbf{s}, \mathbf{y} and λ from the equations $A^T\mathbf{y} - \mathbf{s} = \mathbf{c}$ and $x_i s_i = \lambda$ as follows. We first replace the coordinates of \mathbf{s} by $s_i = \lambda/x_i$. The linear system becomes $A^T\mathbf{y} - \lambda\mathbf{x}^{-1} = \mathbf{c}$. This condition means that $\mathbf{x}^{-1} = (\frac{1}{x_1}, \dots, \frac{1}{x_n})^T$ lies in the linear space $\mathcal{L}_{A,\mathbf{c}}$ spanned by \mathbf{c} and the rows of A . The result of the elimination says that \mathbf{x} lies in the central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$. For \mathbf{x} in the Zariski-dense set $\mathcal{L}_{A,\mathbf{c}}^{-1} \cap (\mathbb{C}^*)^n$, one can reconstruct values of $\lambda, \mathbf{y}, \mathbf{s}$ for which $(\mathbf{x}, \mathbf{y}, \mathbf{s}, \lambda)$ is a solution to the equations $A^T\mathbf{y} - \mathbf{s} = \mathbf{c}$, $x_i s_i = \lambda$. This shows that $\mathcal{L}_{A,\mathbf{c}}^{-1}$ is indeed the projection of the set of these solutions onto the \mathbf{x} -coordinates. \square

The linear space $\{A\mathbf{x} = \mathbf{b}\}$ has dimension $n - d$, and we write $I_{A,\mathbf{b}}$ for its linear ideal. The central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$ is an irreducible variety of dimension $d + 1$, and we write $J_{A,\mathbf{c}}$ for its prime ideal. Both $I_{A,\mathbf{b}}$ and $J_{A,\mathbf{c}}$ are ideals in $K[x_1, \dots, x_n]$. We argue the following is true:

Lemma 11. *The prime ideal of polynomials that vanish on the central curve \mathcal{C} is $I_{A,\mathbf{b}} + J_{A,\mathbf{c}}$. The degree of both \mathcal{C} and the central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$ coincides with the Möbius number $|\mu(A, \mathbf{c})|$.*

Proof. The intersection of the affine space $\{A\mathbf{x} = \mathbf{b}\}$ with the central sheet is the variety of the ideal $I_{A,\mathbf{b}} + J_{A,\mathbf{c}}$. This ideal is prime because \mathbf{b} and \mathbf{c} are generic over $\mathbb{Q}(A)$. The intersection is the central curve. In Proposition 12 we show that the degree of the central sheet is $|\mu(A, \mathbf{c})|$, so here it only remains to show that this is the degree of the central curve as well. For a generic vector $(\mathbf{b}, c_0) \in \mathbb{R}^{d+1}$, we consider the hyperplane arrangement induced by $\{x_i = 0\}$ in the affine space $\{ \binom{A}{\mathbf{c}}\mathbf{x} = \binom{\mathbf{b}}{c_0} \}$. The number of bounded regions of this hyperplane arrangement equals the Möbius number $|\mu(A, \mathbf{c})|$, as seen in (11).

Each of these bounded regions contains a unique point maximizing $\sum_i \log |x_i|$. This point is the *analytic center* of that region. Each such analytic center lies in $\mathcal{L}_{A,\mathbf{c}}^{-1}$, and thus on the central curve by Lemma 10. This shows that the intersection of the central curve with the plane $\{\mathbf{c}^T\mathbf{x} = c_0\}$ contains at least $|\mu(A, \mathbf{c})|$ points.

Bézout's Theorem implies that the degree of a variety $V \subset \mathbb{C}^n$ is an upper bound for the degree of its intersection $V \cap H$ with an affine subspace H , provided that $n + \dim(V \cap H) = \dim(V) + \dim(H)$. We use this theorem for two inequalities; first, that the degree of $\mathcal{L}_{A,\mathbf{c}}^{-1}$ bounds the degree of the central curve \mathcal{C} , and, second, that the degree of \mathcal{C} bounds the number of its intersection points with $\{\mathbf{c}^T\mathbf{x} = c_0\}$. To summarize, we have shown:

$$|\mu(A, \mathbf{c})| \leq \#(\mathcal{C} \cap \{\mathbf{c}^T\mathbf{x} = c_0\}) \leq \deg(\mathcal{C}) \leq \deg(\mathcal{L}_{A,\mathbf{c}}^{-1}) = |\mu(A, \mathbf{c})|.$$

From this we conclude that $|\mu(A, \mathbf{c})|$ is the degree of the primal central curve \mathcal{C} . \square

At this point we are left with the problem of computing the degree of the homogeneous ideal $J_{A,\mathbf{c}}$ and a set of generators. Luckily, this has already been done for us in the literature. The following proposition was proved by Proudfoot and Speyer [24] and it refines an earlier result of Terao [31]. See also [6] for related results. The paper [30] begins the challenging task of extending these results from linear programming to semidefinite programming.

Proposition 12 (Proudfoot-Speyer [24]). *The degree of the central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$, regarded as a variety in complex projective space, coincides with the Möbius number $|\mu(A, \mathbf{c})|$. Its prime ideal $J_{A,\mathbf{c}}$ is generated by a universal Gröbner basis consisting of all homogeneous polynomials*

$$(18) \quad \sum_{i \in \text{supp}(v)} v_i \cdot \prod_{j \in \text{supp}(v) \setminus \{i\}} x_j,$$

where $\sum v_i x_i$ runs over non-zero linear forms of minimal support that vanish on $\mathcal{L}_{A,\mathbf{c}}$.

Proof. The construction in [24] associates the ring $K[x_1, \dots, x_n]/J_{A,\mathbf{c}}$ to the linear subspace $\mathcal{L}_{A,\mathbf{c}}$ of K^n . Theorem 4 of [24] says that the homogeneous polynomials (18) form a universal Gröbner bases for $J_{A,\mathbf{c}}$. As argued in [24, Lemma 2], this means that the ring degenerates to the Stanley-Reisner ring of the broken circuit complex $\text{Br}(M_{A,\mathbf{c}})$. Hence, by our discussion in Section 3, or by [24, Prop. 7], the Hilbert series of $K[x_1, \dots, x_n]/J_{A,\mathbf{c}}$ is the rational function (12), and the degree of $J_{A,\mathbf{c}}$ equals $|\mu(A, \mathbf{c})|$ as seen in (13). The ideal $J_{A,\mathbf{c}}$ is radical, since its initial ideal is square-free, and hence it is prime because its variety $\mathcal{L}_{A,\mathbf{c}}^{-1}$ is irreducible. \square

The polynomials in (18) correspond to the circuits of the matroid $M_{A,\mathbf{c}}$, of which there are at most $\binom{n}{d+2}$. If the matrix A is generic, then $M_{A,\mathbf{c}}$ is the uniform matroid and every $(d+2)$ -subset of $\{1, \dots, n\}$ forms a circuit. In this case, by (14), its Möbius number equals

$$|\mu(A, \mathbf{c})| = \binom{n-1}{d}.$$

For arbitrary matrices A , this binomial coefficient furnishes an upper bound on the Möbius number $|\mu(A, \mathbf{c})|$. We are now prepared to conclude with the main theorem of this section. The analogous equations for the dual central curve are given in Proposition 24 in Section 7.

Theorem 13. *The degree of the primal central curve of (1) is the Möbius number $|\mu(A, \mathbf{c})|$ and is hence at most $\binom{n-1}{d}$. The prime ideal of polynomials that vanish on the primal central path is generated by the circuit polynomials (18) and the d linear polynomials in $A\mathbf{x} - \mathbf{b}$.*

Proof. This is an immediate consequence of Lemmas 10 and 11 and Proposition 12. \square

It is convenient to write the circuit equations (18) in the following determinantal representation. Suppose that A has format $d \times n$ and its rows are linearly independent. Then the linear forms of minimal support that vanish on $\mathcal{L}_{A,\mathbf{c}}$ are the $(d+2) \times (d+2)$ -minors of the $(d+2) \times n$ matrix $\begin{pmatrix} A \\ \mathbf{c} \\ \mathbf{x} \end{pmatrix}$. This gives the following concise description of our prime ideal $J_{A,\mathbf{c}}$:

$$(19) \quad J_{A,\mathbf{c}} = I_{\text{num},d+2} \left(\begin{pmatrix} A \\ \mathbf{c} \\ \mathbf{x}^{-1} \end{pmatrix} \right)$$

where $\mathbf{x}^{-1} = (x_1^{-1}, \dots, x_n^{-1})$ and the operator $I_{\text{num},d+2}$ extracts the numerators of the $(d+2) \times (d+2)$ -minors of the matrix. For example, if the leftmost $(d+1) \times (d+2)$ submatrix of the

matrix $\binom{A}{\mathbf{c}}$ has full rank $d + 1$, then, as in (9), one generator of the ideal $J_{A,\mathbf{c}}$ equals

$$\det \begin{pmatrix} A_1 & A_2 & \dots & A_{d+2} \\ c_1 & c_2 & \dots & c_{d+2} \\ x_1^{-1} & x_2^{-1} & \dots & x_{d+2}^{-1} \end{pmatrix} \cdot \prod_{i \in \mathcal{I}} x_i,$$

where \mathcal{I} is the unique circuit contained in $\{1, 2, \dots, d + 2\}$. Note that there are $\binom{n}{d+2}$ such minors but they need not be distinct and some of them may be zero.

Example 14. Let $d = 4, n = 6$ and A the matrix in Example 2. The linear ideal is

$$I_{A,\mathbf{b}} = \langle x_1 + x_2 + x_3 - b_1, x_4 + x_5 + x_6 - b_2, x_1 + x_4 - b_3, x_2 + x_5 - b_4 \rangle.$$

The central sheet $\mathcal{L}_{A,\mathbf{c}}^{-1}$ is the quintic hypersurface whose defining polynomial is

$$(20) \quad f_{A,\mathbf{c}}(\mathbf{x}) = \det \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ c_1 & c_2 & c_3 & c_4 & c_5 & c_6 \\ x_1^{-1} & x_2^{-1} & x_3^{-1} & x_4^{-1} & x_5^{-1} & x_6^{-1} \end{pmatrix} \cdot x_1 x_2 x_3 x_4 x_5 x_6.$$

The primal central curve is the plane quintic defined by the ideal $I_{A,\mathbf{b}} + \langle f_{A,\mathbf{c}} \rangle$. This ideal is prime for general choices of \mathbf{b} and \mathbf{c} . However, this may fail for special values: the quintic on the left in Figure 1 is irreducible but that on the right decomposes into a quartic and a line. For a concrete numerical example we set $b_1 = b_2 = 3$ and $b_3 = b_4 = b_5 = 2$. Then the transportation polygon P is the regular hexagon depicted in Figure 1. Its vertices are

$$(21) \quad \begin{pmatrix} 0 & 1 & 2 \\ 2 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 2 & 1 \\ 2 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 2 & 0 \\ 1 & 0 & 2 \end{pmatrix}, \begin{pmatrix} 2 & 0 & 1 \\ 0 & 2 & 1 \end{pmatrix}, \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \end{pmatrix}.$$

Consider the two transportation problems (1) given by $\mathbf{c} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 3 \end{pmatrix}$ and $\mathbf{c}' = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 2 \end{pmatrix}$.

In both cases, the last matrix in (21) is the unique optimal solution. Modulo the linear ideal $I_{A,\mathbf{b}}$ we can write the quintics $f_{A,\mathbf{c}}$ and $f_{A,\mathbf{c}'}$ as polynomials in only two variables x_1 and x_2 :

$$\begin{aligned} f_{A,\mathbf{c}} &= 3x_1^4 x_2 + 5x_1^3 x_2^2 - 2x_1 x_2^4 - 3x_1^4 - 22x_1^3 x_2 - 15x_1^2 x_2^2 + 8x_1 x_2^3 + 2x_2^4 \\ &\quad + 18x_1^3 + 45x_1^2 x_2 - 12x_2^3 - 33x_1^2 - 22x_1 x_2 + 22x_2^2 + 18x_1 - 12x_2, \\ f_{A,\mathbf{c}'} &= (x_2 - 1) \cdot (2x_1^4 + 4x_1^3 x_2 + x_1^2 x_2^2 - x_1 x_2^3 - 12x_1^3 - 14x_1^2 x_2 + x_1 x_2^2 \\ &\quad + x_2^3 + 22x_1^2 + 10x_1 x_2 - 5x_2^2 - 12x_1 + 6x_2). \end{aligned}$$

Both quintics pass through all intersection points of the arrangement of six lines. The cost matrix \mathbf{c} exemplifies the generic behavior, when the quintic curve is irreducible. On the other hand, the central path for \mathbf{c}' is a segment on the horizontal line $x_2 = 1$ in Figure 1. \diamond

Remark 15. When \mathbf{b} or \mathbf{c} is not generic, various aspects of the above analysis break down. If \mathbf{b} is not generic, then the hyperplane arrangement $\{x_i = 0\}_{i \in [n]} \subset \{\mathbf{A}\mathbf{x} = \mathbf{b}\}$ may not be simple, that is, it may have a vertex at which more than $n - d$ hyperplanes meet. This vertex will maximize $\mathbf{c}^T \mathbf{x}$ over more than one adjoining region of the arrangement. In particular, the central curve passes through this vertex more than once and is singular at this point.

If the cost function \mathbf{c} is maximized at a (non-vertex) face of a region of the hyperplane arrangement $\{x_i = 0\}_{i \in [n]} \subset \{\mathbf{A}\mathbf{x} = \mathbf{b}\}$, then the central curve meets this face in its analytic center and does not pass through any of the vertices of the hyperplane arrangement contained in the affine span of this face. For example, see Figure 2. Another potential problem is that

for non-generic \mathbf{c} the curve defined by the equations of Theorem 13 may be reducible, as happens for the cost vector \mathbf{c}' in Example 14. The central curve will then be whatever component of these solutions passes through the region of interest. In particular, its degree and equations are no longer independent of the sign conditions on \mathbf{x} . Fortunately, the Möbius number $|\mu(A, \mathbf{c})|$ is always an upper bound for the degree of the central curve. \diamond

In the remainder of this section we consider the question of what happens to the central sheet, and hence to the central path, when the cost function \mathbf{c} degenerates to one of the unit vectors e_i . Geometrically this means that the cost vector becomes normal to one of the constraint hyperplanes, and the curve reflects this by breaking into irreducible components. In addition to the nice geometry, our interest in these degenerations comes from the observation that they seem to produce curves with high curvature. What follows is independent of the rest of the paper and can be skipped upon first reading.

To set up our degeneration in proper algebraic terms, we work over the field $K\{\{t\}\}$ of Puiseux series over the field $K = \mathbb{Q}(A)(\mathbf{b}, \mathbf{c})$ that was used above. The field $K\{\{t\}\}$ comes with a natural t -adic valuation. Passing to the special fiber represents the process of letting the parameter t tend to 0. Our cost vector \mathbf{c} has its coordinates in the Puiseux series field:

$$(22) \quad \mathbf{c} = (t^{w_1}, t^{w_2}, \dots, t^{w_{n-1}}, 1)$$

Here $w_1 > w_2 > \dots > w_{n-1} > 0$ are any rational numbers. We are interested in the special fiber of the central sheet $\mathcal{L}_{A, \mathbf{c}}^{-1}$. This represents the limit of the central sheet as t approaches 0. This induces a degeneration of the central curve $\mathcal{L}_{A, \mathbf{c}}^{-1} \cap \{\mathbf{A}\mathbf{x} = \mathbf{b}\}$. We wish to see how, in that limit, the central curve breaks into irreducible curves in the affine space $\{\mathbf{A}\mathbf{x} = \mathbf{b}\}$.

The ideal defining the special fiber of $J_{A, \mathbf{c}}$ is denoted $\text{in}(J_{A, \mathbf{c}}) = J_{A, \mathbf{c}}|_{t=0}$. By a combinatorial argument as in [24], the maximal minors in (19) have the Gröbner basis property for this degeneration. Hence we obtain the prime ideal of the flat family by simply dividing each such minor by a non-negative power of t . This observation implies the following result:

Theorem 16. *The central sheet $\mathcal{L}_{A, \mathbf{c}}^{-1}$ degenerates into a reduced union of central sheets of smaller linear programming instances. More precisely, the ideal $\text{in}(J_{A, \mathbf{c}})$ is radical, and it has the following representation as an intersection of ideals that are prime when A is generic:*

$$(23) \quad \text{in}(J_{A, \mathbf{c}}) = \bigcap_{i=d}^{n-1} \left(I_{\text{num}, d+1} \left(\begin{array}{cccc} A_1 & A_2 & \cdots & A_i \\ x_1^{-1} & x_2^{-1} & \cdots & x_i^{-1} \end{array} \right) + \langle x_{i+2}, x_{i+3}, \dots, x_n \rangle \right)$$

Proof sketch. The Gröbner basis property says that $\text{in}(J_{A, \mathbf{c}})$ is generated by the polynomials obtained from the maximal minors of (19) by dividing by powers of t and then setting t to zero. The resulting polynomials factor, and this factorization shows that they lie in each of the ideals on the right hand side of (23). Conversely, each element in the product of the ideals on the right hand side is seen to lie in $\text{in}(J_{A, \mathbf{c}})$. To complete the proof, it then suffices to note that $\text{in}(J_{A, \mathbf{c}})$ is radical because its generators form a square-free Gröbner basis. \square

Example 17. Let $n = 6$ and $d = 3$. The matrix A might represent the three-dimensional Klee-Minty cube. The decomposition of the initial ideal in (23) has three components:

$$\text{in}(J_{A, \mathbf{c}}) = \langle x_5, x_6 \rangle \cap \langle \det \begin{pmatrix} x_1 A_1 & x_2 A_2 & x_3 A_3 & x_4 A_4 \\ 1 & 1 & 1 & 1 \end{pmatrix}, x_6 \rangle \cap I_{\text{num}, 4} \left(\begin{array}{ccccc} A_1 & A_2 & A_3 & A_4 & A_5 \\ x_1^{-1} & x_2^{-1} & x_3^{-1} & x_4^{-1} & x_5^{-1} \end{array} \right).$$

For general A , the ideal $J_{A, \mathbf{c}}$ defines an irreducible curve of degree 10, namely the central path, in each of the 3-planes $\{\mathbf{A}\mathbf{x} = \mathbf{b}\}$. The three curves in its degeneration above are

irreducible of degrees 1, 3 and 6 respectively. The first is one of the lines in the arrangement of six facet planes, the second curve is the central path inside the facet defined by $x_6 = 0$, and the last curve is the central path of the polytope obtained by removing that facet. \diamond

In general, we can visualize the degenerated central path in the following geometric fashion. We first flow from the analytic center of the polytope to the analytic center of its last facet. Then we iterate and flow from the analytic center of the facet to the analytic center of its last facet, which is a ridge of the original polytope. Then we continue inside that ridge, etc.

5. THE GAUSS CURVE OF THE CENTRAL PATH

The total curvature of the central path is an important quantity for the estimation of the running time of interior point methods in linear programming [11, 22, 28, 34, 37]. In this section we relate the algebraic framework developed so far to the problem of bounding the total curvature. The relevant geometry was pioneered by Dedieu, Malajovich and Shub [11]. Following their approach, we consider the *Gauss curve* associated with the primal central path. The Gauss curve is the image of the central curve under the Gauss map, and its arc length is precisely the total curvature of the central path. Moreover, the arc length of the Gauss curve can be bounded in terms of its degree. An estimate of that degree, via the multihomogeneous Bézout Theorem, was the workhorse in [11]. Our main result here is a more precise bound, in terms of matroid invariants, for the degree of the Gauss curve of the primal central curve. As a corollary we obtain a new upper bound on its total curvature.

We begin our investigation by reviewing definitions from elementary differential geometry. Consider an arbitrary curve $[a, b] \rightarrow \mathbb{R}^m$, $t \mapsto f(t)$, whose parameterization is twice differentiable and whose derivative $f'(t)$ is a non-zero vector for all parameter values $t \in [a, b]$. This curve has an associated *Gauss map* into the unit sphere S^{m-1} , which is defined as

$$\gamma : [a, b] \rightarrow S^{m-1}, \quad t \mapsto \frac{f'(t)}{\|f'(t)\|}.$$

The image $\gamma = \gamma([a, b])$ of the Gauss map in S^{m-1} is called the *Gauss curve* of the given curve f . In our situation, the curve f is algebraic, with known defining polynomial equations, and it makes sense to consider the *projective Gauss curve* in complex projective space \mathbb{P}^{m-1} . By this we mean the Zariski closure of the image of the Gauss curve under the double-cover map $S^{m-1} \rightarrow \mathbb{P}^{m-1}$. If $m = 2$, so that \mathcal{C} is a non-linear plane curve, then the Gauss curve traces out several arcs on the unit circle S^1 , and the projective Gauss curve is the entire projective line \mathbb{P}^1 . Here, the line \mathbb{P}^1 comes with a natural multiplicity, to be derived in Example 22.

If $m = 3$ then the Gauss curve lies on the unit sphere S^2 and the projective Gauss curve lives in the projective plane \mathbb{P}^2 . Since a curve in 3-space typically has parallel tangent lines, the Gauss curve is here expected to have singularities, even if f is a smooth curve.

The *total curvature* K of our curve f is defined to be the arc length of its associated Gauss curve γ ; see [11, §3]. This quantity admits the following expression as an integral:

$$(24) \quad K := \int_a^b \left\| \frac{d\gamma(t)}{dt} \right\| dt.$$

The *degree* of the Gauss curve $\gamma(t)$ is defined as the maximum number of intersection points, counting multiplicities, with any hyperplane in \mathbb{R}^m , or equivalently, with any equator in S^{m-1} . This (geometric) degree is bounded above by the (algebraic) degree of the projective Gauss curve in \mathbb{P}^{m-1} . The latter can be computed exactly, from any polynomial

representation of \mathcal{C} , using standard methods of computer algebra. Throughout this section, by *degree* we mean the degree of the image of γ in \mathbb{P}^{m-1} multiplied by the degree of the map that takes \mathcal{C} onto $\gamma(\mathcal{C})$. From now on we use the notation $\deg(\gamma(\mathcal{C}))$ for that number.

Proposition 18. [11, Corollary 4.3] *The total curvature of any real algebraic curve \mathcal{C} in \mathbb{R}^m is bounded above by π times the degree of its projective Gauss curve in \mathbb{P}^{m-1} . In symbols,*

$$K \leq \pi \cdot \deg(\gamma(\mathcal{C})).$$

Remark 19. In higher dimensions, the degree of the map from \mathcal{C} onto $\gamma(\mathcal{C})$ is typically equal to 1, in which case our definition of $\deg(\gamma(\mathcal{C}))$ is exactly that used in [11]. However, that the extra factor is needed can be seen by considering the case $m = 2$ of non-linear plane curves: the Gauss curve $\gamma(\mathcal{C})$ is just \mathbb{P}^1 , but it has a non-reduced structure coming from the map. \diamond

We now present our main result in this section, which concerns the degree of the projective Gauss curve $\gamma(\mathcal{C})$, when \mathcal{C} is the central curve of a linear program in primal formulation. As before, A is an arbitrary real matrix of rank d having n columns, but the cost vector \mathbf{c} and the right hand side \mathbf{b} are generic over $\mathbb{Q}(A)$. The curve \mathcal{C} lives in an $(n - d)$ -dimensional affine subspace of \mathbb{R}^n , which we identify with \mathbb{R}^{n-d} , so that $\gamma(\mathcal{C})$ is a curve in \mathbb{P}^{n-d-1} .

Let $M_{A,\mathbf{c}}$ denote the matroid of rank $d + 1$ on the ground set $[n]$ associated with the matrix $\begin{pmatrix} A \\ \mathbf{c} \end{pmatrix}$. We write (h_0, h_1, \dots, h_d) for the h-vector of the broken circuit complex of $M_{A,\mathbf{c}}$, as defined in (12). In the generic case, $M_{A,\mathbf{c}} = U_{d+1,n}$ is the uniform matroid, as in Example 9. In this case, the coordinates of the h-vector are $h_i = \binom{n-d+i-2}{i}$. For special matrices A , this simplicial complex gets replaced by a pure shellable subcomplex of the same dimension, so the h-vector (weakly) decreases in each entry. Hence, the following always holds:

$$(25) \quad h_i \leq \binom{n-d+i-2}{i} \quad \text{for } i = 0, 1, \dots, d.$$

As indicated, this inequality holds with equality when $M_{A,\mathbf{c}}$ is the uniform matroid.

Theorem 20. *The degree of the projective Gauss curve of the primal central curve \mathcal{C} satisfies*

$$(26) \quad \deg(\gamma(\mathcal{C})) \leq 2 \cdot \sum_{i=1}^d i \cdot h_i.$$

In particular, we have the following upper bound which is tight for generic matrices A :

$$(27) \quad \deg(\gamma(\mathcal{C})) \leq 2 \cdot (n - d - 1) \cdot \binom{n-1}{d-1}.$$

The difference between the bound in (26) and the degree of $\gamma(\mathcal{C})$ can be explained in terms of singularities the curve \mathcal{C} may have on the hyperplane at infinity. The relevant algebraic geometry will be seen in the proof of Theorem 20, which we shall present after an example.

Example 21. In the following two instances we have $d = 3$ and $n = 6$.

- (1) First assume that A is a generic 3×6 -matrix. The arrangement of six facet planes creates 10 bounded regions. The primal central curve \mathcal{C} has degree $\binom{6-1}{3} = 10$. It passes through the $\binom{6}{3} = 20$ vertices of the arrangements. In-between it visits the 10 analytic centers of the bounded regions. Here the curve \mathcal{C} is smooth and its genus is 11. This number is seen from the formula (29) below. The corresponding Gauss curve in \mathbb{P}^2 has degree $2 \cdot 10 + 2 \cdot \text{genus}(\mathcal{C}) - 2 = 40$, as given by the right hand side of (27). Hence the total curvature of the central curve \mathcal{C} is bounded above by 40π .

(2) Next consider the Klee-Minty cube in 3-space. Normally, it is given by the constraints

$$0 \leq z_1 \leq 1, \quad \epsilon z_1 \leq z_2 \leq 1 - \epsilon z_1, \quad \text{and} \quad \epsilon z_2 \leq z_3 \leq 1 - \epsilon z_2.$$

To see this in a primal formulation (1), we use z_1, z_2, z_3 to parametrize the affine space $\{A\mathbf{x} = \mathbf{b}\}$. The facets of the cube then correspond to the intersection of the coordinate hyperplanes with this affine space. This is given by the matrices

$$\begin{pmatrix} A \\ \mathbf{c} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 2\epsilon & 0 & 1 & 1 & 0 & 0 \\ 2\epsilon^2 & 0 & 2\epsilon & 0 & 1 & 1 \\ c_1 & c_2 & c_3 & c_4 & c_5 & c_6 \end{pmatrix} \quad \text{and} \quad \mathbf{b} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Here ϵ is a small positive real constant. The above 4×6 -matrix is not generic, and its associated matroid $M_{A,\mathbf{c}}$ is not uniform. It has exactly one non-basis, and so the h-vector equals $(h_0, h_1, h_2, h_3) = (1, 2, 3, 3)$. The central curve \mathcal{C} has degree $\sum_{i=0}^3 h_i = 9$. In the coordinates used above, the curve is defined by the 5×5 -minors of the 5×6 -matrix which is obtained from the 4×6 -matrix $\begin{pmatrix} A \\ \mathbf{c} \end{pmatrix}$ by adding one row consisting of reciprocal facet equations:

$$(z_1^{-1}, (1 - z_1)^{-1}, (z_2 - \epsilon z_1)^{-1}, (1 - z_2 - \epsilon z_1)^{-1}, (z_3 - \epsilon z_2)^{-1}, (1 - z_3 - \epsilon z_2)^{-1}).$$

According to Theorem 20, the degree of the Gauss curve $\gamma(\mathcal{C})$ in \mathbb{P}^2 is bounded above by $2 \sum_{i=1}^3 i \cdot h_i = 34$. A computation using Macaulay2 [15] reveals that $\text{degree}(\gamma(\mathcal{C})) = 32$ and that the total curvature is bounded by 32π . \diamond

Proof of Theorem 20. For the proof we shall use the *generalized Plücker formula for curves*:

$$(28) \quad \text{deg}(\gamma(\mathcal{C})) = 2 \cdot \text{deg}(\mathcal{C}) + 2 \cdot \text{genus}(\mathcal{C}) - 2 - \kappa.$$

The formula in (28) is obtained from [23, Thm. (3.2)] by setting $m = 1$ or from [16, Eq. (4.26)] by setting $k = 0$. The quantity κ is a non-negative integer and it measures the singularities of the curve \mathcal{C} . We have $\kappa = 0$ whenever the projective curve \mathcal{C} is smooth, and this happens in our application when $M_{A,\mathbf{c}}$ is the uniform matroid. In general, we may have singularities at infinity because here the real affine curve \mathcal{C} has to be replaced by its closure in complex projective space \mathbb{P}^{n-d} , which is the projectivization of the affine space defined by $A\mathbf{x} = \mathbf{b}$. The degree and genus on the right hand side of (28) refer to that projective curve in \mathbb{P}^{n-d} .

The references above actually give the degree of the *tangent developable* of the projective curve \mathcal{C} , but we see that this equals the degree of the Gauss curve. The tangent developable is the surface obtained by taking the union of all tangent lines at points in \mathcal{C} . The projective Gauss curve $\gamma(\mathcal{C})$ is obtained from the tangent developable by intersecting it with a hyperplane, namely, the hyperplane at infinity, representing the directions of lines.

In the formula (28), the symbol $\text{genus}(\mathcal{C})$ refers to the arithmetic genus of the curve. We shall now compute this arithmetic genus for primal central curve \mathcal{C} . For this we use the formula for the Hilbert series of the central sheet due to Terao, in Theorem 1.2 on page 551 of [31]. See the recent work of Berget [6] for a nice proof of a more general statement.

As seen in the proof of Proposition 12, the Hilbert series of the coordinate ring of the central sheet equals

$$\frac{h_0 + h_1 z + h_2 z^2 + \cdots + h_d z^d}{(1 - z)^{d+2}}.$$

The central curve \mathcal{C} is obtained from the central sheet by intersection with a general linear subspace of dimension $n - d$. The (projective closure of the) central sheet is arithmetically Cohen-Macaulay since it has a flat degeneration to a shellable simplicial complex, as shown by Proudfoot and Speyer [24]. We conclude that the Hilbert series of the central curve \mathcal{C} is

$$\frac{h_0 + h_1z + h_2z^2 + \cdots + h_dz^d}{(1 - z)^2} = \sum_{m \geq d} \left[\left(\sum_{i=0}^d h_i \right) \cdot m + \sum_{j=0}^d (1 - j)h_j \right] z^m + O(z^{d-1}).$$

The parenthesized expression is the Hilbert polynomial of the projective curve \mathcal{C} . The degree of \mathcal{C} is the coefficient of m , and using (13) we again find this to be the Möbius number:

$$\text{degree}(\mathcal{C}) = |\mu(A, \mathbf{c})| = \sum_{i=0}^d h_i.$$

The arithmetic genus of the curve \mathcal{C} is one minus the constant term of its Hilbert polynomial:

$$(29) \quad \text{genus}(\mathcal{C}) = 1 - \sum_{j=0}^d (1 - j)h_j.$$

We now see that our assertion (26) follows directly from the generalized Plücker formula (28).

For fixed d and n , the degree and genus of \mathcal{C} are maximal when the matrix A is generic. In this case, h_i equals the right hand side of (25), and we need to sum these binomial coefficients times two. Hence, our second assertion (27) follows from the identity

$$\sum_{i=0}^d i \cdot \binom{n - d + i - 2}{i} = (n - d - 1) \cdot \binom{n - 1}{d - 1}.$$

This completes the proof of Theorem 20. \square

Example 22. Let $d = n - 2$ and suppose A is generic. The primal central curve \mathcal{C} is a plane curve. Our h-vector equals $(1, 1, \dots, 1)$. The proof of Theorem 20 reveals that the degree of \mathcal{C} is $d + 1 = n - 1$ and the genus of \mathcal{C} is $\binom{d}{2}$. The Gauss curve $\gamma(\mathcal{C})$ is the projective line \mathbb{P}^1 , but regarded with multiplicity $\text{deg}(\gamma(\mathcal{C})) = (d + 1)d$. This number is the degree of the projectively dual curve \mathcal{C}^\vee . The identity (28) specializes to the Plücker formula for plane curves, which expresses the degree of \mathcal{C}^\vee in terms of the degree and the singularities of \mathcal{C} . \diamond

We close this section by showing how to compute the Gauss curve for a non-planar instance.

Example 23. Let $n = 5, d = 2$ and $A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 4 \end{pmatrix}$. The primal central curve has degree 6 and its equations are obtained by clearing denominators in the 4×4 -minors of

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 4 \\ c_1 & c_2 & c_3 & c_4 & c_5 \\ (z_1 - g_1)^{-1} & (-2z_1 + z_2 - g_2)^{-1} & (z_1 - 2z_2 + z_3 - g_3)^{-1} & (z_2 - 2z_3 - g_4)^{-1} & (z_3 - g_5)^{-1} \end{pmatrix}.$$

The c_i and g_j are random constants representing the cost function and right hand side of (1). To be precise, the vector $\mathbf{g} = (g_1, g_2, g_3, g_4, g_5)^T$ satisfies $\mathbf{A}\mathbf{g} = \mathbf{b}$ as in Sections 6 and 7. The linear forms in $\mathbf{z} = (z_1, z_2, z_3)$ come from the change of coordinates $B^T\mathbf{z} - \mathbf{x} = \mathbf{g}$ where B is a 3×5 matrix whose rows span the kernel of A . Writing \mathbf{I} for the ideal of these polynomials,

the following one-line command in the computer algebra system Macaulay2 [15] computes the defining polynomial of the Gauss curve in \mathbb{P}^2 :

```
eliminate({z1,z2,z3},I+minors(1,matrix{{u,v,w}}*diff(matrix{{z1},{z2},{z3}},gens I)))
```

The output is a homogeneous polynomial of degree 16 in the coordinates u, v, w on \mathbb{P}^2 . Note that $\deg(\gamma(\mathcal{C})) = 16$ is consistent with Theorem 20 because $h = (h_0, h_1, h_2) = (1, 2, 3)$. \diamond

6. THE PRIMAL-DUAL FORMULATION AND THE AVERAGE TOTAL CURVATURE

This short section offers a dictionary that translates between the primal and the dual central curve. We begin by offering an algebraic representation of the primal-dual central curve that is more symmetric than that given in the Introduction. This will allow us to take our main results so far, previously only stated for the primal LP, and show that they hold verbatim also for the dual LP. As primary application, we then derive a matroid-theoretic refinement of the Main Theorem in [11] on the average total curvature of the central path.

Let \mathcal{L}_A denote the row space of the matrix A and \mathcal{L}_A^\perp its orthogonal complement in \mathbb{R}^n . We also fix a vector $\mathbf{g} \in \mathbb{R}^n$ such that $A\mathbf{g} = \mathbf{b}$. By eliminating \mathbf{y} from the system (5) in Theorem 1, we see that the primal-dual central path $(\mathbf{x}^*(\lambda), \mathbf{s}^*(\lambda))$ has the following symmetric description:

$$(30) \quad \mathbf{x} \in \mathcal{L}_A^\perp + \mathbf{g}, \quad \mathbf{s} \in \mathcal{L}_A + \mathbf{c} \quad \text{and} \quad x_1s_1 = x_2s_2 = \cdots = x_ns_n = \lambda.$$

The implicit (*i.e.* λ -free) representation of the primal-dual central curve is simply obtained by erasing the very last equality “ $= \lambda$ ” in (30). Its prime ideal is generated by the quadrics $x_i s_i - x_j s_j$ and the affine-linear equations defining $\mathcal{L}_A^\perp + \mathbf{g}$ in \mathbf{x} -space and $\mathcal{L}_A + \mathbf{c}$ in \mathbf{s} -space.

The symmetric description of the central path in (30) lets us write down the statements from Section 4 for the dual version. For example, we derive equations for the dual central curve in \mathbf{s} -space or \mathbf{y} -space as follows. Let B be any $(n - d) \times n$ matrix whose rows span the kernel of A . In symbols, $\mathcal{L}_B = \mathcal{L}_A^\perp$. The (dual) central curve in \mathbf{s} -space is obtained by intersecting the d -dimensional affine space $\mathcal{L}_A + \mathbf{c} = \{\mathbf{s} \in \mathbb{R}^n : B\mathbf{s} = B\mathbf{c}\}$ with the central sheet $\mathcal{L}_{B,\mathbf{g}}^{-1}$ in (17). To obtain the central curve in \mathbf{y} -space, we substitute $s_i = \sum_{j=1}^d a_{ji}y_j - c_i$ in the equations defining $\mathcal{L}_{B,\mathbf{g}}^{-1}$. This gives dual formulations of Theorems 13 and 20:

Corollary 24. *The degree of the dual central curve of (2) equals the Möbius number $|\mu(B, \mathbf{g})|$ and is hence at most $\binom{n-1}{d-1}$. The prime ideal of polynomials that vanish on the central path is generated by the circuit polynomials (18), but now associated with the space generated by the rows of B and the vector \mathbf{g} , and the $n - d$ linear equations in \mathbf{s} given by $B\mathbf{s} = B\mathbf{c}$.*

Corollary 25. *The degree $\deg(\gamma(\mathcal{C}))$ of the Gauss image of the dual central curve \mathcal{C} is at most $2 \cdot \sum_i i \cdot h_i$, where $h = h(\text{Br}(M_{B,\mathbf{g}}))$. This implies the bound $\deg(\gamma(\mathcal{C})) \leq 2 \cdot (d-1) \cdot \binom{n-1}{d}$.*

Remark 26. Theorem 20 and Corollary 25 give a strengthening of Theorem 1.1 of [11]. Megiddo and Shub [21] proved the lower bound $(d - 1)\pi/2$ for the total curvature of the central path of a d variable LP with $d + 1$ constraints. For such instances, our upper bound is tight up to a constant. Our main contribution is that we adjust the upper bound to the specific matroid of the constraint matrix A . \diamond

Dedieu, Malajovich and Shub [11] investigated the average curvature of the central paths over all the bounded regions of the associated hyperplane arrangement for the linear program and show this to be at most $2\pi d$ when A is generic. Combining (11), Proposition 18, Theorem 20, and their dual versions seen in this section, we obtain the following statement:

Theorem 27. *The average total curvature of the primal central path over the bounded regions of the hyperplane arrangement $\{x_i = 0\}_{i \in [n]}$ in the affine space $\mathcal{L}_A^\perp + \mathbf{g}$ is at most*

$$(31) \quad 2\pi \cdot \frac{\sum_{i=0}^d i \cdot h_i(A, \mathbf{c})}{|\mu(A)|}.$$

If A is a generic $d \times n$ matrix, then this average total curvature is at most $2\pi(n - d - 1)$.

Similarly, the average total curvature of the dual central path over the bounded regions of the hyperplane arrangement $\{s_i = 0\}_{i \in [n]}$ in the affine space $\mathcal{L}_A + \mathbf{c}$ is at most

$$(32) \quad 2\pi \cdot \frac{\sum_{i=0}^d i \cdot h_i(B, \mathbf{g})}{|\mu(B)|}.$$

In the generic case, this average total curvature is bounded above by $2\pi(d - 1)$.

Note that the h-vector entries in the numerator of (31) refer to the rank $d+1$ matroid $M_{A,\mathbf{c}}$ while the Möbius number in the denominator refers to the rank d matroid M_A . Likewise, the h-vector entries in the numerator of (32) refer to the rank $n - d + 1$ matroid $M_{B,\mathbf{g}}$ while the Möbius number in the denominator refers to the rank $n - d$ matroid M_B .

Example 28. We consider two examples to demonstrate the finer bounds provided by Theorem 27. First, recall the primal LP formulation of the Klee-Minty Cube in Example 21. We had already seen that our formulas predict the total curvature is bounded by 32π . We can calculate $|\mu(A)|$ for the 3×6 defining matrix. The matroid in question is not uniform (it only has 14 bases). We recover $\mu(A)$ through the computation and evaluation of the Tutte polynomial as shown in 16. In this case we obtain $|\mu(A)| = 5$ thus Formula 31 says the average curvature is no more than $\frac{32\pi}{5}$. This is a better bound than the generic case formula.

Second, recall the DTZ snake of Example 3 and Figure 2 which is given in dual LP form. The curve has degree four and its projective closure \mathcal{C} is smooth in \mathbb{P}^2 . So, we have $\deg(\gamma(\mathcal{C})) = 12$, and Proposition 18 gives an upper bound 12π on the total curvature of the entire central curve in \mathbb{R}^2 . The Möbius number $|\mu(B)|$ is in this case 10, thus we get from formula (32) that the average curvature is 1.2π . Once more, our bound is better than the generic case prediction. It is interesting to see the large discrepancy between the curvature within a single feasible region (calculated in Example 3) and the average total curvature. \diamond

7. GLOBAL GEOMETRY OF THE CENTRAL CURVE

In this section we return to the central path in its primal-dual formulation, and we study several geometric properties. We shall show how the central curve connects the vertices of the hyperplane arrangement with the analytic centers of its bounded regions. This picture behaves well under duality, as the vertices of the two arrangements are in natural bijection.

For an algebraic geometer, it is natural to replace each of the affine spaces in (30) by a complex projective space and to study the closure \mathcal{C} of the central curve in $\mathbb{P}^n \times \mathbb{P}^n$. Algebraically, we use homogeneous coordinates $[x_0 : x_1 : \cdots : x_n]$ and $[s_0 : s_1 : \cdots : s_n]$. Writing \mathbf{x} and \mathbf{s} for the corresponding column vectors of length $n + 1$, we represent

$$\mathcal{L}_A^\perp + \mathbf{g} \quad \text{by} \quad \{\mathbf{x} \in \mathbb{P}^n : (-\mathbf{b}, A) \cdot \mathbf{x} = 0\} \quad \text{and} \quad \mathcal{L}_A + \mathbf{c} \quad \text{by} \quad \{\mathbf{s} \in \mathbb{P}^n : (-B\mathbf{c}, B) \cdot \mathbf{s} = 0\}.$$

The projective primal-dual central curve \mathcal{C} is an irreducible curve in $\mathbb{P}^n \times \mathbb{P}^n$. Its bi-homogeneous prime ideal in $K[x_0, x_1, \dots, x_n, s_0, s_1, \dots, s_n]$ can be computed by the process of saturation. Namely, we compute it as the saturation with respect to $\langle x_0 s_0 \rangle$ of the ideal

generated by the above linear forms together with the bi-homogeneous forms $x_i s_i - x_j s_j$. The resulting ideal is irreducible because the vectors \mathbf{c} and \mathbf{g} are generic.

Example 29. Let $d = 2, n = 4$. Fix 2×4 -matrices $A = (a_{ij})$ and $B = (b_{ij})$ such that $\mathcal{L}_B = \mathcal{L}_A^\perp$. We start with the ideal J in $K[x_0, \dots, x_4, s_0, \dots, s_4]$ generated by

$$\begin{aligned} & a_{11}(x_1 - g_1 x_0) + a_{12}(x_2 - g_2 x_0) + a_{13}(x_3 - g_3 x_0) + a_{14}(x_4 - g_4 x_0), \\ & a_{21}(x_1 - g_1 x_0) + a_{22}(x_2 - g_2 x_0) + a_{23}(x_3 - g_3 x_0) + a_{24}(x_4 - g_4 x_0), \\ & b_{11}(s_1 - c_1 s_0) + b_{12}(s_2 - c_2 s_0) + b_{13}(s_3 - c_3 s_0) + b_{14}(s_4 - c_4 s_0), \\ & b_{21}(s_1 - c_1 s_0) + b_{22}(s_2 - c_2 s_0) + b_{23}(s_3 - c_3 s_0) + b_{24}(s_4 - c_4 s_0), \\ & s_1 x_1 - s_2 x_2, \quad s_2 x_2 - s_3 x_3, \quad s_3 x_3 - s_4 x_4. \end{aligned}$$

The central curve \mathcal{C} is irreducible in $\mathbb{P}^4 \times \mathbb{P}^4$. It has degree $(3, 3)$ unless A is very special. The prime ideal of \mathcal{C} is computed as the saturation $(J : \langle x_0 s_0 \rangle^\infty)$. We find that this ideal has two minimal generators in addition to the seven above. These are cubic polynomials in \mathbf{x} and in \mathbf{s} , which define the primal and dual central curves. They are shown in Figure 4. \diamond

Returning to the general case, we note that, since our curve \mathcal{C} lives in $\mathbb{P}^n \times \mathbb{P}^n$, its *degree* is now a pair of integers $(d_{\mathbf{x}}, d_{\mathbf{s}})$. These two integers can be defined geometrically: $d_{\mathbf{x}}$ is the number of solutions of a general equation $\sum_{i=0}^n \alpha_i x_i = 0$ on the curve \mathcal{C} , and $d_{\mathbf{s}}$ is the number of solutions of a general equation $\sum_{i=0}^n \beta_i s_i = 0$ on \mathcal{C} .

Corollary 30. Let \mathbf{c} and \mathbf{g} be generic vectors in \mathbb{R}^n and let $(d_{\mathbf{x}}, d_{\mathbf{s}})$ be the degree of the projective primal-dual central curve $\mathcal{C} \subset \mathbb{P}^n \times \mathbb{P}^n$. This degree is given by our two Möbius numbers, namely $d_{\mathbf{x}} = |\mu(A, \mathbf{c})|$ and $d_{\mathbf{s}} = |\mu(B, \mathbf{g})|$. These numbers are defined in (15).

Proof. The projection from the primal-dual central curve onto its image in either \mathbf{x} -space or \mathbf{s} -space is birational. For instance, if \mathbf{x} is a general point on the primal central curve then the corresponding point \mathbf{s} is uniquely obtained by solving the linear equations $x_i s_i = x_j s_j$ on $\mathcal{L}_A + \mathbf{c}$. Likewise, given a general point \mathbf{s} on the dual central curve we can recover the unique \mathbf{x} such that $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}$. This implies that the intersections in $\mathbb{P}^n \times \mathbb{P}^n$ that define $d_{\mathbf{x}}$ and $d_{\mathbf{s}}$ are equivalent to intersecting the primal or dual central curve with a general hyperplane in \mathbb{P}^n , and the number of points on that intersection is the respective Möbius number. \square

Next we discuss the geometry of this correspondence between the primal and dual curves at their special points, namely vertices and analytic centers of the relevant hyperplane arrangements. These special points are given by intersecting the primal-dual curve \mathcal{C} with certain bilinear equations. The sum of our two Möbius numbers, $d_{\mathbf{x}} + d_{\mathbf{s}}$, is the number of solutions of a general bilinear equation $\sum_{i,j} \gamma_{ij} x_i s_j = 0$ on the primal-dual central curve \mathcal{C} . Two special choices of such bilinear equations are of particular interest, namely, the bilinear equation $x_0 s_0 = 0$ and the bilinear equation $x_i s_i = 0$ for any $i \geq 1$. Note that the choice of the index i does not matter for the second equation because $x_i s_i = x_j s_j$ holds on the curve.

Let us first observe what happens in $\mathbb{P}^n \times \mathbb{P}^n$ when the parameter λ becomes 0. The corresponding points on the primal-dual curve \mathcal{C} are found by solving the equation $x_1 s_1 = 0$ on \mathcal{C} . Its points are the solutions of the n equations $x_1 s_1 = x_2 s_2 = \dots = x_n s_n = 0$ on the n -dimensional subvariety $(\mathcal{L}_A^\perp + \mathbf{g}) \times (\mathcal{L}_A + \mathbf{c})$ of $\mathbb{P}^n \times \mathbb{P}^n$. This intersection now contains many points in the product of affine spaces, away from the hyperplanes $\{x_0 = 0\}$ and $\{s_0 = 0\}$. We find the points by solving the linear equations $x_{i_1} = \dots = x_{i_d} = 0$ on $\mathcal{L}_A^\perp + \mathbf{g}$ and $s_{j_1} = \dots = s_{j_{n-d}} = 0$ on $\mathcal{L}_A + \mathbf{c}$, where $\{i_1, \dots, i_d\}$ runs over all bases of the matroid $M(\mathcal{L}_A)$




FIGURE 4. Correspondence of vertices and analytic centers in the two projections of a primal-dual central curve. Here both curves are plane cubics.

and $\{j_1, \dots, j_{n-d}\}$ is the complementary basis of the dual matroid $M(\mathcal{L}_A)^* = M(\mathcal{L}_B)$. These points represent vertices in the hyperplane arrangements \mathcal{H} and \mathcal{H}^* , where

$$\mathcal{H} \text{ denotes } \{x_i = 0\}_{i \in [n]} \text{ in } \mathcal{L}_A^\perp + \mathbf{g} \quad \text{and} \quad \mathcal{H}^* \text{ denotes } \{s_i = 0\}_{i \in [n]} \text{ in } \mathcal{L}_A + \mathbf{c}.$$

The vertices come in pairs corresponding to complementary bases, so the points with parameter $\lambda = 0$ on the primal-dual central curve \mathcal{C} are the pairs (\mathbf{x}, \mathbf{s}) where \mathbf{x} is a vertex in the hyperplane arrangement \mathcal{H} and \mathbf{s} is the complementary vertex in the dual arrangement \mathcal{H}^* .

Imposing the equation $x_0 s_0 = 0$ means setting $\lambda = \infty$ in the parametric representation of the central curve. The points thus obtained have the following geometric interpretation in terms of bounded regions of the hyperplane arrangements \mathcal{H} and \mathcal{H}^* . We recall that the *analytic center* of the polytope $P = \{A\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ is the unique point in the interior of P that maximizes the concave function $\sum_{i=1}^n \log(x_i)$. The algebraic characterization of the analytic center is that the gradient of $\sum_{i=1}^n \log(x_i)$, which is \mathbf{x}^{-1} , is orthogonal to the affine space $\mathcal{L}_A^\perp + \mathbf{g} = \{A\mathbf{x} = \mathbf{b}\}$. This means that the vector \mathbf{x}^{-1} lies in the row span \mathcal{L}_A of A . Let \mathcal{L}_A^{-1} denote the coordinate-wise reciprocal of \mathcal{L}_A , regarded as a subvariety in \mathbb{P}^n .

Proposition 31. *The intersection $\mathcal{L}_A^{-1} \cap (\mathcal{L}_A^\perp + \mathbf{g})$ defines a zero-dimensional reduced subscheme of the affine space $\mathbb{P}^n \setminus \{x_0 = 0\}$. All its points are defined over \mathbb{R} . They are the analytic centers of the polytopes that form the bounded regions of the arrangement \mathcal{H} .*

Proof. The analytic center of each bounded region is a point in the variety $\mathcal{L}_A^{-1} \cap (\mathcal{L}_A^\perp + \mathbf{g})$, by the gradient argument in the paragraph above. This gives us $|\mu(A)|$ real points of intersection on $\mathcal{L}_A^{-1} \cap (\mathcal{L}_A^\perp + \mathbf{g})$. By replacing $\mathcal{L}_{A,c}^{-1}$ with \mathcal{L}_A^{-1} in Proposition 12, we know that the degree of \mathcal{L}_A^{-1} is $|\mu(A)|$. This shows that these real points are all the intersection points (over \mathbb{C}) and they occur with multiplicity one. This argument closely follows the proof of Lemma 11. \square

Naturally, the dual statement holds verbatim, and we shall now state it explicitly.

Proposition 32. *The intersection $(\mathcal{L}_A^\perp)^{-1} \cap (\mathcal{L}_A + \mathbf{c})$ defines a zero-dimensional reduced subscheme of the affine space $\mathbb{P}^n \setminus \{s_0 = 0\}$. All its points are defined over \mathbb{R} . They are the analytic centers of the polytopes that form the bounded regions of the dual arrangement \mathcal{H}^* .*

We summarize our discussion with the following theorem on the global geometry of the primal-dual central curve. Figure 4 serves as an illustration for the case $n = 4$ and $d = 2$.

Theorem 33. *The primal central curve in \mathbf{x} -space passes through each vertex of the arrangement \mathcal{H} as the dual central curve in \mathbf{s} -space passes through the corresponding vertex of \mathcal{H}^* . As the primal curve passes through the analytic center of each bounded region in \mathcal{H} , the dual curve reaches the hyperplane $\{s_0 = 0\}$. Similarly, as the dual curve reaches the analytic center of each bounded region in \mathcal{H}^* , the primal curve meets the hyperplane $\{x_0 = 0\}$.*

Remark 34. The picture painted by Theorem 33 resembles the unpublished results of Adler and Haimovich [1, 8, 17] on *co-optimal paths*. They considered an LP with two objective functions \mathbf{c} and \mathbf{c}' , and they studied the parametric objective $\mathbf{c} + \rho\mathbf{c}'$. As ρ runs from $-\infty$ to ∞ , the optimal solutions form a path of edges in the arrangement \mathcal{H} . Adler and Haimovich showed that the average length (in edges) of this co-optimal path in a cell of \mathcal{H} , conditional on the path being nonempty in that cell, is at most $O(\min(d, n - d))$. This bound is close to the curvature bounds in [11]. It would be interesting to explore possible connections between co-optimal edge paths and the degenerations of central curves constructed in Theorem 16. \diamond

The primal central curve misses precisely one of the antipodal pairs of unbounded regions of \mathcal{H} . It corresponds to the region in the induced arrangement at infinity that contains the point representing the cost function \mathbf{c} . For a visualization see the picture of the central curve in Figure 3. Here a projective transformation of \mathbb{P}^2 moves the line from infinity into \mathbb{R}^2 .

The points described in Propositions 31 and 32 are precisely those points on the primal-dual central curve \mathcal{C} for which the parameter λ becomes ∞ . Equivalently, in its embedding in $\mathbb{P}^n \times \mathbb{P}^n$, these are solutions of the equation $x_0s_0 = 0$ on the curve \mathcal{C} . Note, however, that for special choices of A , the projective curve \mathcal{C} will pass through points with $x_0 = s_0 = 0$. Such points, which lie on the hyperplanes at infinity in both projective spaces, are entirely independent of the choice of \mathbf{c} and \mathbf{g} . Indeed, they are the solutions of the equations

$$(33) \quad \mathbf{s} \in \mathcal{L}_A = \ker B, \quad \mathbf{x} \in \mathcal{L}_A^\perp = \ker A, \quad \text{and} \quad x_1s_1 = x_2s_2 = \cdots = x_ns_n = 0.$$

The solutions to these equations form the *disjoint support variety* in $\mathbb{P}^{n-1} \times \mathbb{P}^{n-1}$, which contains pairs of vectors in the two spaces \mathcal{L}_A and \mathcal{L}_A^\perp whose respective supports are disjoint.

Example 35. Figure 5 shows the primal-dual central curve for the matrices

$$A = \begin{pmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & 1 & -1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}.$$

The disjoint support variety consists of the point $(\mathbf{x}, \mathbf{s}) = ([0 : 0 : 1 : 1], [1 : -1 : 0 : 0])$ in $\mathbb{P}^3 \times \mathbb{P}^3$. After we identify $\mathbb{P}^3 \times \mathbb{P}^3$ with $\{(\mathbf{x}, \mathbf{s}) \in \mathbb{P}^4 \times \mathbb{P}^4 : x_0s_0 = 0\}$, this point lies on the primal-dual central curve and appears as point “c” in Figure 5. \diamond

When studying the global geometry of the primal-dual central curve, it is useful to start with the case when the constraint matrix A is generic. In that case, our matroids are uniform, namely $M(\mathcal{L}_A) = U_{d,n}$ and $M(\mathcal{L}_B) = U_{n-d,n}$, and the disjoint support variety (33) is empty. This condition ensures that the intersections of the curve \mathcal{C} with both the hypersurfaces




FIGURE 5. A primal-dual curve that intersects its disjoint support variety.

$\{x_0s_0 = 0\}$ and $\{x_1s_1 = 0\}$ in $\mathbb{P}^n \times \mathbb{P}^n$ is reduced, zero-dimensional and fully real. The number of points in these intersections is the common number of bases in the two matroids:

$$d_{\mathbf{x}} + d_{\mathbf{s}} = \binom{n-1}{d} + \binom{n-1}{d-1} = \binom{n}{d} = \binom{n}{n-d}.$$

The intersection points of \mathcal{C} with $\{x_0s_0 = 0\}$ are the pairs (\mathbf{x}, \mathbf{s}) where either \mathbf{x} is an analytic center in \mathcal{H} and \mathbf{s} lies at infinity in the dual central curve, or \mathbf{x} lies at infinity in the primal central curve and \mathbf{s} is an analytic center in \mathcal{H}^* . The intersection points of \mathcal{C} with $\{x_1s_1 = 0\}$ are the pairs (\mathbf{x}, \mathbf{s}) where \mathbf{x} is a vertex in \mathcal{H} and \mathbf{s} is a vertex in \mathcal{H}^* . Figure 4 visualizes the above correspondences for the case $n = 4$ and $d = 2$. If we now degenerate the generic matrix A into a more special matrix, then some of the above points representing vertices and analytic centers degenerate to points on the disjoint support variety (33), as in Example 35.

In Theorem 33 we did not mention the degree of the primal or dual central curve. For the sake of completeness, here is a brief discussion of the geometric meaning of the degree $d_{\mathbf{x}}$:

Remark 36. Consider the intersection of the primal central path with a level set $\{\mathbf{c}^T \mathbf{x} = c_0\}$ of the objective function \mathbf{c} . Varying c_0 produces a family of parallel hyperplanes. Each hyperplane meets the curve in precisely $d_{\mathbf{x}}$ points, all of which have real coordinates. These points are the analytic centers of the $(n-d-1)$ -dimensional polytopes obtained as the bounded regions of the induced arrangement of n hyperplanes $\{x_i = 0\}$ in the affine space $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, \mathbf{c}^T \mathbf{x} = c_0\}$. We can see $d_{\mathbf{x}}$ as the number of $(n-d-1)$ -dimensional bounded regions in the restriction of the arrangement \mathcal{H} to a general level hyperplane $\{\mathbf{c}^T \mathbf{x} = c_0\}$. In particular, this gives a one-dimensional family of hyperplanes all of whose intersection points with the central curve are real, as suggested by the left diagram in Figure 3. \diamond

A main theme in this paper was that projective algebraic geometry provides an alternative view on optimality and duality in optimization, as well as powerful tools for analyzing interior point methods. This parallels the discussion of semidefinite programming in [27]. See also [30] for a statistical perspective on analytic centers and central curves in the semidefinite context.

8. ACKNOWLEDGMENTS

We are grateful to Andrew Berget, Antoine Deza, Ragni Piene, Kristian Ranestad, Franz Rendl, Raman Sanyal, Kim Chuan Toh, Michael J. Todd, Yuriy Zinchenko, and the referees for many useful ideas, suggestions, clarifications, and references. Bernd Sturmfels and Cynthia Vinzant acknowledge support by the U.S. National Science Foundation (DMS-0757207 and DMS-0968882). Jesús De Loera was supported by the NSF grant DMS-0914107 and a UC Davis Chancellor Fellow award. He is also grateful to IPAM UCLA, where he was based during the writing of this article.

REFERENCES

- [1] I. Adler: The expected number of pivots needed to solve parametric linear programs and the efficiency of the self-dual simplex method, Technical Report, University of California at Berkeley, CA, 1983.
- [2] A. Bachem and W.Kern: *Linear Programming Duality. An Introduction to Oriented Matroids*, Universitext, Springer Verlag, Berlin, 1992.
- [3] D. Bayer and J. Lagarias: The nonlinear geometry of linear programming. I. Affine and projective scaling trajectories. *Trans. Amer. Math. Soc.* **314**, 499–526, 1989.
- [4] D. Bayer and J. Lagarias: The nonlinear geometry of linear programming. II. Legendre transform coordinates and central trajectories. *Trans. Amer. Math. Soc.* **314**, 527–581, 1989.
- [5] R. Benedetti and M. Dedò: A geometric inequality for the total curvature of plane curves. *Geom. Dedicata* **22**, 105–115, 1987
- [6] A. Berget: Products of linear forms and Tutte polynomials, *European J. Combin.* **31**, 1924–1935, 2010.
- [7] A. Björner: The homology and shellability of matroids and geometric lattices, in *Matroid Applications*, 226–283, Encyclopedia Math. Appl., 40, Cambridge Univ. Press, 1992.
- [8] K.H. Borgwardt: Probabilistic analysis of the simplex method (pp. 21–34) in *Mathematical developments arising from linear programming*. Edited by J. C. Lagarias and M. J. Todd. Contemporary Mathematics, 114. American Mathematical Society, Providence, RI, 1990.
- [9] E. Brugallé and L. López de Medrano: Inflection points of real and tropical plane curves, *Journal of Singularities*, to appear.
- [10] T. Brylawski and J. Oxley: The Tutte polynomial and its applications, in *Matroid Applications*, 123–225, Encyclopedia Math. Appl., 40, Cambridge Univ. Press, 1992.
- [11] J.P. Dedieu, G. Malajovich, and M. Shub: On the curvature of the central path for linear programming theory, *Found. Comput. Math.* **5**, 145–171, 2005.
- [12] A. Deza, T. Terlaky and Y. Zinchenko: A continuous d -step conjecture for polytopes, *Discrete Comput. Geom.* **41**, 318–327, 2009.
- [13] A. Deza, T. Terlaky, and Y. Zinchenko: Polytopes and arrangements: Diameter and curvature. *Operations Research Letters*, **36** (2), 215–222, 2008.
- [14] I. V. Dolgachev: *Classical Algebraic Geometry: A Modern View*, Cambridge University Press, 2012.
- [15] D. Grayson and M. Stillman: Macaulay 2, a software system for research in algebraic geometry, available at www.math.uiuc.edu/Macaulay2/.
- [16] P. Griffiths: On Cartan’s method of Lie groups and moving frames as applied to uniqueness and existence questions in differential geometry, *Duke Math. J.* **41**, 775–814, 1974.
- [17] M. Haimovich: The simplex algorithm is very good!: On the expected number of pivot steps and related properties of random linear programs, Technical report, Columbia University, New York, 1983.
- [18] F. Klein: Eine neue Relation zwischen den Singularitäten einer algebraischer Curve. *Math. Ann.*, **10**, 199–210, 1876.
- [19] E.L. Lawler: *Combinatorial Optimization: Networks and Matroids*, Holt, Rinehart and Winston, New York, Montreal, London, 1976.
- [20] A. Lewis, P. Parrilo and M. Ramana: The Lax conjecture is true, *Proceedings Amer. Math. Soc.* **133**, 2495–2499, 2005.
- [21] N. Megiddo, M. Shub: Boundary behavior of interior point algorithms in linear programming, *Math. of Operations Research*, **14**, No. 1, 97–146, 1989.

- [22] R. Monteiro and T. Tsuchiya: A strong bound on the integral of the central path curvature and its relationship with the iteration complexity of primal-dual path-following LP algorithms, *Mathematical Programming, Series A*, **115**, 105–149, 2008.
- [23] R. Piene: Numerical characters of a curve in projective space, in Real and complex singularities (Proc. Ninth Nordic Summer School Sympos. Math., Oslo, 1976), pp. 475–495, Sijthoff and Noordhoff, 1977.
- [24] N. Proudfoot and D. Speyer: A broken circuit ring, *Beiträge Algebra Geom.* **47**, 161–166, 2006.
- [25] J. Renegar: Hyperbolic programs, and their derivative relaxations, *Found. Comput. Math.* **6**, 59–79, 2006.
- [26] C. Roos, T. Terlaky, and J-Ph. Vial: *Theory and Algorithms for Linear Optimization: An Interior Point Approach*. Springer, New York, second edition, 2006.
- [27] P. Rostalski and B. Sturmfels: Dualities in convex algebraic geometry, *Rend. Mat. Appl. (7)* **30** no. 3-4, 285–327, 2010.
- [28] G. Sonnevend, J. Stoer, and G. Zhao: On the complexity of following the central path of linear programs by linear extrapolation. II. *Mathematical Programming, Series B*, **52**, 527–553, 1991.
- [29] R.P. Stanley: *Combinatorics and Commutative Algebra*. Second edition, Progress in Mathematics, 41, Birkhäuser, Boston, 1996.
- [30] B. Sturmfels and C. Uhler: Multivariate Gaussians, semidefinite matrix completion, and convex algebraic geometry, *Annals of the Institute of Statistical Mathematics* **62**, 603–638, 2010.
- [31] H. Terao: Algebras generated by reciprocals of linear forms, *Journal of Algebra* **250**, 549–558, 2002.
- [32] M. Todd and Y. Ye: A lower bound on the number of iterations of long-step and polynomial interior-point linear programming algorithms, *Annals of Operations Research* **62** 233–252, 1996.
- [33] R.J. Vanderbei: *Linear Programming: Foundations and Extensions*. Third edition. International Series in Operations Research & Management Science, 114. Springer, New York, 2008.
- [34] S. Vavasis and Y. Ye: A primal-dual interior-point method whose running time depends only on the constraint matrix. *Mathematical Programming, Series A*, **74**, 79–120, 1996.
- [35] V. Vinnikov: Self-adjoint determinantal representations of real plane curves. *Math. Ann.* **296** (3), 453–479, 1993.
- [36] T. Zaslavsky: Facing up to arrangements: face-count formulas for partitions of space by hyperplanes. *Mem. Amer. Math. Soc.*, issue 1, no. 154, 1975.
- [37] 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**, 85–103, 1993.

E-mail address: deloera@math.ucdavis.edu, bernd@math.berkeley.edu, vinzant@umich.edu

UNIVERSITY OF CALIFORNIA AT BERKELEY AND DAVIS; UNIVERSITY OF MICHIGAN AT ANN ARBOR