Lattice Subdivisions and Tropical Oriented Matroids Featuring $\Delta_{n-1} \times \Delta_{d-1}$

LindsayCPiechnik

Submitted in partial fulfillment of the

requirements for the degree

of Doctor of Philosophy

in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2011

ABSTRACT

Lattice Subdivisions and Tropical Oriented Matroids
Featuring $\Delta_{n-1} \times \Delta_{d-1}$

LindsayCPiechnik

Subdivisions of products of simplices, and their applications, appear across mathematics. In this thesis, they are the tie between two branches of my research: polytopal lattice subdivisions and tropical oriented matroid theory. The first chapter describes desirable combinatorial properties of subdivisions of lattice polytopes, and how they can be used to address algebraic questions. Chapter two discusses tropical hyperplane arrangements and the tropical oriented matroid theory they inspire, paying particular attention to the previously uninvestigated distinction between the generic and non-generic cases. The focus of chapter three is products of simplices, and their connections and applications to ideas covered in the first two chapters.

Contents

To *everyone who has loved me anyway* .

Chapter 1

Idealic Triangulations

1.1 Introduction

There are many connections between combinatorial properties of lattice polytopes and properties of algebraic objects. A direct translation for these ties is provided by the coordinate/degree map taking a lattice point to a corresponding Laurent monomial.

$$\begin{array}{ccc} \text{Laurent monomial} & & \text{lattice point} \\ \boldsymbol{x}^{\boldsymbol{n}} := x_1^{n_1} \cdot \ldots \cdot x_d^{n_d} \in \Bbbk[x_1^{\pm 1}, \ldots, x_d^{\pm 1}] & \longleftrightarrow & v_n = (n_1, \ldots, n_d) \in \mathbb{Z}^d \end{array}$$

For example,

$$x^3y^4 \quad \longleftrightarrow \quad (3,4) \in \mathbb{Z}^2$$

This map allows commutative algebraists to study various algebraic objects via lattice polyhedral complexes, particularly lattice polytopes and their decompositions. Figure 1.1 shows such a correspondence, the lattice polyhedron corresponding to the ideal generated by $x^3y^4 \subset k[x, y]$.

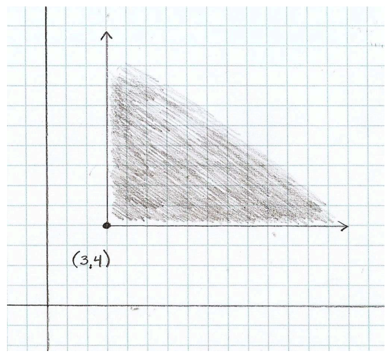

Figure 1.1.1: This shows the lattice polyhedron corresponding to the ideal generated by x^3y^4. $\mathcal{I} =< x^3y^4 >$ corresponds to the staircase diagram consisting of a positive orthant positioned at the point $(3, 4)$.

These algebraic applications include investigating properties of the graded semigroup ring $R_P = \Bbbk[\sigma_P \cap \mathbb{Z}^{d+1}]$, where P is a lattice polytope and σ_P is the pointed polyhedral cone obtained by embedding $P \subset \mathbb{Z}^d$ at height one $\subset \mathbb{Z}^{d+1}$ as $P \times 1 \subset \mathbb{R}^{d+1}$. Figure 1.1 depicts an example, in which P is the 1-dimensional line segment $[0, 1]$. In this context, the polytope P is integrally closed if and only if the domain R_P is generated in degree one. I will say more about how R_P is obtained from P in Section 1.2.1.

People also examine the closely related notion of normality. For normality, one considers the subring $\tilde{R}_P \subseteq R_P$ generated by the degree one piece, and calls P *normal* if $\tilde{R}_P$ is normal (i.e. integrally closed in its quotient field). That is, P is normal if $\Bbbk[\sigma_P \cap \Lambda]$ is generated in degree one, where $\Lambda \subseteq \mathbb{Z}^{d+1}$ is the sublattice generated by $(P \times \{1\}) \cap \mathbb{Z}^{d+1}$ [9, Def. 2.59]. A hierarchy of properties on both sides of this bijection can be found in Section 1.5.1.

My primary research interest in this area is the tie between the Gröbner

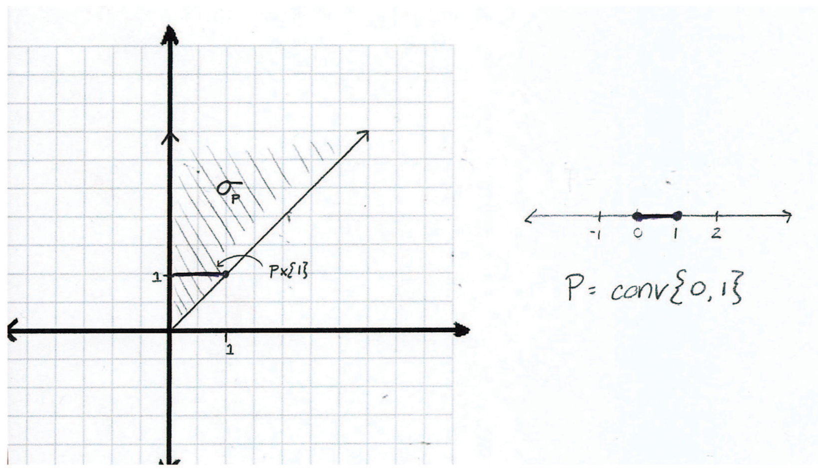

Figure 1.1.2: Here we see a 1-dimensional polytope P, the segment $[0,1]$, on the numberline and the 2-dimensional pointed cone σ_P in the plane.

basis of the defining ideal $\mathfrak{I}_P$ of $\tilde{R}_P = \Bbbk[x_1 \dots x_r]/\mathfrak{I}_P$ (where $r = |P \cap \mathbb{Z}^d|$) and regular triangulations of P. In particular, the Gröbner basis corresponding to a regular unimodular triangulation consists of binomials corresponding to minimal non-faces of the triangulation. I will explain this correspondence in more detail, but most simply stated, the degree of each binomial in the Gröbner basis is equal to the size of the corresponding non-face. Therefore, triangulations provide degree bounds for Gröbner bases, and the existence of such a triangulation guarantees the existence of a quadratic Gröbner basis, which makes the search for quadratic triangulations of particular interest. However, there is a hierarchy of other combinatorial covering and triangulation properties corresponding to algebraic properties via this degree map. This hierarchy also offers a rich source of research questions and corresponding techniques for addressing them when quadratic triangulations can not be found.

1.2 Background

1.2.1 Algebraic Basics

While my work focuses on the combinatorial side of the correspondence between special triangulations and algebraic properties, it is good to have a basic understanding of the algebraic objects being discussed. Therefore, I will give the algebraic objects of interest an independent introduction. (For more on related algebraic topics discussed here, see [16], [19], [29].)

Take a field $\mathbb{k}$ (for example $\mathbb{C}$) and consider $\mathbb{k}[x_1, ..., x_n]$ the ring of polynomials with coefficients in $\mathbb{k}$. An *ideal* in $\mathbb{k}[x_1, ..., x_n]$ is a subset $\mathcal{I}$ of $\mathbb{k}[x_1, ..., x_n]$ which is closed under addition of elements in $\mathcal{I}$ and closed under multiplication by polynomials in the ring (i.e. $g + f \in \mathcal{I}$ for all $f, g \in \mathcal{I}$ and $fg \in \mathcal{I}$ for all $f \in \mathcal{I}$ and $g \in k[x_1, ..., x_n]$).

The Hilbert basis theorem states that any ideal can be obtained by taking all linear combinations of polynomial multiples of some finite set of polynomials, $g_1, ...g_k$, meaning linear combinations of the form $f_1g_1 + ... + f_kg_k$ where the f_i are arbitrary polynomials in $\mathbb{k}[x_1, ..., x_n]$. The resulting ideal, generated by $g_1, ...g_k$, is denoted $\langle g_1, ..., g_k\rangle$.

In a one-variable polynomial ring there is a natural ordering of monomials according to degree. However, determining relative order of monomials in a multivariable setting is not obvious. The question of determining the leading term of a polynomial requires having some such ordering. Having such an ordering is also useful for dividing polynomials in the multivariable setting. In the multivariable case there are many potential term orders, and defining leading terms and divisibility between polynomials requires making a choice. Often a weight vector $\boldsymbol{\omega}$ is used to define an order on the monomials in a multivariable polynomial ring. This ordering is used to determine leading terms and to define the special ideal $\mathcal{I}_{\boldsymbol{\omega}}$, known as the initial ideal. Given an integer vector $\boldsymbol{u} \in \mathbb{Z}^d$, the notation $\boldsymbol{x}^{\boldsymbol{u}}$ denotes the monomial $x_1^{u_1} x_2^{u_2} ... x_d^{u_d}$. In general, a total order $\succ$ on $\mathbb{N}^n$ is a *term order* on the monomials of $\Bbbk[x_1, ..., x_n]$ given by $\boldsymbol{x}^{\boldsymbol{u}} < \boldsymbol{x}^{\boldsymbol{v}}$ for all $\boldsymbol{u} < \boldsymbol{v} \subset \mathbb{N}^n$ if 0 is the unique minimal element and $\boldsymbol{u} < \boldsymbol{v}$ implies $\boldsymbol{u} + \boldsymbol{a} < \boldsymbol{v} + \boldsymbol{a}$ for all $\boldsymbol{a}, \boldsymbol{u}, \boldsymbol{v} \in \mathbb{N}^n$.

Common examples of term orders, which will appear in coming sections, include lexicographic, degree lexicographic, and graded reverse lexicographic. In *lexicographic order*, the variables are ordered by $x_1 > x_2 > ... > x_n$ and $x_i > x_j^k$ for all k whenever $i < j$. For *degree lexicographic order*, the variables are ordered in the same manner, but monomials are first sorted by total degree and then by their variable composition within degree. For example, in purely lexicographic order $x_1 x_2^3 x_3 > x_2^2 x_3^4$ because the first monomial contains x_1 and the second does not. However, in degree lexicographic order, $x_2^2 x_3^4 > x_1 x_2^3 x_3$ as $x_2^2 x_3^4$ has total degree 6 and $x_1 x_2^3 x_3$ has total degree 5. *Graded reverse lexicographical order* again first compares monomials by total degree, but the smallest degree is taken first and ties are broken with a reverse ordering of the variables (this is often described as sorting from right to left rather than left to right).

A total order $\succ_\omega$ can be defined by a real *weight vector* $\boldsymbol{\omega} = (\omega_1, ..., \omega_d) \in \mathbb{R}^d$ on the set of monomials in $\Bbbk[x_1, ..., x_n]$ by declaring $\boldsymbol{x}^{\boldsymbol{u}} \succ_\omega \boldsymbol{x}^{\boldsymbol{v}}$ if $\langle \boldsymbol{u}, \boldsymbol{\omega} \rangle > \langle \boldsymbol{v}, \boldsymbol{\omega} \rangle$ breaking ties by lexicographical order (i.e. a component-wise comparison of $\boldsymbol{u}$ and $\boldsymbol{v}$ where the greater of the first non-equal position determines the greater monomial).

Having an ordering on the monomials of a multivariable polynomial ring is not just a nice tool for determining a preferred order for the terms of a polynomial and deciding when one polynomial divides another, it is also fundamental for defining Gröbner bases. Gröbner bases are generating sets of ideals which satisfy special properties with respect to a given term order, and have many applications. Their fundamental use is determining membership in a monomial ideal. However, they can also be viewed as a multivariable non-linear generalization of the Euclidean algorithm, [1] Gaussian elimination, [2] and integer programming problems.

Given an ordering $\succ_\omega$, the *initial term* of a polynomial f, denoted $in_{\succ_\omega}(f)$ is the maximal term of the monomials of f with respect to the *term order* $\succ_\omega$. The *initial ideal* of an ideal $\mathcal{I}$ in $\mathbb{k}[x_1, ..., x_n]$ with respect to a term order $\succ_\omega$, denoted $in_{\succ_\omega}(\mathcal{I})$, can be defined as the ideal generated by the initial monomials of polynomials in $\mathcal{I}$. A finite set of polynomials $G = \{g_1, ...g_k\} \subset \mathcal{I}$ is a *Gröbner basis* of $\mathcal{I}$ with respect to $\succ_\omega$ if $in_{\succ_\omega}(I)$ is generated by $\{in_{\succ_\omega}(g_1), ..., in_{\succ_\omega}(g_k)\}$. Monomials of $\mathbb{k}[x_1, ..., x_n]$ not contained in $in_{\succ_\omega}(\mathcal{I})$ are called *standard*. These standard monomials form a basis for $R/\mathcal{I}$.

Monomials needn't be contained in G to be in $in_{\succ_\omega}(\mathcal{I})$; in particular, we see that G is just a set of monomials such that the leading term of every polynomial in $\mathcal{I}$ with respect to $\succ_\omega$ is divisible by some monomial in G. Note that I have made no claim that G's cardinality is minimal. In general, Gröbner bases are not unique. However, for every ideal $\mathcal{I}$ and term order $\succ$ there is a unique reduced Gröbner basis [43]. A *reduced Gröbner basis* is a Gröbner basis G such that for each $g \in G$, the coefficient of $in_\succ(g)$ is

[1]Recall that the Euclidean algorithm is used to compute greatest common divisors in the single variable setting. Gröbner bases calculations take on this role in the multivariable setting.

[2]Gaussian elimination is used to simplify systems of linear equations. Gröbner bases calculations generalize this to the non-linear case.

1, the set $\{in_{\succ}(g)\colon g \in G\}$ minimally generates $in_{\succ}(\mathcal{I})$ (i.e. nothing can be removed), and for each $g \in G$ the only term of g appearing in $in_{\succ}(\mathcal{I})$ is $in_{\succ}(g)$. [3] My interest is in the properties of the terms of the polynomials $g \in G$.

Of particular relevance to my work is the case of toric ideals. To define a toric ideal, consider a $d \times n$ integer matrix A with columns $\boldsymbol{a}_1, ...\boldsymbol{a}_n$, a vector $\boldsymbol{u} = (u_1, ...u_n) \in \mathbb{Z}^n$, and $A\boldsymbol{u} = u_1\boldsymbol{a}_1 + ... + u_n\boldsymbol{a}_n$. For any vector $\boldsymbol{u} \in \mathbb{Z}^n$ the support of $\boldsymbol{u}$ is $\mathrm{supp}(\boldsymbol{u}) = \{i \mid u_i \neq 0\}$, and $\boldsymbol{u}$ can be expressed uniquely as $\boldsymbol{u} = \boldsymbol{u}^+ - \boldsymbol{u}^-$, where $\boldsymbol{u}^+$ and $\boldsymbol{u}^-$ are non-negative and have disjoint support. The *toric ideal of* A is given by:

$$\mathcal{I}_A = \langle \boldsymbol{x}^{\boldsymbol{u}^+} - \boldsymbol{x}^{\boldsymbol{u}^-} \mid A\boldsymbol{u} = 0 \rangle$$

Thinking of A as a point (or vector) configuration, the binomials $x^{\boldsymbol{u}^+} - x^{\boldsymbol{u}^-}$ in $\mathcal{I}_A$ correspond to dependencies $\boldsymbol{u} = \boldsymbol{u}^+ - \boldsymbol{u}^-$ of A.

Gröbner bases of toric ideals can also be interpreted in terms of lattice points of the polyhedra given by $A\boldsymbol{u} = \boldsymbol{b}$ for $\boldsymbol{u} \geq 0$. This wonderful translation is the tie to my work in triangulations.

1.2.2 Necessary Triangulation Information

Making use of the connections between triangulations of lattice polytopes and algebraic properties of associated rings and ideals requires some background in lattice polytopes and their triangulations.

A *lattice polytope* in $\mathbb{R}^d$ is the convex hull of finitely many points in the lattice $\mathbb{Z}^d$. Two lattice polytopes are said to be *lattice equivalent* if they

[3] There are algorithms for obtaining a reduced Gröbner basis of an ideal $\mathcal{I}$ given a term order $\succ$ (see [43]).

are related by a lattice preserving affine map. That is, a map of the form $\boldsymbol{x} \to A\boldsymbol{x} + c$, which preserves the lattice. Up to this equivalence, one can assume the lattice polytopes discussed here are d-dimensional. There are books devoted to the study of convex and lattice polytopes, including [48] and [6].

A d-dimensional *lattice simplex* is the convex hull of $d+1$ points in $\mathbb{Z}^d$. The *standard simplex* Δ^d is the convex hull of the origin $\mathbf{0}$ and each of the standard unit vectors $\boldsymbol{e}_i$ $(1 \le i \le d)$. A *unimodular simplex* is a lattice simplex which is lattice equivalent to the standard simplex. Equivalently, a unimodular simplex can be described as a d-dimensional lattice polytope of minimal possible Euclidean volume, $1/d!$.

For my purposes, a lattice *subdivision* of a d-dimensional lattice polytope P is a finite collection of lattice polytopes $\mathcal{S}$ satisfying the 3 following properties:

1. every face of a member of $\mathcal{S}$ is itself a member of $\mathcal{S}$,
2. any two elements of $\mathcal{S}$ intersect in a common face (possibly empty),
3. the union of the polytopes in $\mathcal{S}$ is P.

The d-dimensional polytopes in $\mathcal{S}$ are called the cells of the subdivision.

A *triangulation* of a lattice polytope is a subdivision in which each cell is a simplex, and a triangulation is *unimodular* if every cell (i.e. simplex) is unimodular. Figure 1.2.1 depicts three triangulations of the 9-point square. The first is not unimodular, but the other two are.

A *full triangulation* is a lattice triangulation that uses every lattice point in P. The right and center triangulations in Figure 1.2.1 are full, but the triangulation on the left is not. Any subdivision can be refined to a full

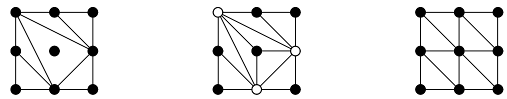

Figure 1.2.1: A non-unimodular triangulation and two unimodular ones. They are all regular, but only the last is quadratic.

triangulation. One procedure for obtaining such a refinement is pulling, which is discussed various places including Section 1.3.

Every unimodular triangulation is full, and in dimension ≤ 2 it is also true that every full triangulation is unimodular. [4] However, this nice property already fails in 3 dimensions, where there are polytopes which fail to admit any unimodular triangulation. The tetrahedron in Figure 1.2.2 contains only its vertices as lattice points. Therefore, its only lattice triangulation is the trivial one, and since its Euclidean volume is $q/6$, this simplex does not have a unimodular triangulation when $q > 1$.

$$\operatorname{conv}\begin{bmatrix} 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & q \end{bmatrix}$$

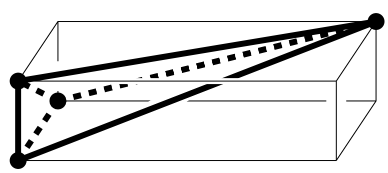

Figure 1.2.2: This family of *Reeve's tetrahedra* is credited to John Reeve. [35]. This class of examples demonstrates the existence of empty tetrahedra of arbitrarily large volume.

Formally, a subdivision is *regular* if the cells are the domains of linearity of a convex piecewise linear function (see [24, Section 14.3]). Less formally, an

[4]This two-dimensional case can be viewed as a result of, or evidence for, Pick's formula, which says that the area of a polygon is one less than its number of interior lattice points plus half the number of lattice points on its boundary.

intuitive notion of a regular triangulation (or subdivision) can be achieved by thinking of a regular triangulation (or subdivision) as one that can be obtained via a convex "folding" of the polytope (Figure 1.2.3 on the left). The 3 triangulations in Figure 1.2.1 are all regular, but the one on the right in Figure 1.2.3 is not.

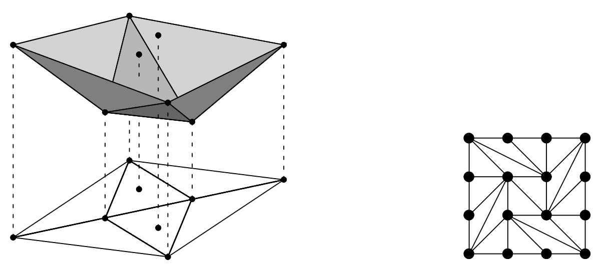

Figure 1.2.3: The image on the left provides a visual of how a regular subdivision can be obtained via a lifting or "folding." The image on the right is a two-dimensional example of a non-regular unimodular triangulation.

One way of constructing a regular subdivision of P is to specify heights (or weights) $\boldsymbol{\omega} \in \mathbb{R}^{\mathcal{A}}$ where $\mathcal{A} = P \cap \mathbb{Z}^d$ is the set of lattice points in P. Letting the polyhedron $\tilde{P} = \operatorname{conv}(\boldsymbol{a} \times [\boldsymbol{\omega}, \infty) \ : \ \boldsymbol{a} \in \mathcal{A})$ in $\mathbb{R}^{d+1}$, the lower/bounded faces of $\tilde{P}$ project to a subdivision of P. These are the domains of linearity of the function $\boldsymbol{x} \mapsto \min\{h \ : \ (\boldsymbol{x}, h) \in \tilde{P}\}$. This lifting procedure is depicted in Figure 1.2.3.

A set of lattice points whose convex hull does not form a face of a given triangulation is called a *non-face* of that triangulation. Of particular interest are *minimal non-faces*. These are collections of points that do not form faces themselves, but for which every proper subset of the set does form a face. Like the list of cells (typically described by the sets of points that form them), the list of minimal non-faces completely characterizes a triangulation.

If all minimal non-faces of a triangulation contain two elements, it is a *flag*

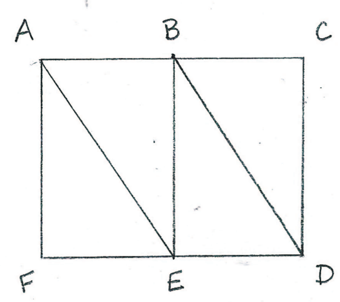

Figure 1.2.4: This shows a triangulation of the 6-point rectangle. Its minimal non-faces are $\{A, D\}$, $\{A, E\}$, $\{A, C\}$, $\{B, D\}$, $\{B, F\}$, $\{C, A\}$, $\{C, E\}$, $\{C, F\}$, and $\{D, F\}$. The triangulation can be reconstructed from this list of non-faces alone.

triangulation. A *quadratic triangulation* is defined to be one which has all of these nice properties, meaning a triangulation that is regular, unimodular, and flag. The rightmost triangulation in Figure 1.2.1 is quadratic. However, in the center triangulation, the 3 white vertices form a minimal non-face, so that triangulation is not quadratic. As already noted, the triangulation on the left isn't unimodular, so it too fails to be quadratic.

1.3 Constructing Nice Triangulations

As I will explain further in the next section, having these nice, special triangulations is desirable. This is why people look for them and have techniques for constructing and searching for them.

1.3.1 Pulling

One useful technique for constructing regular triangulations is known as pulling. Two different versions of pullings are used and appear in publication. The first is known as strong and the second is known as weak, but both have desirable properties.

Given a subdivision S of a polytope P and lattice point $\boldsymbol{m} \in P \cap \mathbb{Z}^d$, the *strong pulling refinement* $pull_{\boldsymbol{m}}S$ is obtained from S by replacing every face $F \in S$ containing $\boldsymbol{m}$ by the pyramids $conv(\boldsymbol{m}, F')$ where F' runs over all faces of F not containing $\boldsymbol{m}$. The following properties of strong pulling refinements make them a useful tool for constructing nice triangultaions.

Lemma 1.3.1. *[21]*

(1) Strong pulling preserves regularity.

(2) Strongly pulling all lattice points in P in some order results in a full triangulation.

(3) If only vertices of P are pulled, then every maximal cell is the join of the first pulled vertex $\boldsymbol{v}_1$ with a maximal cell in the pulling subdivisions of the facets not containing $\boldsymbol{v}_1$.

The important message of this lemma is that every regular lattice subdivision of a lattice polytope has a regular refinement which is a full triangulation. If you start with the trivial triangulation, strong pulling can be used to construct a full regular triangulation of any lattice polytope.

Proof. (1) Take a regular subdivision S of P, given by weights $\boldsymbol{\omega} \in \mathbb{R}^{\mathbf{A}}$, where ($\mathbf{A} = P \cap \mathbb{Z}^d$) and $\tilde{P} = conv(\boldsymbol{a} \times [\omega_{\boldsymbol{a}}, \inf) \mid \boldsymbol{a} \in \mathbf{A})$, and $\boldsymbol{m} \in \mathbf{A}$. Set $\omega'_{\boldsymbol{m}} = min\{h \mid (\boldsymbol{m}, h) \in \tilde{P}\} - \epsilon$ and $\omega'_{\boldsymbol{a}} = \omega_{\boldsymbol{a}}$ for all $\boldsymbol{a} \in \mathbf{A} \setminus \{\boldsymbol{m}\}$. For small enough $\epsilon > 0$, the strong pulling $pull_{\boldsymbol{m}}(S)$ is induced by the weights $\boldsymbol{\omega}'$.

(2) Each face of $pull_{\boldsymbol{m}}$ containing $\boldsymbol{m}$ is a pyramid with apex $\boldsymbol{m}$. Given $Q \in S$ with apex $\boldsymbol{n}$, any face of $pull_{\boldsymbol{m}}$ contained in Q that contains $\boldsymbol{n}$ still has $\boldsymbol{n}$ as an apex. Strongly pulling all lattice points makes all lattice points vertices of the subdivision. Each cell has each of its vertices as apices. This makes them all simplices, and we have a full triangulation.

(3) Applying the argument used in the proof of (2) to the trivial subdivision of P yields a subdivision with $\boldsymbol{v}_1$ as an apex of each cell.

□

The second notion of pulling comes up when considering subdivisions of point configurations. Basically, in such settings each face $F \in S$ remembers which points yield F as their convex hull, which may or may not be all lattice points in F.

To illustrate the difference between weak and strong pullings, consider two regular subdivisions $\mathcal{S}_1$ and $\mathcal{S}_2$ of $\mathcal{A} = \{1, 2, 3, 4\} \subset \mathbb{Z}^1$. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be the subdivisions realized by the weight vectors $\boldsymbol{\omega}_1 = (0, 0, 0, 1)$ and $\boldsymbol{\omega}_2 = (0, 1, 0, 1)$ respectively. Both consist of the segments $[1, 3]$ and $[3, 4]$. In $\mathcal{S}_1$ $[1, 3]$ is the convex hull of $1, 2$, and 3, but in $\mathcal{S}_2$ it appears as the convex hull of just 1 and 3. When considering all subdivisions of $\mathcal{A}$, it makes sense to differentiate between $\mathcal{S}_1$ and $\mathcal{S}_2$, since $\mathcal{S}_1$ is not a triangulation and $\mathcal{S}_2$ is a triangulation that refines $\mathcal{S}_1$. Weak pullings do this by removing all points in the pyramid $\operatorname{conv}(\boldsymbol{m}, F')$, but not the points in the base F' itself, from the list of points to be pulled in the future when pulling at $\boldsymbol{m}$. In the context of toric algebra (described in 1.2.1), the weak pulling triangulation corresponds to taking the reverse lexicographic term order. Figure 1.3.1 illustrates an example of a strong and a weak pulling.

Like strong pullings, weak pullings also preserve regularity, and pulling

Figure 1.3.1: The left pulls $\boldsymbol{m}_1$, then $\boldsymbol{m}_2$, then $\boldsymbol{m}_3$ strongly, and the right pulls them weakly in the same order. Since $\boldsymbol{m}_2$ is contained in a face containing $\boldsymbol{m}_1$, pulling it weakly after $\boldsymbol{m}_1$ has no impact. This is not the case with $\boldsymbol{m}_3$, since $\boldsymbol{m}_3$ is not contained in a proper face also containing $\boldsymbol{m}_1$.

all lattice points of a polytope weakly results in a regular triangulation. In fact, when all lattice points of P are vertices of the subdivision S, weak and strong pullings agree. In Figure 1.3.1, the fact that $\boldsymbol{m}_2$ and $\boldsymbol{m}_3$ were not vertices of the original subdivision (the empty 6-point rectangle) explains why their weak and strong pullings do not agree. When this is the case (as in the remainder of the work presented here), the term pulling can and will be used unambiguously without specification.

Not all full triangulations obtained by pulling are unimodular, so polytopes for which every weak pulling triangulation is unimodular have a name. They are called *compressed* [21]. This definition of compressed is from Stanley [42], but there are other equivalent descriptions of such polytopes. Given the usefulness of the property and the fact that more classes of polytopes satisfy this condition than one would expect, it is worth saying a bit more.

Another useful characterization of compressed polytopes is based on their "width," a notion that will appear again later. Given a lattice polytope P with facet defining inequalities $\langle \boldsymbol{y}_i, \boldsymbol{x} \rangle \leq c_i$, for primitive integral $\boldsymbol{y}_i$, the *width* of P with respect to the i-facet of P is $max\langle \boldsymbol{y}_i, P\rangle - min\langle \boldsymbol{y}_i, P\rangle$. This can also be referred to as the width with respect to $\boldsymbol{y}_i$. A polytope is width one with respect to a given facet if it lies entirely between the hyperplane containing that facet and the next parallel lattice hyperplane in that direction. Armed with this definition, 3 equivalent characterizations of compressed polytopes can be stated.

Theorem 1.3.2. [*Santos MSRI 1997 (unpublished)*] *[31] [45]* [5]

Given a lattice polytope P, the following are equivalent.

1. *P is compressed.*

2. *P has width one with respect to each of its facets.*

3. *P is lattice equivalent to the intersection of a unit cube C_n with some affine space.*

Proof. For the purposes of the proof, I introduce formal notation for the P's minimal linear description, $P = \{\boldsymbol{x} \in \mathbb{R}^d \colon \boldsymbol{a}_i^T\boldsymbol{x} \geq b_i, i = 1, .., n\}$ and name the lattice $\mathcal{L}$. With this description 2 can be restated more technically by saying that for each i there is at most 1 non-zero $m_i \in \mathbb{R}$ such that $\{\boldsymbol{x} \in \mathcal{L} \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i + m_i\} \cap P$ is non-empty.

For (1) $\to$ (2) let P be compressed and assume that there is some i for which both m and m' satisfy the condition. That is both $\{\boldsymbol{x} \in \mathcal{L} \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i + m\} \cap P$ and $\{\boldsymbol{x} \in \mathcal{L} \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i + m'\} \cap P$ are non-empty. With out loss of generality, assume $m > m'$. Take $\boldsymbol{p}_m \in \{\boldsymbol{x} \in \mathcal{L} \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i + m\} \cap P$ and $\boldsymbol{p}_{m'} \in \{\boldsymbol{x} \in \mathcal{L} \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i + m'\} \cap P$. Consider the two pulling triangulations obtained by pulling $\boldsymbol{p}_m$ first and by pulling $\boldsymbol{p}_{m'}$ first, and taking the same ordering of the lattice points of the facet $F = \{\boldsymbol{x} \in \mathbb{R}^d \colon \boldsymbol{a}_i^T\boldsymbol{x} = b_i\} \cap P$. Given a simplex σ in the pulling triangulation of F, since $m > m'$ the ratio

[5] This theorem has 3 sources. It was first an unpublished result of Santos's from MSRI in 1997, but the version of the proof presented here is closer to Sullivant's in [45].

of volumes $\mathrm{Vol}(\boldsymbol{p}_m \cup \sigma)$ / $\mathrm{Vol}(\boldsymbol{p}_{m'} \cup \sigma)$ must be greater than 1. This means the pulling triangulation of P that pulls $\boldsymbol{p}_m$ first can not be unimodular. This contradicts the assumption that P was compressed.

For (2) $\rightarrow$ (3) assume P satisfies 2. P is a lattice polytope, so this condition ensures that each lattice point in P is in fact a vertex of P. It is clear that, to have a proper interior lattice point, P would have to have lattice width greater than one in some direction. Having a non-vertex lattice point on some facet of P would imply that the facet itself was not lattice width one in some direction, but clearly by induction, if P is facet width one, then each of its facets must be as well. This means that each lattice point of P is a vertex of P. To prove that P is lattice equivalent to an intersection of a unit cube with some affine space, it is only necessary to consider the case where P does not lie in an affine subspace. For if P does lie in an affine subspace, a unimodular change of coordinates projecting to a lower dimensional space would do the job. Assuming P does not lie in an affine subspace implies that in the technical description of 2 there is exactly 1 non-zero m_i for each i. These m_i's are used in the linear transformation. Consider $\varphi \colon \mathbb{R}^d \to \mathbb{R}^n$ defined by

$$\varphi(\boldsymbol{x}) = ((\boldsymbol{a}_1^T \boldsymbol{x} - b_1)/m_1, ..., (\boldsymbol{a}_n^T \boldsymbol{x} - b_n)/m_n).$$

Since φ maps every vertex of P to a 0/1 vector, $\varphi(P)$ is a 0/1 polytope. φ sends each of the facet defining inequalities $\boldsymbol{a}_i^T \boldsymbol{x} \geq b_i$ to $\boldsymbol{y}_i \geq \boldsymbol{0}$. Therefore, a point $p \in \varphi(P)$ if and only if $\boldsymbol{p}$ is contained in the n-cube and in the affine span of the image of P's vertices, and 3 is satisfied.

For (3) $\rightarrow$ (1) assume lattice polytope P satisfies 3 and let $\varphi(P) = Q$ be the image of the corresponding affine transformation. Since the transformation maps lattice points of P to integer points in Q, and P and Q are otherwise isomorphic, P is compressed if and only if Q is. Therefore, it is enough to show that any integral polytope Q of the form $Q = C_n \cap \{\boldsymbol{x} \colon A\boldsymbol{x} = \boldsymbol{b}\}$ is compressed. The proof follows by induction on Q's dimension. When $dim(Q) = 0$ there is nothing to show. For $dim(Q) = d$, choose an ordering of Q's vertices. Take the first vertex $\boldsymbol{p}$ and construct the pulling triangulation, by taking the pulling triangulation of each facet of Q not containing P and coning these triangulations over $\boldsymbol{p}$. We know that the orthogonal distance between each simplex σ of the facet triangulations and $\boldsymbol{p}$ is 1, so the normalized volume of each $p \cup \sigma$ is the normalized volume of σ. Each facet of Q is a $(d-1)$-polytope of the form $C_n \cap \{\boldsymbol{x} \colon A\boldsymbol{x} = \boldsymbol{b}, x_i = 0\}$ for some i, and therefore by induction compressed. Each triangulation of each facet is unimodular, meaning each simplex σ has normalized volume 1. This means each simplex of the pulling triangulation of Q has normalized volume 1, and is hence unimodular. Since the choice of pulling triangulation was arbitrary, every pulling triangulation of Q must be unimodular, which means Q is compressed.

□

Some well-known classes of polytopes are compressed (see [21]). Their nice properties also prove useful in triangulating some larger polytopes. For example, facet unimodular polytopes are an example class of polytopes which can be subdivided into compressed polytopes with nice results.

A matrix comprised of vectors in $\mathbb{Z}^d$ for which every $(d \times d)$ -minor is -1, 0, or 1 is called *unimodular*. Let $\mathbf{A} = \{\boldsymbol{n}_1, ...\boldsymbol{n}_r\}$ be a collection of vectors spanning $\mathbb{R}^d$ that form such a matrix. This set subdivides $\mathbb{R}^d$ into regular lattice polytopes by inducing an infinite arrangement of hyperplanes, $\{\boldsymbol{x} \in \mathbb{R}^d \mid \langle \boldsymbol{n}_i, \boldsymbol{x} \rangle = k\}$ for $i \in [1, r]$ and $k \in \mathbb{Z}$. Lattice subdivisions of

$\mathbb{R}^d$ that can be obtained this way are called lattice dicings [15]. A polytope P whose primitive facet normals [6] form a unimodular matrix is called *facet unimodular*. Induction shows that each face of a facet unimodular polytope is itself facet unimodular in its lattice. The lattice dicing via the described hyperplanes subdivides P into what is known as the *canonical subdivision* of a facet unimodular polytope. As one would hope, this canonical subdivision subdivides faces canonically as well, yielding the following straightforward result.

Theorem 1.3.3. *[21]*

Every facet unimodular polytope, $P \subset \mathbb{R}^d$ has a regular unimodular triangulation.

Proof. By construction, each dicing cell is width one with respect to each facet direction, so each cell is compressed and any pulling refinement of this canonical subdivision must be unimodular.

□

Similar methods have been used to find unimodular triangulations for a number of classes of polytopes (see [21] for more examples). We will see it again in Section 3, where it is applied to triangulate dilated polytopes.

[6] A vector is primitive if the greatest common divisor of its entries is 1. A primitive facet normal is a vector orthogonal to the facet scaled to satisfy this condition.

1.4 Push-Forward and Pull-Back Subdivisions

Projections are another powerful tool in the search for triangulations. When applicable, push-forward and pull-back subdivsions can simplify the search for a triangulation of a given polytope to a lower dimensional question of triangulating a projection of the polytope. This is a useful technique, as the complexity of searching for triangulations increases greatly as dimension increases.

1.4.1 Chimney Polytopes and Pull-Back Subdivsions

Chimney polytopes and pull-back subdivisions offer a method for recursively constructing a unimodular triangulation of a polytope, by finding unimodular triangulations of prisms over a unimodular triangulation of a projection of the original polytope. Explaining this requires definitions of both chimney polytopes and pull-back subdivisions.

Given a lattice polytope $Q \subset \mathbb{R}^d$, and linear functionals l and u such that $l \leq u$ on Q, the *chimney polytope* associated with Q, l, and u is

$$P(Q, \boldsymbol{l}, \boldsymbol{u}) = \{(\boldsymbol{x}, y) \in \mathbb{R}^d \times \mathbb{R} \mid \boldsymbol{x} \in Q, \boldsymbol{l}(x) \leq y \leq \boldsymbol{u}(\boldsymbol{x})\}.$$

Q is called the *base* of P, and is itself a lattice polytope. Figure 1.4.1 offers a visual. The key to using projections to find unimodular triangulations is the fact that a chimney polytope has a unimodular triangulation if its base has one. This is where the pull-back subdivision is needed.

To define the pull-back subdivision, consider a lattice polytope $P \subset \mathbb{R}^d$, a projection $\pi \colon \mathbb{R}^d \to \mathbb{R}^{d'}$, with $P' = \pi(P)$, and a subdivision $\mathcal{S}'$ of P'. Intersecting the infinite prism $\pi^{-1}(\sigma)$ over each cell $\sigma \in \mathcal{S}'$ with P gives the *pull-back subdivision* $\pi^*\mathcal{S}'$ of P.

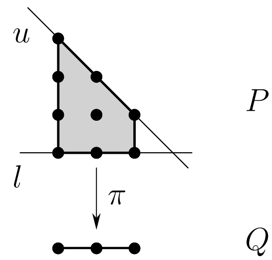

Figure 1.4.1: This shows a chimney construction for which $l = 0$ and $u(x) = 3 - x$ on $Q = [0, 2]$.

Theorem 1.4.1. *[21] Given integral linear functionals $\boldsymbol{l}, \boldsymbol{u}$ such that $\boldsymbol{l} \leq \boldsymbol{u}$ along Q and the corresponding chimney polytope $P = P(Q, \boldsymbol{l}, \boldsymbol{u})$, P has a (regular) unimodular triangulation if Q does.*

Proof. Let $\mathcal{T}$ be the triangulation of Q. This defines a pull-back subdivision of P into chimney polytopes over the simplices of $\mathcal{T}$. Weights demonstrating that this pull-back subdivision is regular can be obtained by "pulling-back" the $\boldsymbol{\omega}$ that showed $\mathcal{T}$ was regular in Q to P. This is done by assigning each lattice point $(\boldsymbol{a}, h)$ in P the weight that its projection, $\boldsymbol{a}$, had in Q.

Any maximal triangulation of such a prism is unimodular, so subdividing each prism $\pi^*\mathcal{T}$ of P in this way yields a unimodular triangulation of P, and using pullings for this ensures that regularity can be preserved. $\square$

This theorem means that if there is a unimodular transformation taking a lattice polytope P to a chimney polytope as described above, projecting to Q and searching for a unimodular triangulation of Q is a way of looking for a regular unimodular triangulation of P. Iterations of this process can be used to prove the existence of regular unimodular triangulations of high-dimensional polytopes by examining their low-dimensional projections.

This inductive pull-back subdivision method has direct application to the class of Nakajima polytopes. Their recursive definition says that a polytope $P \subset \mathbb{R}^d$ is *Nakajima* if P is a single lattice point in $\mathbb{Z}^d$ or

$$P = \{(\boldsymbol{x}, x_d) \mid 0 \leq x_d \leq \ell(\boldsymbol{x}) \; \forall \, \boldsymbol{x} \in F\},$$

where F is a facet of P, and itself a Nakajima polytope, and ℓ is an integral linear functional on $\mathbb{Z}^{d-1}$ that takes non-negative values on $F \cap \mathbb{Z}^{d-1}$.

Theorem 1.4.2. *[21] Every Nakajima polytope has a regular unimodular triangulation.*

Proof. Let P be a Nakajima polytope. If $\dim(P) \leq 2$ the statement is clear. Consider $d = \dim(P) \geq 3$. In this case, projecting P in its last coordinate yields a Nakajima polytope of dimension $d - 1$. Induction says that this projection has a regular unimodular triangulation T', and so the pull-back of T' gives a regular unimodular triangulation of P. $\square$

1.4.2 Push-Forward Subdivision

This method of pulling-back triangulations of projections to find regular unimodular triangulations of higher dimensional lattice polytopes can be generalized for polytopes with multiple functionals bounding them from above and/or below. The difference in this case is that care must be taken in the triangulation of Q to make sure that it respects the projection of the ridges formed by the intersection of upper and/or lower facets of P. This process is formally described using the push-forward subdivision.

Given a lattice polytope $P \subset \mathbb{R}^d$ with subdivision $\mathcal{S}$ and a projection $\pi \colon \mathbb{R}^d \to \mathbb{R}^{d'}$, the *push-forward subdivision* $\pi_*\mathcal{S}$ of $P' = \pi(P)$ is the common refinement of all faces of $\mathcal{S}$. With this definition, the formal generalization of Theorem 1.4.1 can be stated.

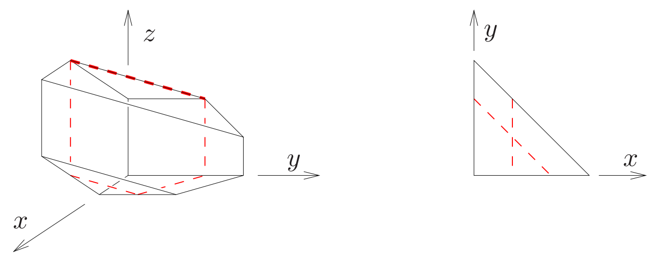

Figure 1.4.2: This figure shows the projections P_{xyz} and P_{xy} in the example described in 1.4.1.

Theorem 1.4.3. *[21]*

If P is a lattice polytope defined by

$$P = \{(\boldsymbol{x}, y) \in Q \times \mathbb{R} \ : \ \boldsymbol{l}_i(\boldsymbol{x}) \leq y \leq \boldsymbol{u}_j(\boldsymbol{x}), \ \textit{for } 1 \leq i \leq r, \ 1 \leq j \leq s\}$$

where $\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r$ and $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_s$ are integral linear functionals such that $\boldsymbol{l}_i \leq \boldsymbol{u}_j$ for $1 \leq i \leq r$, $1 \leq j \leq s$ along the lattice polytope Q, and $\mathcal{S}$ is a (regular) subdivision of P whose push-forward to Q has a (regular) unimodular refinement, then $\mathcal{S}$ has a (regular) unimodular refinement.

For an example of this, consider the 4-dimensional polytope P defined by the following inequalities in x, y, z, and w.

$$\begin{array}{rcccl} 0 & \leq & x & & \\ 0 & \leq & y & \leq & 3 - x \\ 0 & \leq & z & & \\ x - 1 & \leq & z & & \\ 0 & \leq & w & \leq & 2 + x - z \\ & & w & \leq & 4 - y - z \end{array} \tag{1.4.1}$$

The inequalities have been ordered so that each variable is bounded from above, or below, by integral linear functions in the previous variables. P can be projected to 3-dimensional P_{xyz} with defining inequalities:

$$\begin{array}{rcccl} 0 & \leq & x & & \\ 0 & \leq & y & \leq & 3-x \\ 0 & \leq & z & \leq & 2+x \\ x-1 & \leq & z & \leq & 4-y\ . \end{array}$$

This projection is the image on the left in Figure 1.4.2. Note that the facets $z \leq 2+x$ and $z \leq 4-y$ do not pull-back to facets of P. These inequalities are implied by $0 \leq w$ and $w \leq 2+x-z$ and $w \leq 4-y-z$. Pushing-forward the trivial subdivision of P subdivides P_{xyz} along the plane $x+y=2$, which is the projection of the ridge formed by the intersection of the upper bounds of w in P's defining inequalities:

$$\left.\begin{array}{rcccl} 0 & \leq & w & \leq & 2+x-z \\ & & w & \leq & 4-y-z \end{array}\right\} x+y=2\ .$$

The intersection of this hyperplane with P_{xyz} is the convex hull of the lattice points $(1,1,0)$, $(0,2,0)$, $(0,2,2)$, and $(2,0,1)$, so this is a lattice subdivision.

Theorem 1.4.3 says that finding a regular unimodular triangulation of P_{xyz} will yield one of P as well. However, P_{xyz} can be projected again, and finding regular unimodular triangulations in 2-dimensions is a much nicer task. Projecting P_{xyz} to the x-y-plane results in Pxy with defining inequalities:

$$\begin{array}{rcccl} 0 & \leq & x & & \\ 0 & \leq & y & \leq & 3-x. \end{array} \tag{1.4.2}$$

P_{xy} is depicted on the right in Figure 1.4.2, where the push-forward subdivsion of P_{xyz} along $x+y=2$ is shown. In order to obtain a regular unimodular triangulation of P, it is only necessary to find one of this lattice triangle that respects its lattice subdivision along the lines $x+y=2$ and $x=1$ (which themselves intersect in the point $(1, 1)$). Thus, Theorem 1.4.3 has reduced the search for a regular unimodular triangulation of this 4-dimensional polytope to a manageable hand computation.

This method of pull-back and push-forward subdivisions has been particularly successful on the class of smooth reflexive polytopes. A *reflexive* lattice polytope is a lattice polytope P which contains exactly one interior lattice point and for which each facet is lattice distance one from that lone interior point. Without loss of generality, the interior point can be assumed to be the origin. A lattice polytope is *smooth* if every cone in its normal fan is unimodular. [7] Given a polytope $P \subset \mathbb{R}^d$ such that $\mathbf{0} \in \operatorname{int}(P)$, the polar polytope $P^\vee$ of P is defined as

$$P^\vee = \left\{ \boldsymbol{u} \in \mathbb{R}^d \mid \langle \boldsymbol{x}, \boldsymbol{u} \rangle \geq -1 \; \forall \boldsymbol{x} \in P \right\}.$$

When P is reflexive, $P^\vee$ is also a lattice polytope, and if P is simplicial and reflexive, $P^\vee$ is smooth.

I first used the notion of allowing multiple upper and lower bounds in the pull-back and push-forward subdivision, as described here, with Haase to show that all smooth reflexive d-polytopes have regular unimodular triangulations for $d \leq 4$. I initially did this by hand computation using Batyrev's classification and descriptions of all 4-dimensional cases [7]. Thanks to Øbro [28], there are explicit representations (up to lattice equivalence) of all simplicial reflexive polytopes of dimension ≤ 8. There are 5, 18, 124, 866, 7622, 72256, and 749892 in dimensions 2, 3, 4, 5, 6, 7, and 8 respectively. Haase and Paffenholz used a computer implementation of the pull-back and

[7] Recall that the normal fan to a polytope $P \subset \mathbb{R}^d$ is the fan that partitions $\mathbb{R}^d$ into cones in bijection with the faces of P. The cones of this fan correspond to the cones normal to each face of P.

push-forward subdivisions (together with my visual inspection of some 2-dimensional projection's subdivisions) to confirm the existence of regular unimodular triangulations for all but 108 of the examples of dimension ≤ 8. These results are summarized in the following theorem.

Theorem 1.4.4. *[21]*

- *All smooth reflexive polytopes in dimensions $d \leq 6$ admit a flag regular unimodular triangulation.*
- *All but at most 3 (out of 72256) of the 7-dimensional smooth reflexive polytopes have a flag regular unimodular triangulation.*
- *All but at most 105 (out of 749892) of the 8-dimensional smooth reflexive polytopes have a regular unimodular triangulation.*

The remaining 108 cases may have such triangulations as well; however, this approach has thus far been unable to construct them. Some of these cases have been successfully projected to 3-dimensional subdivisions, which were unable to be projected further. There is no reason to believe that these projected 3-dimensional subdivisions may not admit nice triangulations themselves, but this has not been confirmed.

1.5 Unimodular Triangulations and Gröbner Bases

Armed with some background in triangulations and algebra, I turn our attention to how special triangulations of lattice polytopes can be used to study properties of algebraic objects, and why people care so much about finding these desirable triangluations.

To this end, I recall the setup of a toric ideal described earlier, with notational adjustments to reflect that which will be used in triangulation discussions. Given a field $\Bbbk$ and the homogenized set of lattice points in polytope P, $\mathcal{A} = (P \times \{1\}) \cap \mathbb{Z}^{d+1})$, consider the polynomial ring $S = \Bbbk[x_{\boldsymbol{a}} : \boldsymbol{a} \in \mathcal{A}]$ with one variable for each lattice point. There is a canonical ring homomorphism ϕ_P from S to the Laurent polynomial ring $\Bbbk[t_1^{\pm 1}, \ldots, t_d^{\pm 1}, t_{d+1}]$ which maps each variable to the corresponding, homogenized, $\boldsymbol{t}$-monomial: $\phi_P(x_{\boldsymbol{a}}) = \boldsymbol{t}^{\boldsymbol{a}} = t_1^{a_1} \cdot \ldots \cdot t_d^{a_d} \cdot t_{d+1}^{d+1}$. The toric ideal $\mathcal{I}_P = \ker \phi_P$ corresponding to this map encodes affine dependencies among the lattice points in P [44, Lemma 4.1]

$$\mathcal{I}_P = \left\langle \; \boldsymbol{x}^{\boldsymbol{m}} - \boldsymbol{x}^{\boldsymbol{n}} \; : \; \boldsymbol{m}, \boldsymbol{n} \in \mathbb{Z}^{\mathcal{A}}_{\geq 0} \; , \; \sum_{\boldsymbol{a} \in \mathcal{A}} m_{\boldsymbol{a}} \boldsymbol{a} = \sum_{\boldsymbol{a} \in \mathcal{A}} n_{\boldsymbol{a}} \boldsymbol{a} \; \right\rangle$$

where $\boldsymbol{x}^{\boldsymbol{n}} = \prod_{\boldsymbol{a} \in \mathcal{A}} \boldsymbol{x}_{\boldsymbol{a}}^{\boldsymbol{n}_{\boldsymbol{a}}}$.

A generic choice of weights $\boldsymbol{\omega} \in \mathbb{R}^{\mathcal{A}}$ induces a regular triangulation $\mathcal{T}_{\boldsymbol{\omega}}$ of P, as described earlier. Such an $\boldsymbol{\omega}$ also induces an ordering of the monomials in S: $\boldsymbol{x}^{\boldsymbol{m}} \prec \boldsymbol{x}^{\boldsymbol{n}} \iff \langle \boldsymbol{\omega}, \boldsymbol{m} \rangle < \langle \boldsymbol{\omega}, \boldsymbol{n} \rangle$. Given a polynomial $f \in S$, the leading term $\operatorname{in}_{\boldsymbol{\omega}} f$ is the monomial deemed greatest in this ordering with a non-zero coefficient in f, and the initial ideal $\operatorname{in}_{\boldsymbol{\omega}} \mathcal{I} = \langle \operatorname{in}_{\boldsymbol{\omega}} f \; : \; f \in \mathcal{I} \rangle$ of an ideal $\mathcal{I}$ is the collection of leading terms for all polynomials in $\mathcal{I}$. A set of polynomials in $\mathcal{I}$ whose leading terms generate this initial ideal, $\operatorname{in}_{\boldsymbol{\omega}} \mathcal{I}$, is called a Gröbner basis of $\mathcal{I}$ with respect to $\boldsymbol{\omega}$.

For a simplicial complex, such as $\mathcal{T}_{\boldsymbol{\omega}}$, the ideal generated by monomials corresponding to minimal non-faces of the complex is known as the Stanley-Reisner ideal (see example in Figure 1.5.1). The correspondence between $\operatorname{in}_{\boldsymbol{\omega}} \mathcal{I}_P$ and $\mathcal{T}_{\boldsymbol{\omega}}$ tells us whether this triangulation induced by $\boldsymbol{\omega}$ is unimodular.

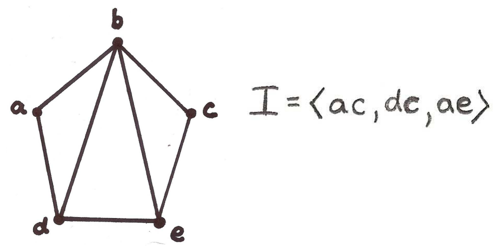

Figure 1.5.1: This image shows an example of a triangulation and its corresponding Stanley-Reisner ideal.

Lemma 1.5.1. *[21] Let $\mathcal{T}_{\boldsymbol{\omega}}$ be a regular unimodular triangulation $\mathcal{T}_{\boldsymbol{\omega}}$ of P. Then*

$$\mathrm{in}_{\boldsymbol{\omega}} \mathcal{I}_P = \left\langle \prod_{\boldsymbol{a}\in N} x_{\boldsymbol{a}} \ : N \text{ is a minimal non-face of } \mathcal{T}_{\boldsymbol{\omega}} \right\rangle.$$

Proof. Given a subset N of $\mathcal{A}$, consider the incidence vector of N, $\boldsymbol{n}(N) \in \{0,1\}^{\mathcal{A}}$. The resulting squarefree monomial $\prod_{\boldsymbol{a}\in N} x_{\boldsymbol{a}}$ will be abbreviated as $\boldsymbol{x}^{\boldsymbol{n}(N)}$.

For any non-face N, the sum $\boldsymbol{b}(N) = \sum_{\boldsymbol{a}\in N} \boldsymbol{a}$ is a lattice point in the cone σ_P (this is true whether or not N is a minimal non-face). Therefore, $\boldsymbol{b}$ belongs to exactly one cone over some face $F = F(N)$ of $\mathcal{T}_{\boldsymbol{\omega}}$ (this is true in general of lattice points in the cone σ_P). This means $\boldsymbol{b}$ can be written as a non-negative and integral linear combination $\sum_{\boldsymbol{a}\in F} m_{\boldsymbol{a}}(N)\boldsymbol{a} = \boldsymbol{b}(N)$ of the vertices of the unimodular simplex F. This yields $f_N = \boldsymbol{x}^{\boldsymbol{n}(N)} - \boldsymbol{x}^{\boldsymbol{m}(N)} \in \mathcal{I}_P$ with leading term $\mathrm{in}_{\boldsymbol{\omega}} f_N = \boldsymbol{x}^{\boldsymbol{n}(N)}$. This would fail if N was a face because in that case $\boldsymbol{b}(N)$ would lie in the cone over the face N defining it, and the construction would produce $f_N = \boldsymbol{x}^{\boldsymbol{n}(N)} - \boldsymbol{x}^{\boldsymbol{n}(N)}$, also known as zero.

Conversely, consider a binomial $f = \boldsymbol{x}^{\boldsymbol{n}} - \boldsymbol{x}^{\boldsymbol{m}} \in \mathcal{I}_P$, such that $\mathrm{in}_{\boldsymbol{\omega}} f = \boldsymbol{x}^{\boldsymbol{n}}$. I argue that the support of N, $\boldsymbol{n}$, is a non-face. By the definition of the regular triangulation $\mathcal{T}_{\boldsymbol{\omega}}$, the exponent vector $\boldsymbol{m}$ of a monomial $\prod_{\boldsymbol{a}\in F} \boldsymbol{x}_{\boldsymbol{a}}^{\boldsymbol{m}_{\boldsymbol{a}}}$ has the least $\boldsymbol{\omega}$-weight among all exponent vectors $\boldsymbol{n}$ with $\sum_{\boldsymbol{a}\in\mathcal{A}} m_{\boldsymbol{a}}\boldsymbol{a} = \sum_{\boldsymbol{a}\in\mathcal{A}} n_{\boldsymbol{a}}\boldsymbol{a}$. No such monomial, supported on a face F of $\mathcal{T}_{\boldsymbol{\omega}}$, can ever be the initial term of a binomial in $\mathfrak{I}_P$, so N contains a minimal non-face and $\boldsymbol{x}^{\boldsymbol{n}(N)}$ divides $\boldsymbol{x}^{\boldsymbol{n}}$. $\square$

The constructive nature of this lemma's proof demonstrates how $\mathcal{T}_{\boldsymbol{\omega}}$ can be recovered from $\operatorname{in}_{\boldsymbol{\omega}} \mathfrak{I}_P$ and conversely how $\operatorname{in}_{\boldsymbol{\omega}} \mathfrak{I}_P$ can be determined from $\mathcal{T}_{\boldsymbol{\omega}}$, as faces of $\mathcal{T}_{\boldsymbol{\omega}}$ correspond to the standard monomials of $\operatorname{in}_{\omega} \mathfrak{I}_P$. Further, a Gröbner basis can be obtained from the triangulation by considering the binomial $\boldsymbol{x}^{\boldsymbol{n}(N)} - \boldsymbol{x}^{\boldsymbol{m}(N)}$ for each for each minimal non-face N where $\boldsymbol{n}(N)$ and $\boldsymbol{m}(N)$ are as described in the proof.

However, when $\mathcal{T}_{\boldsymbol{\omega}}$ is not unimodular, the formula must be modified. In this case, the correspondence is not with the ideal itself, but the radical of the ideal. Given a ring S and ideal $\mathfrak{I} \in S$, $\operatorname{Rad}(\mathfrak{I}) = \{r \in S \mid r^n \in \mathfrak{I}$ for some positive $n \in \mathbb{Z}\,\}$. For the non-unimodular case, Lemma 1.5.1 must be modified to:

$$\operatorname{Rad}(\operatorname{in}_{\boldsymbol{\omega}} \mathfrak{I}_P) = \left\langle \prod_{\boldsymbol{a}\in N} \boldsymbol{x}_{\boldsymbol{a}} \;:\; N \text{ is a minimal non-face of } \mathcal{T}_{\boldsymbol{\omega}} \right\rangle.$$

While $\mathcal{T}_{\boldsymbol{\omega}}$ can still be recovered from $\operatorname{in}_{\boldsymbol{\omega}} \mathfrak{I}_P$, the reverse is not true.

Theorem 1.5.2. *[44] Given that $\mathcal{A}$ generates the lattice $\mathbb{Z}^{d+1}$, the initial ideal $\operatorname{in}_{\omega} \mathfrak{I}_P$ is squarefree if and only if the regular triangulation $\mathcal{T}_{\boldsymbol{\omega}}$ of P is unimodular.*

This appears as Corollary 8.9 in [44], and follows from [22, Thm. 5.3].

Proof. When $\mathcal{T}_{\omega}$ is unimodular, the $\operatorname{in}_{\boldsymbol{\omega}} \mathfrak{I}_P$ is squarefree by the last lemma (since a monomial whose exponents are given by an incidence vector is always squarefree).

Consider the case where $\mathcal{T}_{\boldsymbol{\omega}}$ is not unimodular. We know some face of the triangulation, $F = \operatorname{conv}(\boldsymbol{a}_0, \ldots, \boldsymbol{a}_d) \in \mathcal{T}_{\boldsymbol{\omega}}$, is a simplex of determinant $D > 1$. Let the sublattice of $\mathbb{Z}^{d+1}$ generated by the vertices of this face F be denoted by Λ. Note that for any $\boldsymbol{m} \in \mathbb{Z}^{d+1}$ we will have $D\boldsymbol{m} \in \Lambda$.

We construct a vector $\boldsymbol{b}$ as a non-negative integral linear combination of $\mathcal{A}$ such that $\boldsymbol{b} \in \operatorname{cone} F \cap \mathbb{Z}^{d+1} \setminus \Lambda$. Start by taking $\boldsymbol{b}' \in \mathbb{Z}^{d+1} \setminus \Lambda$. $\boldsymbol{b}'$ is an integral linear combination of $\mathcal{A}$, by construction. Add a multiple of $D \sum_{\boldsymbol{a} \in \mathcal{A}} \boldsymbol{a}$ large enough that the resulting coefficients are non-negative; then, add a big multiple of $\sum_{\boldsymbol{a} \in F} \boldsymbol{a}$ so the result is a point in $\operatorname{cone} F$.

Now, consider all $\boldsymbol{n} \in \mathbb{Z}^{\mathcal{A}}_{\geq 0}$ such that $\sum_{\boldsymbol{a} \in \mathcal{A}} n_{\boldsymbol{a}} \boldsymbol{a} = \boldsymbol{b}$ and pick the one whose $\boldsymbol{\omega}$-weight is minimal. We know $\boldsymbol{b} \notin \Lambda$, by construction, so $\boldsymbol{x}^{\boldsymbol{n}}$ is not supported on F, but $\boldsymbol{x}^{\boldsymbol{n}}$ is still never a leading term (i.e. $\boldsymbol{x}^{\boldsymbol{n}} \notin \operatorname{in}_{\boldsymbol{\omega}} \mathcal{I}_P$).

However, $D\boldsymbol{b} \in \operatorname{cone} F \cap \Lambda$, so there is some $\boldsymbol{m}$ such that $D\boldsymbol{b} = \sum_{i=0}^{d} m_i \boldsymbol{a}_i$ and $\boldsymbol{x}^{D\boldsymbol{n}} - \boldsymbol{x}^{\boldsymbol{m}} \in \mathcal{I}_P$. Since it is supported on the face F, $\boldsymbol{x}^{\boldsymbol{m}}$ can't be the leading term, which means we must have $(\boldsymbol{x}^{\boldsymbol{n}})^D \in \operatorname{in}_{\boldsymbol{\omega}} \mathcal{I}_P$. Clearly $(\boldsymbol{x}^{\boldsymbol{n}})^D$ can't be squarefree for $D \geq 1$. Therefore, $\operatorname{in}_{\boldsymbol{\omega}} \mathcal{I}_P$ is not squarefree. $\square$

This theorem yields a construction for regular unimodular triangulations, and shows how the regular unimodular triangulations in my research yield Gröbner bases of their corresponding toric ideals.

Lemma 1.5.1 also implies the correspondence between quadratic triangulations and quadratic Gröbner bases that I referenced earlier. (The version presented here combines Corollaries 8.4 and 8.9 in [44].)

Theorem 1.5.3. *[21] If P has a quadratic triangulation $\mathcal{T}$, then the defining ideal $\mathcal{I}_P$ of the projective toric variety $X_P \subset \mathbb{P}^{r-1}$ has a quadratic Gröbner basis.*

In which case, the corresponding initial ideal is :

$\mathrm{in}(\mathcal{I}_P) = \langle x_{\boldsymbol{a}} x_{\boldsymbol{b}} \mid \boldsymbol{ab} \text{ is not an edge in } \mathcal{T} \rangle$.

This is a powerful tool. For example, combined with Theorem 1.4.4, it guarantees that all toric fano d-folds of dimension ≤ 6 have quadratic triangulations.

1.5.1 Covering Properties

The results seen in Section 1.5 are not the only ties between combinatorial properties of polytopes and properties of corresponding algebras. There are many links appearing in hierarchies connecting properties of convex geometry and algebra [21].

The combinatorial properties include, in decreasing strength:

1. $P \cap \mathbb{Z}^d$ is totally unimodular
2. P is compressed
3. P has a regular unimodular triangulation
4. P has a unimodular triangulation
5. P has a unimodular binary cover (a $\mathbb{Z}_2$ cycle generating $H_d(P, \partial P; \mathbb{Z}_2)$ formed by unimodular simplices)
6. P has a unimodular cover

7. C has a free Hilbert cover (every lattice point is a $\mathbb{Z}_{\geq 0}$-linear combination of linearly independent lattice points in $P \times \{1\}$)

8. C has the integral Carathéodory property (every lattice point is a $\mathbb{Z}_{\geq 0}$-linear combination of $\dim C$ many lattice points in $P \times \{1\}$)

9. P is integrally closed

Most of these properties translate directly into an algebraic language expressing facts about R_P and $\mathcal{I}_P$, but there are other properties of interest which do not fit as neatly into the larger hierarchy.

(1') P has a quadratic triangulation

(2') $\mathcal{I}_P$ has a quadratic Gröbner basis

(3') R_P is a Koszul algebra ($\Bbbk$ has a linear free resolution as an R_P-module)

(4') $\mathcal{I}_P$ is generated by quadrics

The two hierarchies are related by the fact that a quadratic triangulation is, in fact, a regular unimodular triangulation. Both hierarchies appear in the forthcoming survey of triangulation results: [21].

Though not all of these properties appear directly in my work, their relative position and relevance in the hierarchy make them worth noting. The quadratic triangulation defined earlier has the nicest properties. The desirability of such a triangulation combined with the fact that understanding them requires unimodularity, regularity, and triangulations, all of which appear in the hierarchy, make them a worthy destination in the first string of definitions I gave. However, that doesn't mean the covering properties should be ignored. When a quadratic triangulation does not exist, the search for nice covers becomes a next target.

1.5.2 Covering Background

Covering properties are typically discussed in terms of cones. A *cone* is the space of all positive linear combinations of its set of generating vectors $\{\boldsymbol{a}_1, \boldsymbol{a}_2, ...\boldsymbol{a}_m\}$; $C = cone\ \{\boldsymbol{a}_1, \boldsymbol{a}_2, ...\boldsymbol{a}_m\} = \{\lambda_1\boldsymbol{a}_1 + \lambda_2\boldsymbol{a}_2 + ... + \lambda_m\boldsymbol{a}_m \in \mathbb{R}^n : \lambda_i \in \mathbb{R}, \lambda_i \geq 0$ for $i = 1, 2, ...m\}$. Here I am interested in the cone over the polytope being investigated. When discussing a d-dimensional polytope the "cone over P", σ_P is the cone whose intersection with its height one plane is the polytope P. Taking $\boldsymbol{a}_1, \boldsymbol{a}_2, ...\boldsymbol{a}_m \in \mathbb{R}$ such that $\{0\}$ is the largest linear subspace of C ensures that C will be *rational*, *pointed*, and *polyhedral*. This is clearly the case in the cone over the polytopes discussed in this work.

A *Hilbert basis* of C is a finite set of integral vectors $\boldsymbol{h}_1, \boldsymbol{h}_2, ...\boldsymbol{h}_k$ such that each integral vector of C can be expressed as a non-negative integral combination of $\{\ \boldsymbol{h}_1, \boldsymbol{h}_2, ...\boldsymbol{h}_k\}$ meaning:

$C \cap \mathbb{Z} = \{\ \lambda_1\boldsymbol{h}_1 + \lambda_2\boldsymbol{h}_2 + ... + \lambda_k\boldsymbol{h}_k \in \mathbb{R}^n :\ \lambda_i \in \mathbb{Z}, \lambda_i \geq 0$ for $i = 1, 2, ...k\}$. A *minimal Hilbert basis* is the finite set of integral vectors in C that cannot be expressed as a sum of 2 other integral vectors in C. The fact that every cone has a Hilbert basis is a result dating to Grodan (1873), but the fact that the minimal Hilbert basis for every pointed cone was unique was shown by van der Corput (1931). This is further discussed in Schrijver's *Theory of Linear and Integer Programing* [39].

Covers of a cone C refer to properties of a collection $\boldsymbol{C} = \{C_1, C_2, ...C_s\}$ of subcones of C. Such a collection $\boldsymbol{C}$ is a *cover* of C if every point of C is contained in some subcone C_i. That is, if the union of the cones of $\boldsymbol{C}$ is C. The difference between a triangulation and a cover is that in a cover each point may be contained in more than one cell. In a cover, one can consider *interior facets* and *generic points* of the cover. A facet F of a subcone C_i is said to be an interior facet if F is not contained in the boundary of C. The set of interior facets of the cones in $\boldsymbol{C}$ is denoted $\boldsymbol{F} = \{F_1, F_2, ...F_r\}$. A point $g_0 \in int(C)$ is said to be generic with respect to $\boldsymbol{C}$ if it is not contained in the boundary of any of the cones of $\boldsymbol{C}$. This is to say it is not on the boundary of C nor is it contained in any interior facet of a cone in $\boldsymbol{C}$. Triangulations correspond to covers in which each generic point is contained in exactly 1 of the C_i and each interior facet F_i is a facet of exactly 2 of the C_i. With this terminology, $\boldsymbol{C}$ is said to be a *binary cover* of C if every generic point $g_0 \in C$ is contained in an odd number of the subcones C_i and every interior facet F_j is a facet of an even number of the C_i. In this context unimodularity refers to the unimodularity of each of the cones of $\boldsymbol{C}$. Here, rather than looking for a proper triangulation of a polytope, the search is for a collection of unimodular simplices contained in the span of the polytope satisfying the desired overlap properties. Triangulations are also binary covers. While points can still be in multiple cones, binary covers are nicer than general covers in that their cells match up along interior (and exterior) facet boundaries.

There are infinite classes of cones of dimension ≥ 4 that do not have unimodular triangulations, but do have binary covers [17]. Unfortunately, most explicitly known examples are too involved for me to provide a simple visual here. [8] However, an example of a binary cover that does not correspond to a triangulation can be obtained by considering a 3-sheeted cover where each sheet is itself a triangulation. In this case each generic and boundary point will be contained in $3 * 1$ cones, and each interior facet point will be contained in $3 * 2$ cones.

[8] For example, the 4-dimensional simplicial cone $c[14, 31, 34, 39]$ has a binary cover, but does not have a unimodular cover [17]. Its binary cover contains 161 facet unimodular cones.

Generally, the lower levels of the hierarchy are only investigated for a polytope or class of polytopes when the more desirable properties can not be confirmed.

1.6 Polytopal Dilations, The KMW Theorem, and $2S(p,q)$

One class of polytopes for which there is interest in binary covers is the class of dilated simplices, $2S(p,q)$. In general, interest in dilations of polytopes ties back to an early result on unimodular triangulations by Knudsen, Mumford, and Waterman [23]. It may not be surprising that, given a polytope P, dilations cP of P inherit desirable properties regarding nice triangulations. However, it is actually true that all polytopes P, even those who fail to have nice triangulations themselves, have dilations with nice triangulations. The statement of their theorem itself is important, but it also raised many interesting questions and sparked further research, including the discussion of $2S(p,q)$ seen here.

1.6.1 Polytopal Dilations

Before discussing how to triangulate a dilated polytope, it is useful to know what happens when a triangulation is dilated. Given a polytope P with a (regular) unimodular triangulation $\mathcal{T}$, dilating each simplex of $\mathcal{T}$ results in a triangulation of cP made up of dilations of unimodular simplices. The dilated standard simplex is $c\Delta^d = \{\boldsymbol{x} \in \mathbb{R}^{d+1}_{\geq 0} \ : \ \sum x_i = c\}$. Slicing this simplex along hyperplanes parallel to its facets (i.e. $x_i = k$ for $i = 1, ..., d+1$ and $k = 1, .., c-1$) divides it into hypersimplices. This is a regular subdivision, and no desirable properties are lost by mapping this subdivision of the standard dilated simplex to a dilation of an arbitrary simplex. Further, the natural nature of this subdivision guarantees that the simplices line up properly on their boundaries. As a result, this canonical subdivision of $c\mathcal{T}$ yields a triangulation of cP which will be regular if $\mathcal{T}$ is.

Theorem 1.6.1. *[21]*

If a lattice polytope P has a (regular) unimodular triangulation $\mathcal{T}$, then for every positive integer c, cP also has a regular unimodular triangulation.

Proof. All cells (the hypersimplices) of the described canonical subdivision of $c\mathcal{T}$ are facet width one with respect to all faces, by construction. Therefore, any pulling refinements of its vertices, yields a (regular) unimodular triangulation.

$\square$

In fact, conditions on non-faces can also be preserved under dilation.

Theorem 1.6.2. *[21]*

If a lattice polytope P has a quadratic triangulation, then so does its dilation cP for every positive integer c.

Constructing these triangulations requires taking a little more care choosing weights which will induce a triangulation with the desired properties, and the overall proof is a bit more involved. It has been omitted as it draws on results regarding classes of polytopes that have not been discussed here.

These theorems guarantee that dilations of polytopes with nice triangulations will have nice triangulations, but in fact it is also true that for large enough c, cP will have nice triangulations even when the original polytope P does not.

1.6.2 The KMW Theorem

The fact that sufficiently large dilations of any polytope have nice triangulations is the content of Knudsen, Mumford, and Waterman's theorem.

Theorem 1.6.3. *[23]*

For every polytope P, there is a factor $c = c(P) \in \mathbb{Z}_{>0}$ such that the dilation cP admits a regular unimodular triangulation.

$c(P)$ is called a KMW-number of P.

The proof of this theorem depends on Knudsen-Mumford triangulations of d-space, which are not discussed here. It is based on induction on the "size" of the polytope with respect to the lattice. Clearly if the polytope is itself a unimodular simplex there is no work to be done. The work of the proof is broken into two cases. The first case is where P's size with respect to the lattice is composite, the second is where it is prime. The full proof, along with necessary background on unimodular triangulations of $\mathbb{R}^d$ can be found in [10].

This theorem has been a jumping-off point for much research to answer the many questions it inspires, including:

- What is the minimum $c(P)$ for a given polytope P?
- Is there a $c(P)$ such that it is a KMW-number for every polytope of dimension d?
- What is the structure of the set of KMW-numbers of a given P — in particular if it is a monoid? [9]

[9] Theorem 1.6.1 makes it clear that the set is closed under multiplication, but it is not clear whether it is closed under addition. No polytope and integer pair P, c are known for which c is a KMW-number for P but $c+1$ is not. It is not clear that this is not possible.

The search for answers to these questions is an active area of research. This work often focuses on specific classes of polytopes whose structure makes them particularly manageable and/or useful. It is known that every line segment and lattice polygon has a regular unimodular triangulation. In the case where $dim(P) = 1$ or 2, P has a KMW-number $c(P) = 1$, but other dimensions are not so straightforward.

The 3-dimensional case is relatively well understood in comparison to higher dimensions, but there are still open questions. In particular, the following has been established (for more on the proofs of these results see [21]):

- There are 3-dimensional polytopes for which 1 and 2 are not KMW-numbers.
- $c = 4$ and $c \geq 6$ are KMW-numbers of every 3-dimensional polytope.

Whether 3 and 5 are KMW-numbers for all 3-dimensional polytopes remains an open question. Proofs in this area are based largely on finding nice, compatible triangulations of dilated simplices, so the search for triangulations of dilated simplices is of particular interest. I have examined the case of $2S(p, q)$.

1.6.3 $2S(p, q)$

White's Theorem [38, 46, 27] says that any empty tetrahedron in $\mathbb{R}^3$ is width one with respect to some (not necessarily a facet normal) lattice direction, meaning it fits between some pair of adjacent hyperplanes. [10] This means every empty 3-simplex is unimodularly equivalent to

$$S(p, q) = \text{conv}\begin{bmatrix} 0 & 0 & 1 & p \\ 0 & 0 & 0 & q \\ 0 & 1 & 0 & 1 \end{bmatrix}$$

for some pair of integers $0 \leq p \leq q$ with $gcd(p, q) = 1$. Since unimodularity is of such interest, note that the volume of $S(p, q)$ is q.

$S(p, q)$ is an empty simplex, as shown in Figure 1.6.1, but dilations $cS(p, q)$ clearly aren't. However, we know exactly where $cS(p, q)$'s non-vertex lattice points are and can determine how many non-vertex lattice points a given $cS(p, q)$ contains. The lattice points of $cS(p, q)$ lie on the bottom edge, top edge, and on the $c - 1$ horizontal lattice planes between the top and bottom edges. In particular, the intersection of each of the intermediate planes with $cS(p, q)$ form parallelograms that can be tiled by translates of the intersection of $2S(p, q)$ with the plane $z = 1$, $C(p, q) = 2S(p, q) \cap \{z = 1\}$ =conv$\{(0, 0, 1), (1, 0, 1), (p, q, 1), (p + 1, q, 1)\}$. Understanding the lattice points of $C(p, q)$ is sufficient for understanding lattice points of arbitrary $cS(p, q)$. Further, since the case of $2S(p, q)$ is the focus of my work, understanding $C(p, q)$ is particularly worthwhile.

$C(p, q)$ has $q + 3$ lattice points, 4 vertices and exactly one in each of the $c - 1$ lattice lines intersecting $C(p, q)$ in the direction $(1, 0, 0)$, and in each of the $c - 1$ lattice lines in the direction $(0, p, q)$. These labelings can be used to create distinct orderings of these $q - 1$ interior lattice points, referred to

[10] More on width can be found in Section 1.3.1.

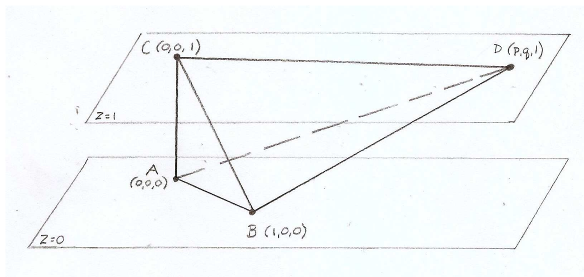

Figure 1.6.1: This shows $S(p,q)$, the empty tetrahedron sandwiched between the $z = 1$ and $z = 0$ planes with vertices $(0,0,0)$, $(1,0,0)$, $(0,0,1)$, and $(p,q,1)$.

as the Y-order and the X-order [21].

Given that these simplices are 3-dimensional, it is known that $c = 4$ and $c \geq 6$ are KMW numbers for $S(p,q)$. It is clear that these simplices are not unimodular for $q > 1$, so the cases of $c = 2, 3, 5$ are of interest. Something is known for the $c = 2$ case.

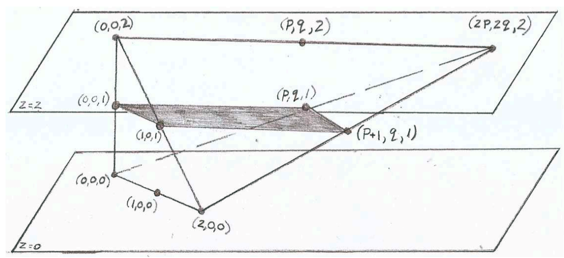

Figure 1.6.2: This figure shows $2S(p,q)$. Its 4 vertices, $(0,0,0)$, $(2,0,0)$,$(0,0,2)$, and $(2p,2q,2)$, as well as the 4 other lattice points on its boundary $(1,0,0)$, $(0,0,1)$,$(p-1,q,1)$, and $(p,q,1)$, are labeled. Its intersection with the $z=1$ plane is indicated by the shaded parallelogram.

Theorem 1.6.4. *[21]*

$2S(p,q)$ has a unimodular triangulation if and only if $p = \pm 1$ mod q; further, this unimodular triangulation is always regular (and can be chosen to have standard boundary).

The standard boundary comment is not particularly interesting for the triangulation of $2S(p,q)$ itself or where it falls on the corresponding hierarchy of algebraic properties. However, this is very useful in the search for nice triangulations of other 3-dimensional polytopes with regular subdivisions whose cells are of the form $2S(p,q)$. The existence of regular unimodular triangulations with standard boundary for each cell guarantees that such a subdivision can be refined to a regular unimodular triangulation, and that each cell can be triangulated in a manner that will match consistently on their boundaries.

However, there is also value in knowing more about the properties of $2S(p,q)$ itself. What happens in the case where there is no unimodular triangulation? I considered this case, implementing a computer search for binary covers of $2S(p,q)$ for examples where $gcd(p,q) = 1$ and $p \neq \pm 1 \mod q$. As described in 1.5.1, the existence of a binary cover, while not as powerful a property as regular unimodular triangulations, is itself desirable, and computer programs can be written to search for them. An algorithm for this is given in [18]. However, my search used a different method.

The code I have was designed to offer a yes or no answer to the question of whether a given (p,q) pair's $2S(p,q)$ has a binary cover. However, it could be altered to actually produce a sample cover when the answer is yes. Since the question is only of interest when $gcd(p,q) = 1$ and $p \neq \pm 1 \mod q$, the code first checks these properties on each (p,q) pair and moves on to the next pairing when the conditions are not satisfied. Once a suitable (p_0,q_0) pair has been determined, the code finds all lattice points in the corresponding $2S(p_0,q_0)$. It then checks all 4-tuples of these points to construct a list of all full dimensional unimodular simplices contained in $2S(p_0,q_0)$. The faces of these simplices are checked for containment in the boundary of $2S(p_0,q_0)$. A system of linear equations over $\mathbb{Z}_2$ is then formed. The equations indicate the incidence relations between triangles and the unimodular siplicies of which they are facets. This provides a context for asking whether there is a collection of the simplices (i.e. a cover) such that each boundary facet occurs an odd number of times, and each interior facet appears an even number (possibly zero) number of times. If the answer to that question is yes, and each point of $2S(p_0, q_0)$ is covered, then $2S(p_0, q_0)$ has a unimodular cover. In this case, the code prints "(p_0, q_0), yes," and moves on to the next (p, q) pair. If the answer is no, the code prints, "(p_0, q_0), no," and stops.

After running for 2 days, all (p, q) pairs through $(3, 49)$, in lexicographical order, had been checked and confirmed to have binary covers. I conjecture that all $2S(p, q)$ have binary covers for (p, q) pairings satisfying this *gcd* condition. Producing a general proof of this is among my current projects.

Chapter 2

Tropical Hyperplanes, Meet Matroids

2.1 Introduction

Tropical oriented matroids were introduced by Federico Ardila and Mike Develin. Their paper aims to establish ties between tropical hyperplane arrangements and oriented matroids, much like those between standard hyperplane arrangements and traditional matroids. There is even discussion of the notion of realizable tropical oriented matroids. Their seminal paper on the subject addresses many properties seen in the traditional setting discussing their tropical analogues. Their conjectures about tropical hyperplane arrangements and tropical oriented matroids themselves, together with ties to products of simplices, make it clear that there is more interesting work to be done with the ideas they introduced. My research has looked at properties of such objects as well as closely related ones. This includes what distinguishes between generic and non-generic cases, as these differences are not discussed in Ardila and Develin's work. I also address other ideas that are of interest moving forward, such as understanding and modeling the structure of the space of all tropical oriented matroids.

2.2 Tropical Hyperplanes Background

2.2.1 Tropical Basics

Working with tropical hyperplane arrangements requires some familiarity with tropical mathematics. (For a complete introductions to the field see [36], [41], and [26]). Tropical geometry is done over the *tropical semiring.* The *tropical semiring* is the ring $R = \{\mathbb{R}, \otimes, \oplus\}$, where $\otimes$ is standard addition and $\oplus$ is taking maximums. To avoid confusion, I will use $\otimes$ and $\oplus$ to represent tropical multiplication and addition, allowing traditional multiplication and addition notation to be used without ambiguity. [1]

$$\text{ex.}$$
$$2 \otimes 3 = 2 + 3 = 5$$
$$2 \oplus 3 = \max\{2, 3\} = 3$$

Alternatively, this ring can also be realized as the image of traditional arithmetic in a power series ring under a degree map sending a power series in $t^{(-1)}$ to its leading exponent. This makes the definition of $\otimes$ and $\oplus$ seem more natural, as the operations can be thought of as the logs of traditional multiplication and addition. The resulting tropical geometry is like a piece-wise linear version of algebraic geometry. Here, geometric questions about

[1] The convention of taking addition to be the max operation rather than the min operation is not universal. However, the resulting rings are equivalent, and the convention I use here matches that of Ardila and Develin's paper.

algebraic varieties can be translated into polyhedral questions about polyhedral fans. The niceness of this perspective has been a driving force behind the recent growth in the field, which has found applications in many areas, including commutative algebra, topology, and phylogenetics. References for these include [4], [13], [26], [36], [41], [12], [2], [14], [40], [32], [25], and [47] .

Given its semiring structure, tropical arithmetic has many of the properties of traditional arithmetic. Both addition and multiplication are commutative,

$$\text{ex.}$$
$$2 \oplus 3 = \max\{2,3\} = \max\{3,2\} = 3 \oplus 2$$
$$2 \otimes 3 = 2 + 3 = 3 + 2 = 3 \otimes 2$$

and the distributive law holds for multiplication over addition.

$$\text{ex.}$$
$$5 \otimes (2 \oplus 3) = 5 + \max\{2,3\} = 8$$
$$(5 \otimes 2) \oplus (5 \otimes 3) = \max\{(5+2),(5+3)\} = 8$$

What is different is that the additive identity is $-\infty$ rather than 0, and there are no additive inverses. This means that there is no notion of subtraction. However, there are multiplicative inverses. The tropical multiplicative identity is 0. This means, for example, that every row of Pascal's triangle is all 0's and the coefficients of each term of all binomial expansions are 0's. As a result, *Freshmen's Dream* holds for all powers. I demonstrate this with the cubic example.

$$(x \oplus y)^3 = (x \oplus y) \otimes (x \oplus y) \otimes (x \oplus y)$$
$$= 0 \otimes x^3 \oplus 0 \otimes x^2 y \oplus 0 \otimes xy^2 \oplus 0 \otimes y^3$$

Since 0 is the multiplicative identity, the coefficients can be ignored, leaving

$$(x \oplus y)^3 = x^3 \oplus x^2 y \oplus xy^2 \oplus y^3$$
$$3(\max\{x, y\}) = \max\ \{3x, 2x + y, x + 2y, 3y\} = \max\{3x, 3y\}$$

and so,

$$(x \oplus y)^3 = x^3 \oplus y^3$$

Monomials and polynomials in the tropical semiring are defined analogously to the traditional setting; however, they behave differently. Given a set of variables $x_1, x_2, ..., x_n$ representing elements of the tropical semiring, a *tropical monomial* is defined to be any product of these variables (allowing positive and negative integer exponents). Each such monomial represents a function from $\mathbb{R}^n$ to $\mathbb{R}$. When evaluated, these monomials behave like linear functionals in traditional arithmetic.

ex.

$$x_1 \otimes x_1 \otimes x_2 \otimes x_3 \otimes x_3 \otimes x_3 \otimes x_3 = x_1^2 \otimes x_2 \otimes x_3^4$$
$$x_1^2 \otimes x_2 \otimes x_3^4 = 2x_1 + x_2 + 4x_3$$

A *tropical polynomial* is then defined as a finite linear combination of tropical monomials, meaning they are of the form

$$p(x_1, x_2, ..., x_n) = \oplus \sum a_1 \otimes x_1^{i_{11}} \otimes x_2^{i_{12}} ... \otimes x_n^{i_{1n}}$$

where each of the a_j are real numbers and each of the i_{jk} are integers. These tropical polynomials also represent functions from $\mathbb{R}^n$ to $\mathbb{R}$. Such polynomials p are evaluated at a point $\boldsymbol{x} \in \mathbb{R}^n$ by taking the maximum over the values of each of the monomials of p evaluated at $\boldsymbol{x}$. That is to say

$$p(x_1, x_2, ..., x_n) = \oplus \sum a_1 \otimes x_1^{i_{11}} \otimes x_2^{i_{12}} ... \otimes x_n^{i_{1n}}$$
$$= max_1^n \, \{(a_j + i_{j1}x_1 + ... + i_{jn}x_n)\}$$

This shows that tropical polynomials are continuous piecewise-linear functions (with a finite number of pieces). Figure 2.2.1 shows the graph of the cubic polynomial $p(x) = 3 \otimes x^3 \oplus 4 \otimes x^2 \oplus 7 \otimes x \oplus 11$. The graph of $p(x)$ is segments of the graphs of the 4 lines: $y = 3x + 3$, $y = 2x + 4$, $y = x + 7$, and $y = 11$. Note that for a general cubic, $p(x) = a \otimes x^3 \oplus b \otimes x^2 \oplus c \otimes x \oplus d$ only when a, b, c, and d satisfy

$$d - c \leq c - b \leq b - a$$

do all four lines, $y = 3x + a$, $y = 2x + b$, $y = x + c$ and $y = d$, actually contribute to the graph.

I explain this to demonstrate how the tropical world differs from standard geometry. However, in the tropical setting, evaluating polynomials at points and graphing their values is not nearly as useful as understanding their roots. Roots of tropical polynomials are not zero sets. In fact, considering any tropical polynomial with a positive constant term, one sees the zero set for many tropical polynomials is empty. For example, the tropical polynomial $x \oplus 3 = 0$ does not have a solution since $max\{x, 3\} \geq 3$ for all x.

Armed with the definition of tropical polynomials, roots of a tropical polynomial and tropical curves can be defined. As described above, a tropical polynomial $p(x) \colon \mathbb{R}^n \to \mathbb{R}$ is the maximum of a finite set of linear functionals. Given such a $p(x)$, the *roots* of $p(\boldsymbol{x})$, making up the set of points in the *hypersurface* $\mathcal{H}(p)$, are defined as all points $\boldsymbol{x} \in \mathbb{R}^n$ for which the maximum is achieved at least twice by the monomials of p. [2] This definition of vanishing

[2]This is also called the nonlinearity locus, as it is the set of points where the function does not behave linearly – the set of points where distinct linear components meet.

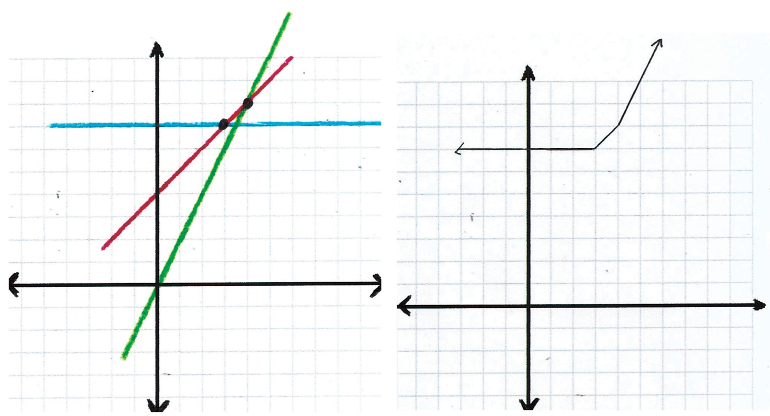

Figure 2.2.1: The image on the left shows the 3 lines contributing to the piecewise linear function given by the tropical quadratic polynomial $p(x) = 0 \otimes x^{\oplus} 4 \otimes x \oplus 7$. The figure on the right is the graph of $p(x)$'s values given by the maximum of the 3 lines for each x value. Note that the roots of $p(x)$ are the values for which this graph is non-linear, $\{3, 4\}$. These are the points, indicated in black, where there is a "tie" for the maximal value of the lines in the image on the left.

is more natural when thinking of tropical operations as the logs of ordinary operations. From this perspective, roots are the points where the leading terms cancel. This is to say that $\mathcal{H}(p)$ is the set of points $\boldsymbol{x} \in \mathbb{R}^d$ such that p is not linear at $\boldsymbol{x}$; informally, this is where the linear functionals defining p meet. For the abstract cubic described earlier this would be the set $\mathcal{H}(p) = \{b-a, c-b, d-c\}$, and for the specific example $p(x)$ depicted in Figure 2.2.1 $\mathcal{H}(p) = \{1, 3, 4\}$. For polynomials in two variables the *tropical curve* $\mathcal{H}(p)$ is a finite piecewise-linear graph with both bounded and unbounded edges. For example, the graph of the general tropical line

$$p(x, y) = a \otimes x \oplus b \otimes y \oplus c,$$

with $a, b, c \in \mathbb{R}$, is the curve $\mathcal{H}(p)$ consisting of all (x, y) for which

$$p\colon (x, y) \to \max\,(a + x, b + y, c)$$

is not linear. It is made up of the 3 rays emanating left, down, and diagonally up in the $x = y$ directions from the point $(x, y) = (c - a, c - b)$. Figure 2.2.1 depicts an example, the tropical line $p(x, y) = 3 \otimes x \oplus 4 \otimes y \oplus 1$.

Tropical lines are 2-dimensional tropical linear spaces. In general, a *tropical hyperplane* is a subset $\mathcal{H}(\ell) \subset \mathbb{R}^n$ such that l is a tropical linear form in n unknowns,

$$\ell(\boldsymbol{x}) = a_1 \otimes x_1 \oplus a_2 \otimes x_2 \oplus ... \oplus a_n \otimes x_n.$$

These tropical hyperplanes are the focus of my tropical work.

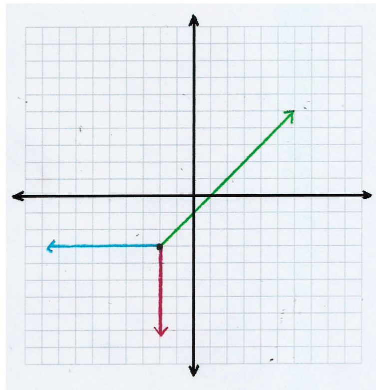

Figure 2.2.2: This is the graph of the tropical line $p(x, y) = 3 \otimes x \oplus 4 \otimes y \oplus 1$. Its apex is located at $(1 - 3, 1 - 4) = (-2, -3)$. The rays of the curve are labeled according to which of the terms they simultaneously maximize. Green depicts where $3 + x = 4 + y > 1$, blue where $4 + y = 1 > 3 + x$ and red where $3 + x = 1 > 4 + y$.

2.2.2 Tropical Hyperplane Arrangements

Tropical d-space $\mathbb{R}^d$ is formed much like ordinary Euclidean geometry. However, it is interpreted differently. Here vector addition is taking coordinate-wise maximums, and scalar multiplication is now realized by adding a constant to each vector. Often in the study of tropical hyperplane arrangements, one benefits from considering *tropical projective* $(d-1)$*-space*, $\mathbb{TP}^{d-1}$. This is obtained from tropical $\mathbb{R}^d$ by modding out by tropical scalar multiplication. The result is the traditional vector quotient space $\mathbb{R}^d/(1, ..., 1)\mathbb{R}$, and can be visualized as real $(d - 1)$-space. This is where the tropical hyperplanes discussed here live. Notice that without projectivizing, $(1, ..., 1)$ would be in the lineality space of every tropical hyperplane.

Like traditional hyperplanes, a tropical hyperplane is determined by the roots of a linear functional, $\ell = \oplus \sum c_i \otimes x_i$. However, in tropical geometry this linear functional is evaluated by taking the maximum over the set $\{c_1 + x_1, ..., c_d + x_i\}$, and a function vanishes when this maximum, whatever that may be, is achieved at least twice on the set. The example of evaluating a linear tropical polynomial at the point $(0, 0, 0)$ seems natural.

ex.

$$\ell(x, y, z) = 4x \oplus y \oplus 3z$$
$$\downarrow$$
$$\ell(0, 0, 0) = max\{4 + 0, 1 + 0, 3 + 0\} = 4$$

However, in the projective tropical setting the point $(0, 0, 0)$ is no more interesting than $(7, 7, 7)$ and since 4 is a unique maximum value, it is not a root. The point $(1, 4, 0)$ is of interest in this example, as here the maximum value of 5 is achieved twice, making it a root of $f(x, y, z) = 4x \oplus y \oplus 3z$.

ex.

$$\ell(x, y, z) = 4x \oplus y \oplus 3z$$
$$\downarrow$$
$$\ell(1, 4, 0) = max\{4 + 1, 1 + 4, 3 + 0\} = 5$$

The graph of a tropical hyperplane in $\mathbb{R}^d$ is a union of $(d+1)$ traditional hyperplanes truncated by their intersections. That is, they form a fan polar to the standard simplex (the simplex formed by the convex hull of $\{\boldsymbol{e}_1, ...\boldsymbol{e}_d\}$). For $\ell(\boldsymbol{x}) = \sum c_i x_i$, the apex of this fan is located at $(-c_1, ... - c_d)$. Two such fans are shown in Figure 2.2.3. Notice that while the coefficients of $f(x, y, z)$ are $(1, 1, 1)$ its apex appears to be the origin. This is a result of the projectivization, which identifies all points that vary by a multiple of $(1, 1, 1)$. Convention is to represent the point $(x_1, ...x_d)$, by its $x_1 = 0$ equivalent, $(0, x_2 - x_1, ..., x_d - x_1)$ in $\mathbb{R}^{d-1}$, as done in 2.2.3.

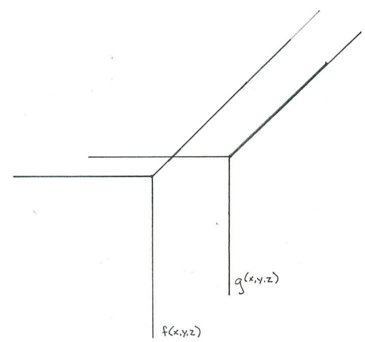

Figure 2.2.3: This shows the graphs of two tropical hyperplanes, $g(x, y, z) = x \oplus y \oplus z$ and $f(x, y, z) = 4x \oplus y \oplus 3z$, in $\mathbb{TP}^{3-1} = \mathbb{TP}^2$.

Unlike traditional hyperplanes whose half-spaces can be labeled by sign, positive or negative, a tropical hyperplane divides $\mathbb{TP}^{d-1}$ into d full dimensional sectors. However, there is a natural labeling of both the full dimensional sectors and their lower dimensional intersections that form the fan that is the tropical hyperplane. The cones of the fan are indexed by the subset of $[d]$ corresponding to the $c_i + x_i$ terms in the defining linear form $\ell(\boldsymbol{x}) = \sum c_i \otimes x_i$ that it maximizes. For the full dimensional sectors this is a single element, as there the maximum is only achieved once. For each $d - k$ dimensional cone this is a k element set, with the apex labeled by the entire set $[d]$. Equivalently, one can think of indexing each cone by the basis vectors corresponding to the face of the standard simplex to which it is polar.

Positions within a tropical hyperplane arrangement are given by ordered tuples of these $[d]$-subsets, just as tuples of $\{+, -, 0\}$ can be used to describe positions in traditional hyperplane arrangements. Some examples are shown in 2.2.6. This similarity between traditional and tropical hyperplane arrangement notation as subset collections foreshadows notational and descriptive parallels that will appear in the discussion of tropical matroid representation.

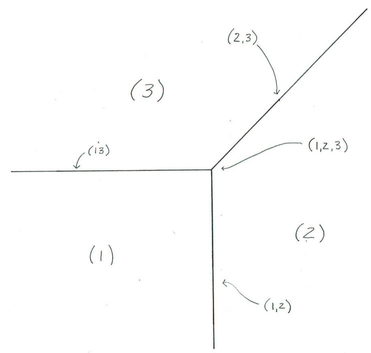

Figure 2.2.4: This shows a tropical hyperplane in $\mathbb{TP}^2$ whose cones have been labeled by 1-tuples corresponding to the faces of the standard tetrahedra to which each is dual.

Just as traditional polytopes can be given as the convex hull of their vertex sets or by a collection of defining hyperplanes [48], tropical polytopes can also be described by either a tropical convex hull of a finite point set, or a union of regions bounded by a set of tropical hyperplanes. However, this requires appropriate definitions of a tropical convex hull and a tropical polytope, as well as understanding how to choose the set of defining tropical hyperplanes.

The *tropical convex hull* of a point set $V = \{\boldsymbol{v}_1, ...\boldsymbol{v}_n\}$ in $\mathbb{TP}^{d-1}$ is the set of all tropical linear combinations $\oplus \sum (c_i \otimes \boldsymbol{v}_1)$ such that $c_i \in \mathbb{R}$, where the scalar multiplication $(c_i \otimes \boldsymbol{v}_i)$ is defined tropically component-wise. Notice that unlike standard polytopes, where the coefficients defining a convex hull of points are taken to be non-negative and summing to one, in the tropical version *all* real linear combinations are allowed. (For more on tropical convexity and tropical polytopes see [13].)

A *tropical polytope* can now be defined as the tropical convex hull of a finite point set. These tropical polytopes are, as one would hope, bounded

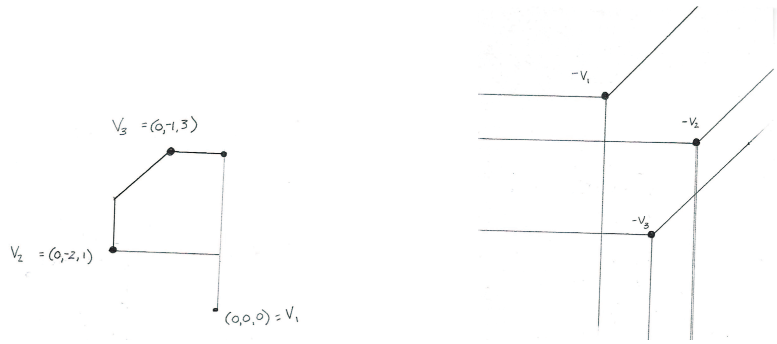

Figure 2.2.5: Here are images of combinatorially equivalent tropical polytopes (in particular triangles). The first is defined as the convex hull of a finite point set $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3\}$. The second is defined as the bounded polyhedral complex obtained by placing a tropical hyperplane at each of $-\boldsymbol{v}_1, -\boldsymbol{v}_2, and - \boldsymbol{v}_3$.

polyhedral complexes, and are explored further along with other aspects of tropical convexity by Develin and Sturmfels in their paper by the same name. There they also establish the relationship between the point set defining a tropical polytope and the tropical hyperplane arrangement determining the same polyhedral complex.

Theorem 2.2.1. *[13]*

If P is a tropical polytope given by the tropical convex hull of a finite point set $V = \{\ \boldsymbol{v}_1, ...\boldsymbol{v}_n\}$ in $\mathbb{TP}^{d-1}$, then P is the union of the bounded regions of the polyhedral decomposition of $\mathbb{TP}^{d-1}$ given by putting an inverted hyperplane at each point, $\boldsymbol{v}_1, ...\boldsymbol{v}_n$.

This inverted hyperplane arrangement is combinatorially equivalent to a tropical hyperplane arrangement with apices $\{\boldsymbol{v}_1, ..., \boldsymbol{v}_n\}$.

Their relationship to tropical hyperplane arrangements alone makes them relevant, but tropical polytopes relate even more closely to the relationship between tropical hyperplane arrangements and products of simplices, which will be examined more in Section 3.3. However, these tropical polytopes themselves are not the focus of my work. Therefore, I will return to a more general discussion of properties of tropical hyperplane arrangements.

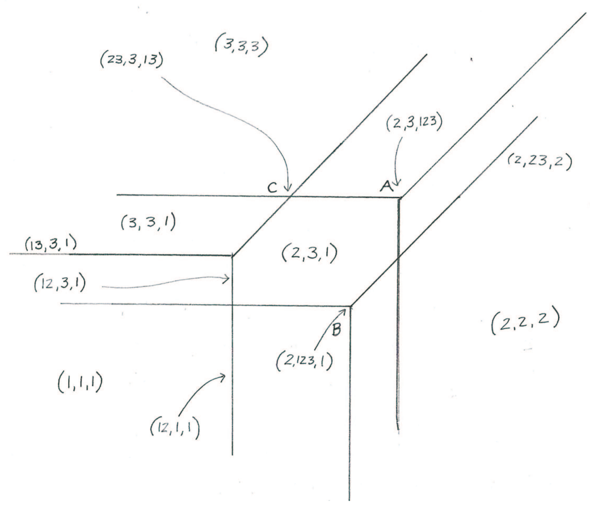

Figure 2.2.6: This shows a partially labeled arrangement of 3 tropical hyperplanes in $\mathbb{TP}^2$.

Points in a specific tropical hyperplane arrangement can be described by their projective coordinates. However, like traditional hyperplane arrangements, often one is only interested in the relative position of a point with respect to each of the hyperplanes. In the traditional setting the relative position of a point in an arrangement of n hyperplanes is described by an n-tuple of $\{+, -, 0\}$ indicating whether the point lies on the positive side, the negative side, or on the hyperplane itself for each of the hyperplanes in the arrangement. In the tropical setting, hyperplanes do not divide d-space into 2 pieces, positive and negative, but into $d + 1$ regions. As a result $\{+, -, 0\}$ is not a sufficient set of choices for describing a point's location with respect to a particular hyperplane. However, there is a tropical analogue to the traditional n-tuple used to indicate the relative position of a point in a tropical hyperplane arrangement.

When one is interested in the purely combinatorial properties of such an arrangement, points are described by their types. The *type* of a point $\boldsymbol{x} \in \mathbb{TP}^{d-1}$ with respect to a tropical hyperplane arrangement $H_1, ..., H_n \in \mathbb{TP}^{d-1}$ is the n-tuple $(A_1, ...A_n)$, where each A_i is the subset of $[d]$ corresponding to the closed sectors of the hyperplane H_i in which $\boldsymbol{x}$ is contained. With respect to algebraic coordinates, for the hyperplane H_i with vertex $\boldsymbol{v}_i = (v_{i1}, ...v_{id})$, A_i indicates which among the $x_j - v_{ij}$ are maximized. It is clear that every point on a face of a given arrangement has the same type. Therefore, these types do not distinguish points within a face, but they do encode all relative information about the faces of the arrangement. Accordingly, the type of any point in a face of the arrangement is called the type of the face. The collection of types of an arrangement determines the arrangement up to combinatorial equivalence (i.e. their relative positions). Figure 2.2.6 depicts a partially labeled arrangement of 3 tropical hyperplanes in $\mathbb{TP}^2$ (see Figure 2.2.4 to recall the standard sector labelings). For reasons that will become obvious in the discussion of tropical oriented matroids, particularly in Section 2.4, I am generally only concerned with distinguishing arrangements up to this equivalence.

2.3 Properties of Tropical Hyperplane Arrangements

Many aspects of the correspondence between tropical hyperplane arrangements and tropical matroid theory have direct analogues in the relationships between standard hyperplane arrangements and traditional matroid theory. However, there is an important difference between traditional hyperplane arrangements and tropical ones. Traditional hyperplanes in d-space either intersect in a $(d-2)$-dimensional linear space (as the planes themselves are $(d-1)$-dimensional linear spaces of d-space), or they are parallel. This is not the case in the tropical setting. All tropical hyperplanes intersect, but they do not all intersect the same way.

Generally, the intersection of two hyperplanes in $\mathbb{TP}^{d-1}$ is a $(d-3)$-dimensional cone. This is not true when an apex of one of the tropical hyperplanes, say H_i, occurs on a proper face of the fan determined by one of the other hyperplanes, say H_j, in the arrangement. The result is that some of the cones in the fan determined by H_i will be proper subsets of H_j's. Arrangements in which this occurs are called *non-generic*. Arrangements that aren't non-generic are called *generic*. The behavior of non-generic tropical hyperplane arrangements differs from that of generic hyperplane arrangements.

In the case of two tropical hyperplanes, the only way an arrangement can fail to be generic is to have a non-generic apex. A *non-generic* apex A of a tropical hyperplane arrangement is an apex of a tropical hyperplane arrangement which is located on a proper face of the fan given by one of the other hyperplanes of the arrangement. Figure 2.3.1 shows two arrangements of three tropical hyperplanes in $\mathbb{TP}^2$, one generic and one non-generic.

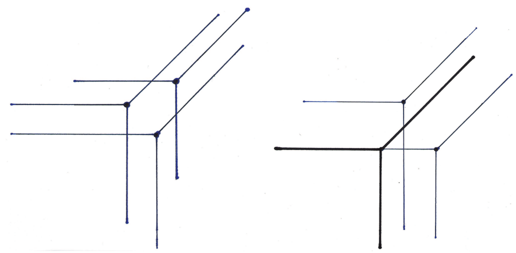

Figure 2.3.1: Here are two arrangements of three tropical hyperplanes in $\mathbb{TP}^2$. The one on the left is generic; the one on the right is non-generic. (The non-generic hyperplane is bold.)

When a tropical hyperplane arrangement includes three or more hyperplanes, the arrangement can be non-generic in another way. In general, any time d hyperplanes in d-space intersect, their intersection will be a point, which is a vertex of an arrangement. An arrangement of tropical oriented hyperplanes in d-space can be non-generic without containing a non-generic apex if there is a vertex of the arrangement determined by the common intersection of $d+1$ tropical hyperplanes. This behavior is no different than a traditional arrangement of three lines intersecting in a common point in the plane. A tropical example of this is depicted in Figure 2.3.2.

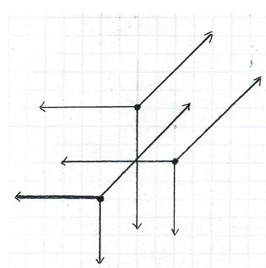

Figure 2.3.2: This is an example of a non-generic arrangement of 3 tropical hyperplanes in 2-space, one without a non-generic apex.

Any zero-dimensional face of a tropical hyperplane arrangement is a vertex of the arrangement, not only the apices. This means there is a definition of non-generic tropical hyperplane arrangements that covers both cases and remains consistent with the definition of the generic tropical hyperplane arrangements stated above. In general, a *non-generic* tropical hyperplane arrangement is a tropical hyperplane arrangement with a non-generic vertex, apex or otherwise. Fortunately, there is a quality that all non-generic vertices have in common.

Lemma 2.3.1. *For a non-generic apex A, in an arrangement of n tropical hyperplanes in $\mathbb{TP}^{d-1}$, the total number of elements summed over its A_i's will be strictly greater than $n + d - 1$.*

Proof. Let M be an arrangement of n tropical hyperplanes in $\mathbb{TP}^{d-1}$ with a non-generic hyperplane. Without loss of generality, assume H_1 is this non-generic hyperplane. By definition, its non-generic apex A occurs on a proper subface of the polyhedral fan given by one of the other hyperplanes, say H_i of M. This means the coordinates (actual coordinates, not type) of A must satisfy equality on at least two of the $\sum c_j x_j = max\{c_1 + x_1, ..., c_d + x_d\}$ defining H_i, say this includes j and k, then A_i contains at least j and k. Since A is H_1's apex, we know $A_1 = [d]$. That gives us at least $d + 2$ entries from 2 positions in A's type. We know there are $n - 2$ positions remaining, and that none of them are empty. Therefore, A's type contains at least $n + d > n + d - 1$ elements.

□

Lemma 2.3.2. *Given a generic tropical hyperplane arrangement M of n hyperplanes in $\mathbb{TP}^{d-1}$, the type A of each apex in M has a total of $n+d-1$ elements summed over its d coordinate sets (the A_i's). In particular, the i^{th} hyperplane will have $A_i = [d]$ and A_j will be a singleton for all $j \neq i$.*

Proof. Without loss of generality, assume that A is the apex of the i^{th} hyperplane. Therefore, by definition, $A_i = [d]$. This means it remains to show that there are exactly $(n-1)$ entries among the $(n-1)$ A_j, such that $j \neq i$. Suppose there are less than $(n-1)$ entries in the remaining A_j. This guarantees the existence of some k such that $A_k = \{\}$. That would mean A failed to define the position of the apex point with respect to H_k. Therefore, A would not actually be the type of the apex. This is a contradiction, so there must be at least $(n-1)$ entries among the remaining A_j's. Suppose there are more than $(n-1)$ entries among the remaining A_j's. This means there is some position, say k such that A_k is not a singleton. That means that the apex with type A maximizes at least two terms of the linear functional, f_k defining H_k. In other words, A describes a root of f_k, and as such, lies on a proper face of H_k. This contradicts the fact that A was a generic apex. As a result, there can not be more than $(n-1)$ entries in the remaining A_j. This proves that there are exactly $(n-1)$ entries in the A_j, $j \neq i$ positions, ensuring exactly $(d+n-1)$ total entries in type A.

□

As mentioned earlier, the apices of a hyperplane arrangement are not its only vertices. Any zero-dimensional intersection of fans given by the hyperplanes of the arrangement is a vertex. The dimension of a face of a hyperplane arrangement in $\mathbb{TP}^{d-1}$ can be defined as a geometric object with respect to the fans of the arrangement.

Lemma 2.3.3. *For a generic vertex A in an arrangement of n tropical hyperplanes in $\mathbb{TP}^{d-1}$, the total number of elements summed over its A_i's will be $n+d-1$.*

Proof. The case where A is an apex was handled in Lemma 2.3.2. It remains to consider the case where A is not an apex. I now consider the general case, where A is not necessarily an apex. No A_i can be empty. This means there are at most $d-2$ elements remaining. An A_i which is a singleton corresponds to the point A lying in a full (d-1)-dimensional chamber of the fan determined by H_i. The presence of a second entry in A_i tells us that A must be contained in some $(d-2)$-dimensional face of H_i. Each additional entry reduces the dimension of the face of H_i containing A (this is consistent with the fact that $A_i = [d]$ when A is the apex of A). This is true for each position A_j in A.

Now recall that the intersection of a k-dimensional linear space and an l-dimensional linear subspace of $(d-1)$-space in general position is a $((d-1)-((d-k)+(d-l)))$-dimensional linear subspace of d-space. For example, the intersection of two lines in the plane (1-dimensional linear subspaces intersecting in 2-space) is a point (a 0-dimensional linear subspace), and the intersection of two planes (2-dimensional linear subspaces) in 3-space is a line (a 1-dimensional linear subspace). More generally, the intersection of a set of m linear subspaces of $(d-1)$-space in generic position, where the dimension of the subspaces are $k_1, ...,$ and k_m respectively, is a $((d-1)-(d-k_1)-(d-k_2)-...-(d-k_m))$-dimensional linear subspace of $(d-1)$-space. For this intersection to be a point, the dimension of linear spaces being intersected must satisfy $(d-k_1)+(d-k_2)+...+(d-k_m) = (d-1)$. Considering this, together with the way the number of entries in a position A_j describes the dimension of the face of H_i containing A, we find that for a point in a generic arrangement of tropical hyperplanes to be a vertex, it must have exactly $(d-1)$ more entries than a point in a full dimensional chamber, or precisely $(n+d-1)$ entries as desired.

□

This pigeonhole characterization of vertices is highly useful when creating programs enumerating all tropical oriented matroids of type (n, d).

2.3.1 Enumeration and Visualization

For reasons that will become more apparent in Section 3.3, understanding the "space" of all arrangements of n tropical hyperplanes in d-space has become a point of increasing interest. To this end, I have code that takes n and d as input and computes, up to combinatorial equivalence, all generic arrangements of n tropical hyperplanes in d-space. The output is a list of the types of all vertices for each distinct arrangement, together with a set of apex coordinates in d-space that realize the arrangement. Clearly, actual coordinate values for apices completely determine an arrangement of tropical hyperplanes, as a tropical hyperplane is completely determined by its apex. However, the types of the apices alone are not enough to determine, even up to combinatorial equivalence, the arrangement. Figure 2.3.1 shows two arrangements of 3 tropical hyperplanes whose apices have the same types, but for which the types of the other vertices differ.

I am also interested in understanding what this space looks like and visualizing what it means for arrangements to be adjacent. Here I define arrangements of tropical hyperplanes to be *adjacent* if there is a unique (non-generic) arrangement "separating" them. That is, a unique arrangement that is real-

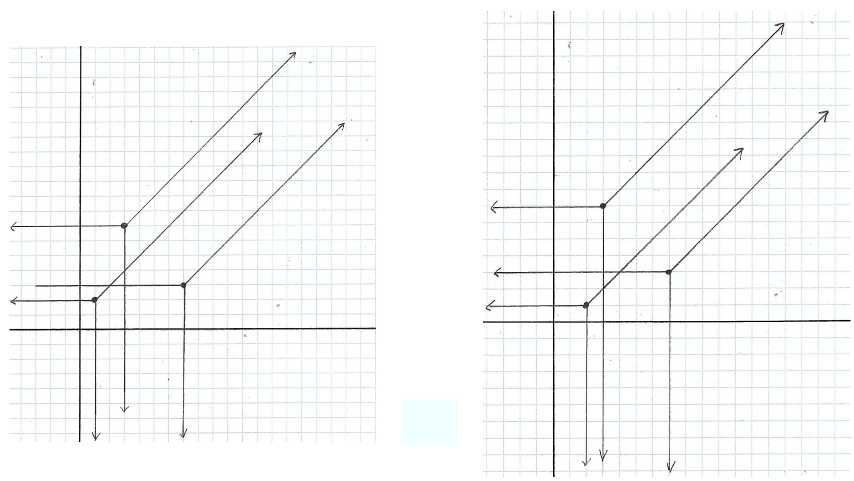

Figure 2.3.3: These are two arrangements of tropical hyperplanes, both of whose apex types are $(123, 3, 3)$, $(2, 123, 2)$, and $(1, 1, 123)$. However, the types of their other vertices differ. The arrangement on the left can be realized by positioning the apices at $(3, 4)$, $(4, 3)$ and $(1, 2)$, while in the figure on the right the third apex is located at $(2, 1)$.

ized by moving one of the hyperplanes towards another in the arrangement in a direction normal to one of the $(d-1)$-dimensional faces of another hyperplane in the arrangement until a non-generic intersection is achieved. This was the inspiration for code that enables me to create rotatable visuals of 3-dimensional arrangements of n tropical hyperplanes. Each hyperplane is colored distinctly, and the planes are truncated by a sphere "at infinity." This allows users to see the 2-dimensional arrangements obtained by projecting in any basis direction, as well as making it easier to distinguish between the hyperplanes determining any given chamber as they use the arrow keys to rotate the image on screen, changing the angle from which the arrangement is viewed.

2.4 Tropical Oriented Matroid Axioms

Tropical oriented matroids are defined via axioms on types, analogous to how traditional oriented matroids can be defined in terms of axioms on their covectors. Ardila and Develin's definition was inspired by both tropical oriented hyperplanes and traditional matroid theory. As such, their discussion of the relationship between tropical hyperplane arrangements and tropical oriented matroids draws many parallels to properties of the relationship between traditional oriented hyperplane arrangements and standard oriented matroid theory.

Understanding the axioms defining a Tropical Oriented Matroid requires a few preliminary definitions.

An (n,d)-*type* is an n-tuple $A = (A_1, .., A_d)$ of non-empty subsets of $[d] = \{1, ...d\}$. The A_i are called the *coordinates* of A; their indices $1, ..., n$ are called the *positions*; and the potential entries $1, ..., d$ are called the *directions*.

The *comparability graph* of two (n, d) types A and B, $CG_{A,B}$ is a semi-directed graph on the vertex set $[d]$. For each position $i \in [n]$, there is an undirected edge between j and k if $j, k \in A_i \cup B_i$ (i.e. both j and k are in position i of both types A and B), and there is an edge directed from j to k if $j \in A_i$ and $k \in B_i$.

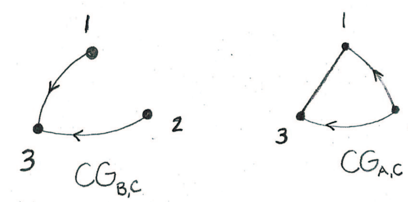

Figure 2.4.1: These are the comparability Graphs $CG_{B,C}$ on types B and C and $CG_{A,C}$ on types A and C seen in Figure 2.2.6.

Generally, a semi-directed graph, known as a *semidigraph*, is a graph in which some edges are directed and some are undirected. A *directed path* from a to b in a semidigraph is a collection of vertices $v_0 = a, v_1, ..., v_k = b$ and a collection of edges $e_1, ..., e_k$, such that the e_i are edges, directed or undirected, from v_{i-1} to v_i, at least one of which is directed. A *directed cycle* in a semidigraph is a directed path from a to itself. A semidigraph is *acyclic* if it has no directed cycles (it may have undirected cycles). This notion of acyclic semidigraphs will be used with comparability graphs in the appropriately named comparability axiom.

Given a type $A = (A_1, .., A_d)$, a *refinement* of A with respect to an ordered partition $P = (P_1, .., P_r)$ of $[d]$ is $A_P = (A_1 \cap P_{m(1)}, ..., A_n \cap P_{m(n)})$ where $m(i)$ is the largest index for which $A_i \cap P_{m(i)}$ is non-empty. A_P is a *total refinement* if each entry of A_P is a singleton.

A *tropical oriented matroid* M with parameters (n, d) can now be defined as a collection of (n, d)-types which satisfy the boundary, elimination, comparability, and surrounding axioms.

Boundary Axiom: for each j in $[d]$, $\boldsymbol{j} = (j, j, ..., j)$ is a type in M.

Elimination Axiom: For any two types A and B in M and any position $j \in [n]$, there is a type C in M such that $C_j = A_j \cup B_j$, and $C_k \in \{A_k, B_k, A_k \cup B_k\}$ for each $k \in [n]$.

Comparability Axiom: For any two types A and B in M the comparability graph $CG_{A,B}$ is acyclic.

Surrounding Axiom: If A is a type of M then every refinement of A is also a type of M.

The spirit of these axioms is natural even if their wording isn't entirely transparent. I will describe the idea of each before moving on to results regarding the properties of these objects. However, I first justify my use of tropical hyperplane arrangements as a valid context for this discussion.

Theorem 2.4.1. *[3] The collection of types in a tropical hyperplane arrangement forms a tropical oriented matroid.*

Proof. The proof considers the axioms in the order presented above for a tropical hyperplane arrangement with apices $\{\boldsymbol{v}_1, ...\boldsymbol{v}_n\}$.

For the boundary axiom, a point with x_j large enough to ensure that $x_j - x_i < v_{kj} - v_{ki}$ for all $k \in [n]$ and $i \neq j \in [d]$ will have the desired type.

For elimination, consider points $\boldsymbol{x}$ and $\boldsymbol{y}$ with types A and B respectively and a position $j \in [n]$. The idea is to construct the tropical line segment containing $\boldsymbol{x}$ and $\boldsymbol{y}$. Each point on this line will satisfy $C_k \in \{A_k, B_k, A_k \cup B_k\}$ for all k. The intersection of this line with the maximal cone of H_j will have $C_j = A_j \cup B_j$, thus satisfying the axiom. Take $a \in A_j$ and $b \in B_j$. Add copies of $(1, ..., 1)$ to obtain coordinates of x and y, such that $x_a - v_{ja} = y_b - v_{jb} = 0$. Now consider the point $\boldsymbol{z} = \boldsymbol{x} + \boldsymbol{y}$. This is the coordinatewise maximum of $\boldsymbol{x}$ and $\boldsymbol{y}$. Therefore, $x_i - v_{ij} = 0$ for $i \in A_j$, maximizing this difference over all i. Similarly, $y_i - v_{ij} = 0$ for $i \in B_j$, also maximizing this difference over all i. As a result, $z_i - v_{ij} = 0$ for all $i \in A_j \cup B_j$, and $y_i - v_{ij} < 0$ for all other i. Call the type of $\boldsymbol{z}$ C. Here we have $C_j = A_j \cup B_j$. Now consider $k \neq j$. We must determine which i maximize $z_i - v_{ki} = max\{x_i - v_{ki}, y_i - v_{kj}\}$. The maximum value of this is $max\{max_i\{x_i - v_{ki}\}, max_i\{y_i - v_{ki}\}\}$. If that maximum comes from $\max\{x_i - v_{ki}\}$, then $z_i - v_{ki}$'s max comes from x and $C_k = A_k$. Similarly, if it comes from $\max\{y_i - v_{ki}\}$, $C_k = A_k$. If the maximums are the same, then $C_k = A_k \cup B_k$. In any case, we have constructed the desired C.

For comparability, we again consider points x and y with types A and B. An edge in their comparability graph $CG_{A,B}$ from i to j means that for some $k \in [n]$ $i \in A_k$ and $j \in B_k$, so $x_i - v_{ki} \geq x_j - v_{kj}$ and $y_j - v_{kj} \geq y_i - v_{ki}$. Combining these inequalities gives $x_i - y_i \geq x_j - y_j$. In the case of a directed edge, one of these inequalities is strict, so $x_i - y_i > x_j - y_j$. Continuing this process along a cycle in $CG_{A,B}$ yields an inequality $0 > 0$, which is not valid. Therefore, $CG_{A,B}$ must be acyclic.

For surrounding, take a point x with type A and a partition $P = (P_1, ..., P_r)$. Let $f(i)$ be the value of j such that $i \in P_j$ and take $\Delta x = \epsilon(f(1), ..., f(d))$. The idea is that for small enough ϵ, $x + \Delta x$ will have type A_P. Pick an index, k. The elements of the k^{th} coordinate of the type of $x + \Delta x$ correspond to i maximizing $(x + \Delta x)_i - v_{ki}$. For small enough ϵ the only candidate i's are those for which $x_i - v_{ki}$ is maximized for this i, which means $i \in A_k$. From these possible i's $(x + \Delta x)_i - v_{ki}$ is maximized if and only if $\Delta x_i = \epsilon f(i)$ is maximized, meaning i is maximal among the A_k with respect to P.

□

Armed with a concrete, visualizable, class of examples of tropical oriented matroids, I now offer a more intuitive description of the ideals presented in the formally stated axioms, in the context of tropical hyperplane arrangements.

It may not seem obvious that any specific type would appear in a tropical oriented matroid. However, all tropical hyperplanes are translates of each other. This means that traveling far enough in any basis direction will eventually land one in the sector corresponding to that direction with respect to each hyperplane. More concretely, heading to infinity in any basis direction, say i, one can reach a position where the i^{th} coordinate maximizes the linear equation determining each of the hyperplanes in the arrangement. The type of such a position is $(i, ..., i)$. This idea is captured in the boundary axiom.

The elimination axiom describes what happens as one walks from one point to another within a tropical hyperplane arrangement. In a traditional hyperplane arrangement traveling along the line containing two points whose corresponding covectors differ at a position, one having a $+$ and the other a $-$, you must pass through a point whose corresponding covector has a 0 in that position. Geometrically, this corresponds to crossing the hyperplane in whose position in which they differ. The idea is a bit more notationally complicated in the tropical case due to the existence of more than two directions, but the idea is the same. To travel between two tropical points whose coordinates differ in some position on the tropical line connecting the two points one must cross through a point whose coordinates at that position are the union of the two points. This is the equivalent of crossing the tropical hyperplane that separates them. If we consider the two direction case, letting our directions be denoted by $+$ and $-$ and letting the union $\{+, -\}$ be represented by 0 we replicate the condition we are accustomed to in the traditional case. Figure 2.4.2 offers an example of this.

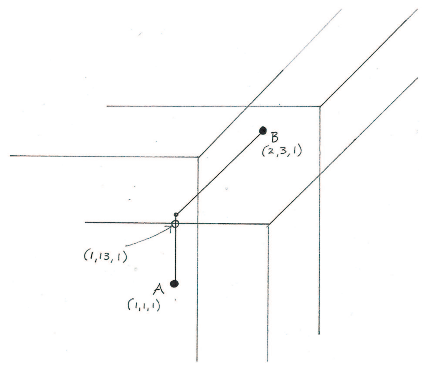

Figure 2.4.2: This illustrates the tropical line connecting points whose second coordinates differ and where the tropical segments containing them cross faces of the arrangement whose second coordinate is the union of their second coordinates, thus depicting how the arrangement satisfies the elimination axiom with respect to these points.

The comparability axiom may be the least intuitively stated of the axioms. An edge from j to k in the comparability graph $CG_{A,B}$ corresponds to the idea that if A has a j in its i^{th} position, and B has a k in its i^{th} position you must travel more in the k direction than in the j direction to get from A to B. A cycle in this graph would amount to contradictory information regarding which direction to travel in order to move from A to B.

The surrounding axiom captures the idea that given a type A for which not all entries are singletons, a refinement of A can be obtained by moving infinitesimally away from A. In particular, if both j and k appear in the i^{th} coordinate of A's type, then a refinement of A can be obtained by moving infinitesimally in the j direction, or the k direction, thus breaking the tie, $x_j - v_{ij} = x_k - v_{ik}$. Figure 2.4.3 illustrates an example of a vertex C in $\mathbb{TP}^{d-1}$ and how moving infinitesimally in all possible directions results in points whose types realize all refinements of C.

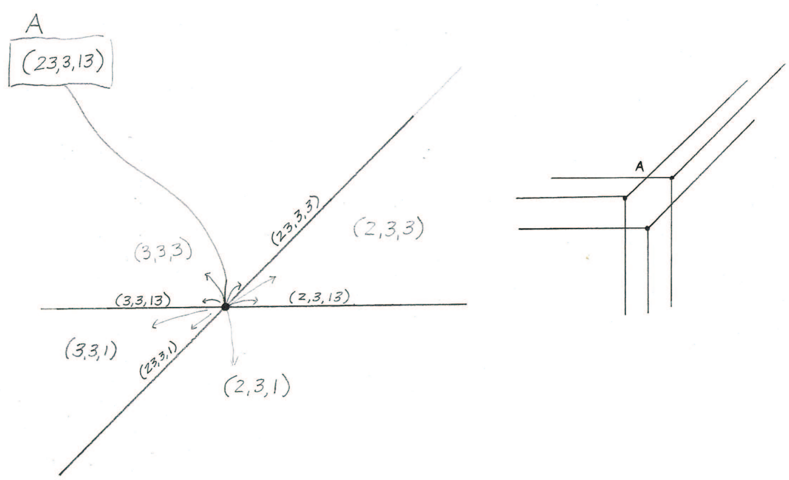

Figure 2.4.3: This image shows how all refinements of C can be obtained by moving infinitesimally in the appropriate direction.

Having both a technical and colloquial understanding of what the axioms are saying, together with a class of examples that can be explicitly written down and played with, is helpful when exploring properties of tropical oriented matroids.

2.5 Characterizations of Tropical Matroids

As traditional matroids can be described by their maximal independent sets or their minimal dependent sets, tropical oriented matroids can be given by either their topes or their vertices alone. The following results establish this. While the properties are nice on their own, they are also practically useful. They provide a straightforward inductive approach to proving properties of arbitrary tropical oriented matroids, namely that a property can be proved by showing that it holds for all vertices of a tropical oriented matroid, and that if that property holds for a type it holds for any refinement of that type.

Lemma 2.5.1. *[3]*

Refinement is transitive: if C is a refinement of B and B is a refinement of A, then C is a refinement of A.

Proof. Given a refinement B of A with respect to ordered partition $(P_1, ..., P_r)$, and a refinement C of B with respect to ordered partition $(Q_1, ...Q_s)$, let $X_{ij} = P_i \cap Q_j$. C is a refinement of A with respect to this order partition $(X_{11}, X_{12}, ...X_{1s}, X_{21}, ...X_{rs})$.

□

Lemma 2.5.2. *[3]*

If A and B are types of a tropical oriented matroid such that $B_i \subseteq A_i$ for all $i \in [n]$, then B is a refinement of A.

Proof. Suppose this is not true, that there exists some tropical oriented matroid M with types A and B such that $B_i \subseteq A_i$ for all $i \in [n]$, but B is not a refinement of A. This means we can't break ties in each A_i in a way that makes the elements in B_i maximal among those in A_i. Algebraically, this says that we can't solve the system of equations given by $x_j = x_k$ for $j, k \in B_i$ and $x_j > x_k$ for $j \in B_i, k \in A_i \backslash B_i$. Linear duality tells us that this means that the system must be inconsistent; that is, that some linear combination of them yields $0 > 0$. However, a collection of such inequalities would produce a cycle in the comparability graph $CG_{A,B}$. This would mean A and B couldn't be types of the same oriented matroid, contradiction.

□

These refinement properties are particularly useful when working with vertices of tropical oriented matroids. Before presenting results regarding minimal and maximal dimensional faces of tropical oriented matroids, it must be clear what is meant by the dimension of a face in a tropical oriented matroid.

The dimension of a face of a tropical oriented matroid has an abstract definition that agrees with the geometric intuition given by tropical hyperplane arrangements. Given a type A of a tropical oriented matroid, consider the undirected graph G_A, with vertex set $[d]$ and an edge between i and j if there is some coordinate A_k such that $i, j \in A_k$. The *dimension* of type A is one less than the number of connected components of G_A. A *vertex* of a tropical oriented matroid is a type A for which G_A is connected (i.e. a type of dimension 0). [3] A *tope* of a tropical oriented matroid is a type $A = \{A_1, A_2, ..., A_n\}$ such that each A_i is a singleton, making G_A edge-free. A tope is a $(d-1)$-dimensional type.

One would hope that this definition of dimension would be consistent with the intuition of dimension offered by the case of tropical hyperplane arrangements, where one identifies vertices as zero-dimensional faces of the arrangement. However, the proof that tropical hyperplane arrangements are tropical oriented matroids does not make it entirely obvious why that is true. So I present a separate statement and justification here.

Lemma 2.5.3. *Zero-dimensional faces of tropical hyperplane arrangements are vertices of the corresponding tropical oriented matroid.*

[3]The proofs and counting arguments presented in Section 2.3 can be phrased with respect to this dimension graph. Some find that approach more intuitive; however, since the results do not depend on the additional structure, I decided to present them independent of the matroid terminology.

Proof. Let A be a zero-dimensional face of a tropical hyperplane arrangement.

I first address the generic setting. There are two cases to consider; first, when a zero-dimensional face is an apex, second when it is not.

First consider the case where A is an apex of the arrangement. Without loss of generality, say A is the apex of H_i. We know $A_i = [d]$. The other A_j are non-empty. Let k be an element of A_j. Each of the d vertices of G_A will have an edge to k. This guarantees that G_A is connected, making A a vertex.

Now consider the non-apex case. Given the generic setting, we know the number of elements summed over the n positions of A is $(n+d-1)$. For G_A to not be connected would mean that there is at least one element of $[d]$ that does not appear in any position of A. This means that there must be two positions with two common elements. However, if this is the case then the point is contained in parallel linear spaces, as corresponding faces of tropical hyperplanes are translates of each other. This means that the point can not be a generic zero-dimensional face, as it lies in overlapping parallel faces of different hyperplanes. Since the arrangement was assumed to be generic, this is a contradiction. Therefore, every element of $[d]$ must appear and G_A must be connected, making A a vertex.

To finish the proof, I explain that the non-generic setting is fundamentally the same.

First, notice that nothing about the argument for the apex case depended on genericity, so it is still valid.

For the non-apex non-generic case, again the key to showing that G_A is connected is proving that all elements of $[d]$ appear over the more than $(n+d-1)$ entries in the positions of A. There are two subcases here. The first is when A is a zero-dimensional face that is not non-generic because it lies in (overlapping) parallel faces of multiple hyperplanes, say H_i and H_j. This means $A_i \subseteq A_j$. In this case, it is enough to consider the subgraph of G_A obtained by replacing A_i (by any element $k \in A_i$. This subgraph of G_A will be connected if and only if G_A was. However, since this subgraph corresponds to the generic tropical hyperplane arrangement obtained by moving A just slightly in the k direction, proving that this subgraph is connected follows from the generic case. The remaining type of non-generic zero-dimensional face of an arrangement of n tropical hyperplanes in d-space is where A is an intersection of $d+1$ or more $(d-1)$-dimensional linear subspaces (without loss of generality, assume an intersection of exactly $d+1$ iterating the construction if necessary for the case of $> d+1$). In particular, this says A is a common intersection of $d+1$ hyperplanes, when a common intersection of any d of them would determine the same zero-dimensional face. In this case, it is again enough to consider a subgraph of G_A. This time take the subgraph obtained by replacing one of the $d+1$ 2-element sets indicating which of the $d+1$ hyperplanes A lies on a proper face of, say A_j, by either element of A_j. Again this graph will be connected if G_A is and it is, like the previous case, connected by the generic argument, completing the proof that A is a vertex.

□

Theorem 2.5.4. *[3]*

A tropical oriented matroid M is completely determined by its topes. In particular, a type A is in M if and only if A satisfies the following conditions:

-A satisfies the compatablity axiom with every tope of M (i.e. given a tope T of M $CG_{A,T}$ is acyclic).

-Every total refinement of A is a tope of M.

Proof. The comparability axiom requirements are addressed first by noting that if A satisfies the conditions, then every refinement B of A will as well. This means the set of total refinements of B is a subset of the set of total refinements of A, and for each T the comparability graph $CG_{B,T}$ is a subgraph of the comparability graph $CG_{A,T}$. Consider a minimal counterexample to refinement, an n-tuple A such that each refinement of A is in M, but A itself is not in M. Since A has a non-trivial refinement, some element of A is not a singleton. Without loss of generality, assume $\{1, 2\} \subset A_1$. Consider the subgraph $G_A \backslash 1$ of A's dimension graph G_A obtained by deleting vertex 1 and all edges incident to it. Some component of this graph contains 2. Without loss of generality, let this subset of $[d]$ be $S = \{2, ..., r\}$ and let $T = \{r + 1, ..., d\}$. Take the refinements B and C of A given respectively by the ordered partitions $(S, 1 \cup T)$ and $(1 \cup T, S)$. $S_1 \neq A_1$ and $T_1 \neq A_1$, so B and C are proper refinements of A and hence in M, by the minimality assumption. Eliminating 1 between B and C yields some element $D \in M$. Showing that $D = A$ confirms that $A \in M$. This is done by position. In position 1, $D_1 = B_1 \cup C_1 = 1 \cup (A_1 \backslash 1) = A_1$, and for $i \neq 1$, $D_i \in \{B_i, C_i, B_i \cup C_i\}$.

S and T don't overlap and A_i cannot contain elements from both S and T. If $A_i \subset S$ or $A_i \subset 1 \cup T$, then $B_i = C_i = A_i$ and $D_i = A_i$ as desired. The remaining case is where A_i contains 1 together with a subset $X \subset S$. The proof that this also guarantees $D_i = A_i$ is a lemma by Ardila and Develin, whose proof follows from constructing a tope refinement of D for which $CG_{A,U}$ contains a cycle, giving the desired contradiction.

□

Theorem 2.5.5. *[3]*

A tropical oriented matroid is completely determined by its vertices. In particular, all types of an oriented matroid are refinements of its vertices.

Notice that this does not say that the apices determine a tropical oriented matroid. As I showed in 2.2.2, while specific coordinates for apices in $\mathbb{TP}^{d-1}$ do determine an arrangement, in general the types of apices alone do not. For an example, see Figure 2.3.1 in Section 2.3.1.

Ardila and Develin provide proof of Theorem 2.5.5 by example. I sketch the idea.

Proof. (sketch)

The idea is to show that any type $A = (A_1, ..., A_n) \in M$ that is not a vertex is a refinement of some type in M. Lemma 2.5.2 shows that it is enough to find a type in M strictly containing A. The fact that A is not a vertex means that its dimension graph G_A is not connected. There are two cases to consider.

The first case is when all numbers appearing in any of the A_i are contained in one connected component of G_A, with some element, say $i \in [d]$ contained in none of them. In this case, eliminating A with $i = \{i, ..., i\}$ in position 1 yields a type $B = \{A_1 \cup i, B_2, ..., B_n\}$ where each B_j is $i, A_j \cup i$ or A_i. If no $B_j = i$, A is a refinement of $B \in M$; otherwise, there is some k such that $B_k = i$. In this case, eliminate B with A in position k and repeat this process as many times as necessary to obtain some type C such that C_j is A_j or $A_j \cup i$ for all j.

In the second case, elements of the A_i appear in multiple components of G_A. This case is again an inductive construction, but it first uses elimination with multiple $i = [i, ..., i]$ to build a type D that agrees with A's intersection with one of the components of G_A. That is to say, each D_j is equal to A_j's intersection with the numbers contained in one connected component of G_A. This type is then "connected" to numbers appearing in other components of G_A via the approach taken in case one. This results in a type E such that $E_j \subset A_j$ for all j such that E_j is connected. This construction guarantees that $E \in M$ and A is a refinement of E.

□

In light of Theorem 2.4.1 and Lemma 2.5.5, I define generic and non-generic abstract tropical oriented matroids based on the combinatorial properties of their vertices. A *generic* vertex of a type-(n, d) tropical oriented matroid is a vertex with exactly $n + d - 1$ entries summed over its A_i's. A vertex of a type-(n, d) tropical oriented matroid is *non-generic* if the number of entries summed over its A_i's is $\geq n + d$. A *non-generic* tropical oriented matroid is a tropical oriented matroid containing a non-generic vertex. A tropical oriented matroid is *generic* if each of its vertices is generic.

These definitions lead to the following corollary of Theorem 2.4.1.

Corollary 2.5.6. *Non-generic tropical hyperplane arrangements are non-generic tropical oriented matroids and generic tropical hyperplane arrangements are generic tropical oriented matroids.*

2.6 Tropical Oriented Matroid Duals

Duality properties are central to traditional matroid theory. Before defining the dual to a tropical oriented matroid, I need to introduce some definitional support structure.

A *semitype*, with parameters (n, d), is given by an n-tuple of subsets of $[d]$, possibly empty. The *completion* $\tilde{M}$ of a tropical oriented matroid M consists of all semitypes which can be obtained from types of M by changing some of the subsets of the coordinates to the empty set. The *reduction* of a collection of semitypes is the subset of all actual types in the collection.

Given a semitype A with parameters (n, d), the *transpose* A^T of A is a semitype with parameters (d, n) with $i \in A_j^T$ if $j \in A$. You can interpret an (n, d)-type A as an $n \times d$ 0/1 matrix with a 1 in position ij if $j \in A_i$ and a 0 otherwise. In this context, A^T is obtained by taking the transpose of the matrix corresponding to A and reading off the resulting (d, n)-type (a d-tuple of subsets of $[n]$). Figure 2.6.1 depicts a type (specifically a vertex), its associated 0/1 matrix, and the semitype (in this case itself a type) obtained by taking its transpose. As I will show in Lemma 2.6.2, the semitype obtained by taking the transpose of a vertex is an actual type.

0	0	1
1	1	1
1	0	0

0	1	1
0	1	0
1	1	0

Figure 2.6.1: This illustrates the 0/1 matrix associated to a vertex type $A = \{3, 123, 1\}$ of a tropical oriented matroid M, and the matrix obtained by taking the transpose of A, which yields the type $\{23, 2, 12\}$ in $\tilde{M}$.

Given these definitions, the dual of a tropical oriented matroid can be defined. The *dual* of a tropical oriented matroid M is the reduction of the collection of semitypes given by the transposes of semitypes in $\tilde{M}$. Figure 2.6.2 shows an arrangement of three tropical hyperplanes in the tropical plane and the arrangement obtained by taking its dual matroid.

In theory, this may seem like a daunting process, but in practice, the following lemma makes the procedure for finding an oriented matroid's dual quite manageable.

Corollary 2.6.1. *Given a tropical oriented matroid M, its dual is made up of all refinements of the transposes of the vertices of M.*

This tells us that to find the dual of a tropical oriented matroid, one must only worry about taking the transpose of the vertices and that refinements of those alone will be enough to produce all types in the reduction of the transposes of $\tilde{M}$. This result follows from the following lemma and Theorem 2.5.5.

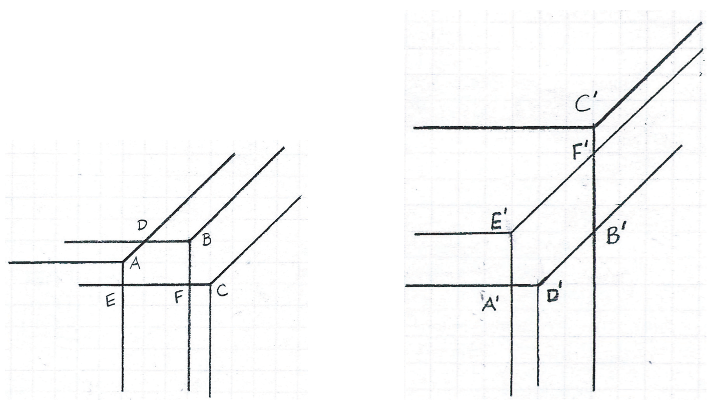

Figure 2.6.2: The image on the left is the tropical hyperplane arrangement corresponding to a tropical oriented matroid M. The image on the right is the arrangement corresponding to M's dual. The vertices of the dual have been labeled to show which vertex of M's transpose they represent.

Lemma 2.6.2. *Given a tropical oriented matroid M, the transpose of any vertex of M will be a vertex in M's dual.*

Proof. First notice that by Lemma 2.3.3 we know that the only candidates for vertices of M's dual are transposes of vertices of M. Translated into the matrix interpretation of types, 2.3.3 says that any vertex of M's dual must have $(d+n-1)$ non-zero matrix entries. This means it must be a transpose of a type of M with $(d+n-1) = (n+d-1)$ non-zero entries or a refinement of a type with at least $(n+d)$ non-zero entries. Since no types of M have more than $(n+d-1)$ non-zero entries, the second option does not occur. Therefore, the vertices of M's dual must come from the transposes of M's vertices. Therefore, it is enough to check that each of these transposes is an honest to goodness type and not just a semitype.

The only way the transpose A^T of a type A of M could fail to be a type is if some row of its $d \times n$ matrix is empty. This would mean that some column of A's matrix was empty, and would imply that there was some $i \in [d]$ such that i was not contained in any of the A_j. I offer two proofs that this does not occur. The second is sleeker, but the first might be more intuitive.

The first proof is based on the incidence matrix presentation of tropical matroid types and has two cases. The first case is where A is an apex of M. In this case we know $A_j = [d]$ for some $j \in [n]$. Therefore, we know $i \in A_j$ for all $i \in [d]$, and each i is represented. The second case is when A is not an apex of M. Assume that i is the unrepresented element of $[d]$. We know that the $(n+d-1)$ non-zero elements of the matrix are spread over the other $(d-1)$ columns of the matrix. We also know that none of the n rows of the matrix are empty (since A was a type in M). This means that there are $(d-1)$ non-empty elements to spread over the $(d-2)$ remaining positions of the rows. We know the i^{th} position of each row is empty and some other position of each row is already occupied, so by the pigeonhole principle, there is some $k \in [d]$ that occurs as a non-singleton entry of at least two rows. This means that these types are not generic with respect to that direction, meaning the point would fail the surrounding axiom. Therefore, A could not be a vertex of a tropical oriented matroid, and the non-zero entries of the rows must be spread over all columns.

The second proof is based on the dimension graph. Again suppose there is some $i \in [d]$ that does not occur in any of the A_j's. That means there is no edge of G_A containing i, guaranteeing that G_A is not connected. This tells us the dimension of A is greater than 0, contradicting the fact that A is a vertex of M.

□

Figure 2.6.2 illustrates an example of how vertices in the original correspond to vertices in the dual. This shows that it is not generally the case that apices in the original will correspond to apices in the dual. However, it is possible to predict which vertices will correspond to apices in the dual when looking at a tropical hyperplane arrangement.

Taken with the definition of generic and non-generic vertices, a straightforward counting argument offers the following immediate corollary to this lemma.

Corollary 2.6.3. *The transpose of a generic vertex of a tropical oriented matroid M will be a generic vertex in M's dual, and the transpose of a non-generic vertex will be a non-generic vertex in the dual.*

As a direct result of Lemma 2.6.2 and Corollary 2.6.1 I state the following corollary.

Corollary 2.6.4. *To prove that the dual to a tropical oriented matroid M is itself a tropical oriented matroid, it is enough to show that the transposes of the vertices of M satisfy the tropical oriented matroid axioms.*

Taken together, Corollaries 2.6.3 and 2.6.4 yield the following:

Corollary 2.6.5. *The dual to a tropical oriented matroid M is generic if and only if M is a generic tropical oriented matroid.*

The fact that the dual to any tropical oriented matroid is again a tropical oriented matroid seems a requirement for an object to have the word matroid in its name, but this is not so obvious. In fact, it appeared as a conjecture in Ardila and Develin's work. It also follows directly from their primary conjecture on tropical oriented matroids.

Conjecture 1. *[3] There is a one-to-one correspondence between the subdivisions of the product of simplices $\Delta_{n-1} \times \Delta_{d-1}$ and the tropical oriented matroids with parameters (n, d).*

Understanding this perspective on tropical oriented matroids and their relationship to tropical hyperplane arrangements requires some background on $\Delta_{n-1} \times \Delta_{d-1}$.

Chapter 3

Products of Simplices

3.1 Introducing Products of Simplices

The *product* of two polytopes $P \subset \mathbb{R}^d$ and $Q \subset \mathbb{R}^n$ is $P \times Q = \{(\boldsymbol{v}, \boldsymbol{w}) \mid \boldsymbol{v} \in P,\ \boldsymbol{w} \in Q\} \in \mathbb{R}^{d+n}$. This is what is meant by product when discussing a *product of two simplices* $\Delta_n \times \Delta_d$. These products are studied by many people for a variety of reasons. While some are interested in them for their own merits, [5], [8], [20], [34], they also have applications to Schubert calculus, [2], hom-complexes, [33], growth series of root lattices, [1], transportation polytopes and Segre embeddings, [43]. They appear here due to their usefulness as building blocks of triangulations and their relationship to tropical oriented matroids.

It is probably not surprising that there is significant interest in products of simplices for their potential applications to triangulating product polytopes. However, it may initially be less obvious how these simplicial products relate to tropical oriented matroids. Ardila and Develin conjectured the existence of a one-to-one correspondence between subdivisions of products of simplices

$\Delta_{n-1} \times \Delta_{d-1}$ and tropical oriented matroids with parameters (n, d). Oh and Yoo established that triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ correspond to tropical oriented matroids by confirming that they satisfy the elimination axiom. Armed with the generic and non-generic matroid distinction, one sees that the proof of the bijection they originally claimed between triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ and tropical oriented matroids failed because it does not hold for non-generic cases. Here, I provide some background in simplicial products and subdivisions along with an application to triangulating product polytopes before clarifying the correspondence between generic and non-generic tropical oriented matroids and subdivisions of $\Delta_{n-1} \times \Delta_{d-1}$ in Section 3.3 and introducing a distinction between the regular and non-regular cases.

3.2 Triangulating $\Delta_{n-1} \times \Delta_{d-1}$

3.2.1 Representing $\Delta_{n-1} \times \Delta_{d-1}$ and its Triangulations

It would be nice if a product of two triangulations was itself a triangulation; however, this is not the case. This is why triangulating products of simplices is so important. While a product of triangulations is not itself a triangulation, a product of two triangulations is a subdivsion into products of simplices. Therefore, understanding products of simplices is important for understanding subdivisions and triangulations of product polytopes. The product of two subdivsions, and hence the product of two triangulations, is defined analogously to a product of polytopes. In particular, given Δ_{n-1} and Δ_{d-1} with vertices $\{\boldsymbol{v}_1, ..., \boldsymbol{v}_n\}$ and $\{\boldsymbol{w}_1, ..., \boldsymbol{w}_d\}$ respectively, the vertices of $\Delta_{n-1} \times \Delta_{d-1}$ are given by the nd points $(\boldsymbol{v}_i, \boldsymbol{w}_j)$ where $i = 1, ..., n$ and $j = 1, ..., d$.

There are two ways to model these graphically: a grid representation and a bipartite graph representation. Figure 3.2.1 depicts the product $\Delta_2 \times \Delta_3$ along with both its grid and bipartite graph representations.

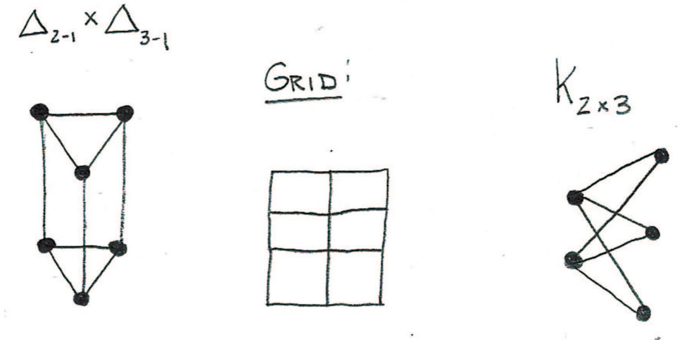

Figure 3.2.1: The product of simplices $\Delta_{2-1} \times \Delta_{3-1}$ is a triangular prism. It is shown here along with its corresponding complete bipartite graph $k_{2,3}$ and 2 by 3 grid.

In the grid representation, the product is represented by an n by d grid with a column for each $\boldsymbol{v}_i$ and a row for each $\boldsymbol{w}_j$. The vertices of the product correspond to the boxes of the grid, and faces of the product correspond to the sub-grids that result from deleting some rows and/or columns from the grid. A cell in a subdivision of the product is represented by marking the boxes corresponding to the included vertices. Thus, a triangulation can be represented by a set of such marked grids, each of which corresponds to a simplex in the subdivision. In the translation between subdivisions of $\Delta_{n-1} \times \Delta_{d-1}$ and tropical oriented matroids, this grid representation is equivalent to the incidence matrix representation discussed in 2.

Recall that the *complete bipartite graph* $K_{n,d}$ is the graph on the vertex set $V_1 \cup V_2$ where $V_1 = [n]$ and $V_2 = [d]$ with an edge connecting each vertex in V_1 with each vertex in V_2, but no edges between vertices of V_1 or V_2. In the bipartite graph representation, vertices of $\Delta_{n-1} \times \Delta_{d-1}$ are represented by edges in the bipartite graph $k_{n,d}$. In this setting, faces of $\Delta_{n-1} \times \Delta_{d-1}$ correspond to induced subgraphs $K_{n,d}$, and facets correspond to subgraphs induced by removing one vertex. A cell C of a sudivision of $\Delta_{n-1} \times \Delta_{d-1}$ is represented by a subgraph of $k_{n,d}$ consisting of edges corresponding to the vertices of C. Therefore, a triangulation $\mathcal{T}$ can be represented by a collection of subgraphs of $k_{n,d}$, one for each simplex in $\mathcal{T}$. Figure 3.2.1 shows an example of a product of simplices and the corresponding complete bipartite graph and grid representations. Figures 3.2.2 and 3.2.3 depict a triangulation of that polytope together with both the subgraphs of the bipartite graph and marked grids corresponding to each simplex.

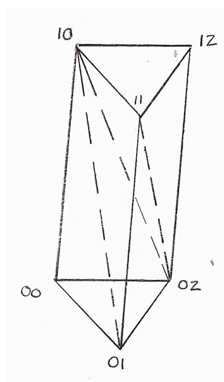

Figure 3.2.2: This shows a triangulation of the product of simplices from Figure 3.2.1.

Further, in the bipartite graph representation, we have the following results regarding properties of the point configuration $\Delta_{n-1} \times \Delta_{d-1}$ and subgraphs of $K_{n,d}$.

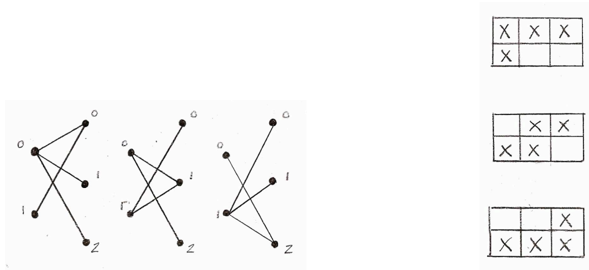

Figure 3.2.3: This shows both the subgraphs of $k_{2,3}$ and the marked grids corresponding to each simplex of the triangulated prism seen in 3.2.2.

Theorem 3.2.1. *[11]*

In the bipartite graph representation of $\Delta_{n-1} \times \Delta_{d-1}$:

(1) A subset of $\Delta_{n-1} \times \Delta_{d-1}$ is affinely independent if and only if the corresponding subgraph has no cycles (this means it is a forest). So affine bases correspond to spanning trees.

(2) A subset of $\Delta_{n-1} \times \Delta_{d-1}$ is affinely spanning if and only if the corresponding subgraph is connected and spanning.

(3) A subset of $\Delta_{n-1} \times \Delta_{d-1}$ is a circuit if and only if the corresponding subgraph is a cycle.

Proof. Consider a subgraph G of $K_{n,d}$ and the corresponding subset of $\Delta_{n-1} \times \Delta_{d-1}$, call it C. For the forward direction of (1), the fact that G, being an acyclic subgraph, corresponds to an independent subset of vertices of $\Delta_{n-1} \times \Delta_{d-1}$ follows from an inductive argument on the size of the set (recall that the number of edges in G corresponds to the number of vertices of C). Any acyclic graph is guaranteed to have at least 1 degree 1 vertex. The subgraph of $K_{n,d}$ that results from deleting this vertex corresponds to a facet of $\Delta_{n-1} \times \Delta_{d-1}$ containing all but one point of C. The inductive hypothesis tells us that the intersection of this facet with C is an independent set, which implies that C is itself independent. The converse statement amounts to the fact that a cycle in $K_{n,d}$ corresponds to an affinely dependent set. We can see this by taking the vertices of $\Delta_{n-1} \times \Delta_{d-1}$ corresponding to the edges in the cycle of $K_{n,d}$ and labeling them with alternating coefficients of $+1$ and -1. This construction yields the desired dependence, as each $(\boldsymbol{v}_i, *)$ will occur with both positive and negative coefficients, as will each $(*, \boldsymbol{w}_j)$, so everything will cancel.

Both (2) and (3) follow from (1).

To see (2), recall that a set is affinely spanning if and only if it contains an affine basis, meaning a full dimensional simplex, and that subgraphs are connected and spanning if and only if they contain a spanning tree. So connected spanning subgraphs correspond to sets containing an affine basis of $\Delta_{n-1} \times \Delta_{d-1}$ and are hence affinely spanning.

(3) follows from the fact that minimal dependent sets correspond to circuits. Minimal dependent sets correspond to subgraphs containing cycles for which the result of removing any edge is an acyclic graph. This is only true when the subgraph itself is a cycle, so circuits of $\Delta_{n-1} \times \Delta_{d-1}$ correspond to cycles in $K_{n,d}$.

□

This circuit terminology clearly has ties to traditional matroid theory. However, it also has connections to results in tropical matroid theory. In light of the translation between tropical oriented matroids and products of unimodular simplices, which will be discussed further in 3.3, Lemma 2.3.3 is equivalent to the following statement.

Lemma 3.2.2. *[11]*

A subset of the grid corresponds to a spanning subset of $\Delta_{n-1} \times \Delta_{d-1}$ if and only if it meets all rows and columns of the grid. In particular, bases are the subsets of cardinality $n + d - 1$ that meet every row and column.

Proof. Theorem 3.2.1 says that spanning subsets correspond to connected spanning subsets in the bipartite graph representation. This means they correspond to subgraphs meeting every vertex of the graph. Vertices of the graph correspond to rows and columns of the graph, so the translation shows that spanning corresponds to meeting each row and column of the grid.

□

3.2.2 Nice Triangulations of $\Delta_{n-1} \times \Delta_{d-1}$

It should not be surprising that people are interested in not only products of subdivisions and triangulations themselves, but also whether properties of subdivisions are "inherited" in their products. To demonstrate ties between these products of simplices and my work with triangulations of lattice polytopes, I consider the case of products of unimodular simplices. This offers a context in which to introduce staircase triangulations, which are particularly useful for proving a more general result regarding the existence of nice triangulations in product polytopes.

A product of unimodular simplices corresponds to the undirected edge polytope of a complete bipartite graph, and hence is itself totally unimodular. To understand this statement it is helpful to know what an edge polytope is. Let G be a finite connected graph, on a vertex set $V(G) = \{1, ..., d\}$, such that G has no loops or multiple edges. Given an edge $e = \{i, j\}$ of G joining vertices $i \in V(G)$ and $j \in V(G)$, define $\rho(e) \in \mathbb{R}^d$ as $\rho(e) = \boldsymbol{e}_i + \boldsymbol{e}_j$, where

$\boldsymbol{e}_i$ is the ith unit coordinate vector in $\mathbb{R}^d$. The *edge polytope* of G, $\mathcal{P}_G$, is the convex hull of $\{\rho(e) \mid e \text{ is an edge of G }\} \subset \mathbb{R}^d$.

Results from Section 1 ensure that a general edge polytope $\mathcal{P}_G$ has a regular unimodular triangulation if and only if there is a squarefree initial ideal of the toric ideal $I_{\mathcal{P}_G}$ of $K[\mathcal{P}_G]$. In the special case of a product of simplices, we know that there is in fact a quadratic triangulation, and hence we know we can obtain a quadratic Gröbner basis for the corresponding ideal. Such a nice triangulation can be obtained for the product $\Delta_{n-1} \times \Delta_{d-1}$ via a staircase triangulation.

The name of the staircase triangulation fits nicely with its visual representation via grids. A *monotone staircase* in the $n \times d$ grid is a subset of $n + d - 1$ boxes, beginning with $(1, 1)$ and ending with (n, d), where each is directly above or to the right of the one that came before it. Another way to describe this is to consider the vertices of $\Delta_{n-1} \times \Delta_{d-1} = \{(\boldsymbol{v}_i, \boldsymbol{w}_j) \mid 1 \leq i \leq n, 1 \leq j \leq d\}$ and think of paths from $(\boldsymbol{v}_1, \boldsymbol{w}_1)$ to $(\boldsymbol{v}_n, \boldsymbol{w}_d)$, where each step increases either the index of $\boldsymbol{v}$ or the index of $\boldsymbol{w}$ by one. Every path of this type will contain $n + d - 1$ vertices of $\Delta_{n-1} \times \Delta_{d-1}$, and hence will determine an $(n+d-2)$-dimensional simplex. The collection of all such paths corresponds to the staircase triangulation. In fact, this is the same triangulation obtained by pulling the vertices in lexicographical order. [1] Figure 3.2.4 depicts the collection of simplices, each given by the grid indicating which vertices are contained in it, comprising the monotone staircase triangulation of $\Delta_3 \times \Delta_4$. A straightforward counting argument demonstrates that the $n \times d$ grid has $\binom{n+d-1}{n-1}$ monotone staircases [11] [24].

[1] For more on pulling triangulations and lexicographical order see Sections 1.3.1 and 1.2.1.

Staircase triangulations can also be described algebraically. Let $\prec$ be a component-wise partial order on the vertices of $\Delta_n \times \Delta_d$ given by ordering the vertices of the factors $v_0 \prec v_1 \prec ... \prec v_n$ and $v'_0 \prec v'_1 \prec ... \prec v'_d$. The collection of totally ordered subsets of this partial ordering agrees with the quadratic triangulation known as the staircase triangulation. This triangulation can be constructed by pulling the vertices (v_i, v'_j) of $\Delta_n \times \Delta_d$ in lexicographical order. [21]

These nice properties of simplicial products extend in the following, more general, theorem.

Theorem 3.2.3. *[21] If P and P' are lattice polytopes such that each admits a (regular) unimodular triangulation, then $P \times P'$ also has a (regular) unimodular triangulation. The set of minimal non-faces will be the lifts of the minimal non-faces from the triangulations of P and P', together with additional non-faces of cardinality two.*

The first sentence tells us that a unimodular triangulation will "lift" to a unimodular triangulation in the product, and that if the original's were regular, the product's will be as well. The second sentence guarantees that no new non-faces of cardinality greater than two will appear in the product. As a result, if we have quadratic triangulations of P and P', we can construct one in the product $P \times P'$ as well.

Proof. The product of a triangulation is a subdivision into a product of simplices; if each of these simplices in the product is unimodular, then any refining triangulation of the subdivision will be unimodular. Further, the non-face condition follows from ordering the lattice points of each of P and P' completely, $v_1 \prec ... \prec p_r$ and $p'_1 \prec ... \prec p'_s$, and pulling the lattice points

(p_i, p'_j) in $P \times P'$ lexicographically. One can then consider a non-face N: either it was a non-face in P or P', or it is a non-face in the staircase triangulation. $\square$

This result is a powerful tool in the search for quadratic triangulations of lattice polytopes discussed in Chapter 1. However, as mentioned earlier, these products of simplices also have direct ties to tropical oriented matroids, so I now return to discussion of their role in my tropical work.

3.3 Tropical Oriented Matroids Meet $\Delta_{n-1} \times \Delta_{d-1}$

Armed with some background on $\Delta_{n-1} \times \Delta_{d-1}$, we turn our attention again tropical oriented matroids. Oh and Yoo initially claimed that a collection of (n, d)-types is a tropical oriented matroid if and only if it corresponds to a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$. However, they have since revised this statement, offering instead only the backwards direction of the bijection.

Theorem 3.3.1. *[30] A collection of all trees in a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$ forms a tropical oriented matroid.*

Oh and Yoo's proof uses, and is phrased with respect to, the bijection between products of the form $\Delta_{n-1} \times \Delta_{d-1}$ and the complete bipartite graph $k_{n,d}$, and appears in [30]. The fact that this statement is not a bijection corresponds to the fact that non-generic tropical oriented matroids correspond not to triangulations of $\Delta_{n-1} \times \Delta_{d-1}$, but to other subdivisions.

Theorem 3.3.2. *Non-generic tropical oriented matroids do not correspond to triangulations of $\Delta_{n-1} \times \Delta_{d-1}$.*

The translation between triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ and tropical oriented matroids is a map between simplices of the triangulation and vertices of the matroid. The proof of Theorem 3.3.2 shows that this correspondence fails in the case of non-generic matroids.

Proof. To show that vertices of a non-generic tropical oriented matroid do not correspond to the simplices of a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$, recall the grid notation for a product of $\Delta_{n-1} \times \Delta_{d-1}$, and a triangulation thereof. In particular, a simplex in this grid notation is represented by a grid with exactly $n + d - 1$ entries. Translating the (n, d)-types of the vertices of a tropical oriented matroid into this grid notation shows that a non-generic vertex will have at least $n + d$ entries and hence not correspond to a simplex in the map taking grids to grids, which provides the translation between simplices of triangulations and vertices of triangulations in the generic case. [2] The grid representation of a type A in a tropical oriented matroid of type-(n, d) is defined as follows. There is an entry in the i^{th} position of the j^{th} row of the grid if $i \in A_j$. Now consider a non-generic tropical oriented matroid M and take B to be a non-generic vertex of M. By definition, B has at least $n + d$ entries in its grid representation. This means B can not correspond to a simplex of a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$. Therefore, the vertices of M do not correspond to a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$.

□

Oh and Yoo's original proof claiming a bijection between tropical oriented matroids and triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ holds in the generic case, and could be restated with the generic qualification as follows.

[2] An empty entry in this grid is equivalent to a 0 in the incidence matrix representation of type A, and an entry in the grid corresponds to a 1 in this matrix.

Theorem 3.3.3. *A collection of (n, d)-types is a generic tropical oriented matroid if and only if it corresponds to a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$.*

3.3.1 The Regular Case

When all triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ are regular, a model for the space of all triangulations is offered by the secondary polytope. Taken with the definition of the secondary polytope, Theorem 3.3.3 leads to the following lemma.

Lemma 3.3.4. *When all triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ are regular, generic tropical oriented matroids of type-(n, d) correspond to vertices of the secondary polytope of $\Delta_{n-1} \times \Delta_{d-1}$.*

The bijection between generic tropical oriented matroids of type-(n, d) and products of simplices $\Delta_{n-1} \times \Delta_{d-1}$ has ties to non-generic tropical oriented matroids as well. In particular, the fact that tropical oriented matroids are in bijection with triangulations of products of simplices means that non-generic tropical oriented matroids do not correspond to triangulations of products of simplices. Such non-generic matroids exist instead as the limit of two or more such triangulations— in the case of two, these are triangulations separated by a "flip." The non-generic nature of the arrangement corresponds to the fact that the flip remains unspecified.

Combining Theorem 3.3.3 and Lemma 3.3.4 with the relationship between generic and non-generic tropical oriented matroids, we obtain the following statement:

Theorem 3.3.5. *When all triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ are regular, non-generic type-(n, d) tropical oriented matroids correspond to faces of dim > 0 of the secondary polytope of $\Delta_{n-1} \times \Delta_{d-1}$.*

Proof. Regular triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ correspond to vertices of the secondary polytope of $\Delta_{n-1} \times \Delta_{d-1}$. As noted above, when all triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ are regular, this means that tropical oriented matroids of type-(n, d) also correspond to vertices of the secondary polytope. Edges between vertices in this polytope correspond to flips between triangulations. In particular, a subdivision that does not make a choice of flip is only a subdivision, not a triangulation, corresponding to an edge (or higher dimensional face of the polytope, if it fails to make a choice on multiple flips). A generic tropical oriented matroid M corresponds to a triangulation of $\Delta_{n-1} \times \Delta_{d-1}$, and hence a vertex of the secondary polytope of $\Delta_{n-1} \times \Delta_{d-1}$. I abuse notation by also labeling this M. Consider an apex A of this arrangement changing relative position with respect to some other position, i. Without loss of generality, say it moves from sector k to sector l (taking you from A with $A_i = \{k\}$ to A' with $A'_i = \{l\}$). This move results in a new tropical oriented matroid, corresponding to a different vertex of the secondary polytope of $\Delta_{n-1} \times \Delta_{d-1}$, call it M'. The elimination axiom guarantees the existence of some $\tilde{A}$ such that $\tilde{A}_i = \{k, l\}$. The tropical oriented matroid containing this $\tilde{A}$ is a non-generic tropical oriented matroid, and corresponds to the face of the secondary polytope containing M and M'.

$\square$

In [11], Santos explains that the only cases of $\Delta_{n-1} \times \Delta_{d-1}$ for which n and d are > 1 and all triangulations are regular are: $\Delta_2 \times \Delta_2$, $\Delta_3 \times \Delta_2$, $\Delta_4 \times \Delta_2$, $\Delta_2 \times \Delta_3$, and $\Delta_2 \times \Delta_4$. The reason this result does not hold for other $\Delta_{n-1} \times \Delta_{d-1}$ is that non-regular triangulations are not represented by the secondary polytope. It remains true that any non-generic tropical hyperplane arrangement corresponds to a subdivision "between" triangulations, in particular the triangulations corresponding to its neighboring generic tropical oriented matroids. This statement does suggest that there is a space of all possible tropical oriented matroids of a given size and that the non-generic ones represent the boundaries separating the generic ones. This is the more general case that is modeled here by the secondary polytope. These preliminary results demonstrate not only an important distinction between the generic and non-generic tropical oriented matroids, but also that the non-generic case merits further investigation.

3.3.2 The Non-Regular Case

The case of non-regular subdivisions and triangulations merits some independent discussion. Personal communication with Develin suggests that non-regular subdivisions of $\Delta_{n-1} \times \Delta_{d-1}$ are in bijection with what he and Ardila defined as tropical pseudohyperplane arrangements. In their paper, Ardila and Develin define *a tropical pseudohyperplane arrangement* as a subset of $\mathbb{TP}^{d-1}$ that is PL-homeomorphic to a tropical hyperplane, and offer the following conjecture of a "tropical representation theorem."

Conjecture 2. *[3]*

Every tropical oriented matroid can be realized by an arrangement of tropical pseudohyperplanes.

The idea here is to use the Cayley trick that puts triangulations of products of simplices of $\Delta_{n-1} \times \Delta_{d-1}$ in bijection with mixed subdvisions of the dilated simplex $n\Delta_{d-1}$. The details of this bijection are outlined in [37]. For example, one can consider the tiling of $\Delta_1 \times \Delta_2$ depicted on the left in Figure 3.3.2. The tetrahedra in this subdivision have types $(123, 1)$, $(23, 13)$, and $(2, 123)$, which are mapped under the bijection to the Minkowski sums

$123 + 1$, $23 + 13$, and $2 + 123$. This corresponds to the mixed subdivision of $2\Delta_2$, seen on the right in the same figure. (This example corresponds to the actual tropical hyperplane arrangement seen in Figure 3.3.2.)

Obtaining a tropical pseudo (or actual) hyperplane arrangement from a subdivision of $n\Delta_{d-1}$ requires taking the combinatorial dual of each chamber of the mixed subdivsion. In the 2-dimensional case there are four possible chambers in the mixed subdivision, 3 rhombuses and 1 triangle. The geometric realization of their duals are depicted in Figure 3.3.2. [3] The lower dimensional faces of this subdivision fit together to form a tropical pseudo-hyperplane arrangement. The simplices in the mixed subdivision (i.e. the triangles in the 2-dimensional examples seen here) correspond to the apices of the pseudo-arrangement, while the other cells indicate how these pseudo-hyperplanes interact throughout the arrangement.

This process allows one to visualize tropical pseudohyperplane arrangements corresponding to triangulations of $\Delta_{n-1} \times \Delta_{d-1}$ which are not regular. An example of this is shown in Figure 3.3.2. It shows a tropical pseudohyperplane arrangement corresponding to a non-regular triangulation of $\Delta_2 \times \Delta_5$, and hence a tropical oriented matroid of type-$(3, 6)$ that should not be realizable by any actual tropical hyperplane arrangement.

Understanding the role of these pseudo arrangements in the space of all tropical oriented matroids is among my current projects. It offers another application of computer code. I could compare the output of code enumerating all tropical oriented matroids of a given (n, d)-type to existing lists of all tropical hyperplane arrangements with parameters n, d. This should help

[3] These geometric realizations are known as their mixed Voronoi subdivions. The *mixed Voronoi subdivision* of a k-simplex divides it into k regions, where regions are defined by associating each point to the nearest vertex of the simplex.

me see how pseudo arrangements fit between actual tropical hyperplane arrangements, allowing me to better model the space of all tropical oriented matroids.

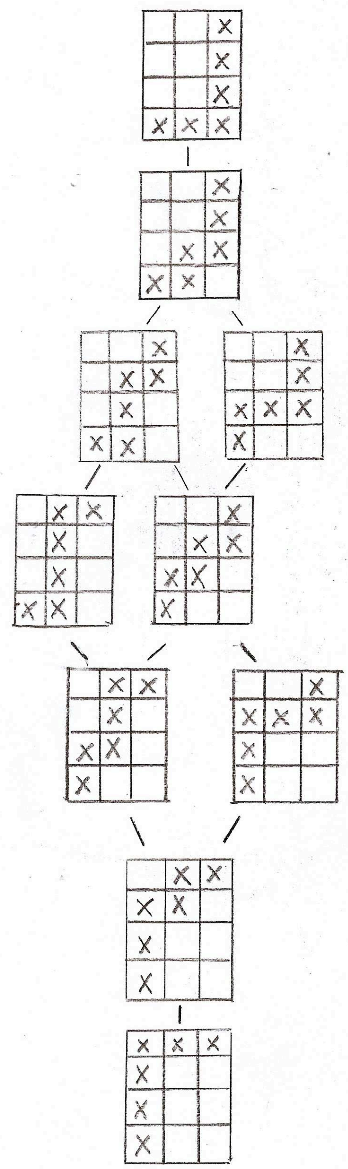

Figure 3.2.4: This depicts the collection of simplices, each given by the grid indicating which vertices it contains, comprising the monotone staircase triangulation of $\Delta_3 \times \Delta_4$. (Note that covering relationships in the poset, implied by this visual, correspond to adjacency of simplices in the triangulation.)

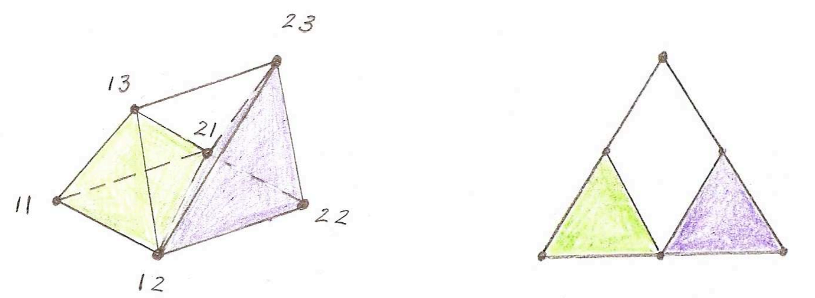

Figure 3.3.1: The image on the left shows a triangulation of $\Delta_1 \times \Delta_2$. Tetrahedron $(123, 1)$ is shown in green. Tetrahedron $(2, 123)$ is shown in purple. The tetrahedron $(23, 13)$ is the convex hull of the edge contained in both and the edge contained in neither. The representation of the same triangulation via a mixed subdivision of $2\Delta_2$ is on the right. The cells are color coded to match their corresponding tetrahedra in the image on the left.

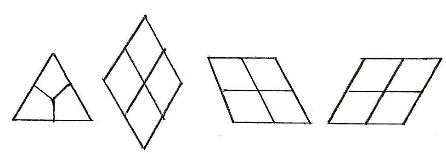

Figure 3.3.2: This shows the 4 possible cells in a mixed subdivision of $2\Delta_2$, 3 rhombuses and 1 triangle.

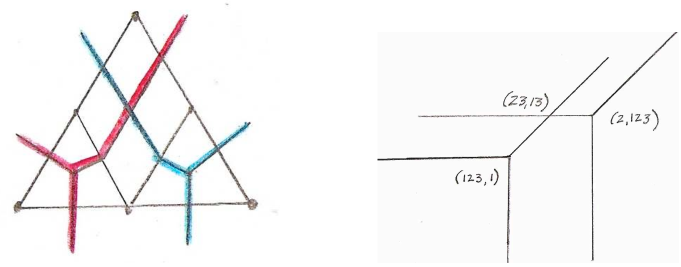

Figure 3.3.3: The image on the left shows the geometric dual of the mixed subdivision of $2\Delta_2$ corresponding to the triangulation in Figure 3.3.2. The edges indicating the faces of the 2 tropical pseudohyperplanes in the corresponding arrangement of tropical pseudohyperplanes are highlighted in red and blue. There is an arrangement of actual tropical hyperplanes with the same types (and hence representing the same tropical oriented matroid), on the right.

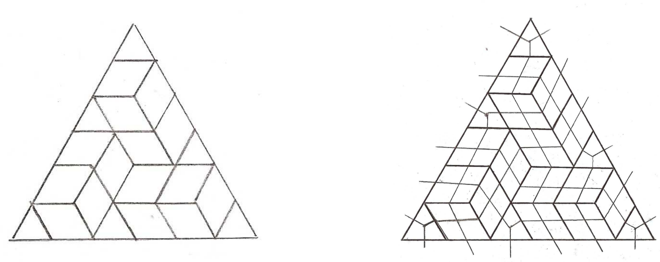

Figure 3.3.4: The image on the left depicts a mixed subdivision of $6\Delta_2$ corresponding to a non-regular triangulation of $\Delta_2 \times \Delta_5$. The image on the right shows the corresponding arrangement of 6 tropical pseudohyperplanes in the plane. This corresponds to a tropical oriented matroid of type-$(3, 6)$.

Bibliography

[1] Federico Ardila, Matthias Beck, Serkan Hosten, Julian Pfeifle, and Kim Seashore. Root polytopes and growth series of root lattices. 2008. [Cited on page 83]

[2] Federico Ardila and Sara Billey. Flag arrangements and triangulations of products of simplices. *Adv. Math.*, 214(2):495–524, 2007. [Cited on pages 45 and 83]

[3] Federico Ardila and Mike Develin. Tropical hyperplane arrangements and oriented matroids. *Math. Z.*, 262(4):795–816, 2009. [Cited on pages 66, 71, 74, 76, 82, and 96]

[4] Federico Ardila and Caroline J. Klivans. The Bergman complex of a matroid and phylogenetic trees. *J. Combin. Theory Ser. B*, 96(1):38–49, 2006. [Cited on page 45]

[5] E. K. Babson and L. J. Billera. The geometry of products of minors. *Discrete Comput. Geom.*, 20(2):231–249, 1998. [Cited on page 83]

[6] Alexander Barvinok. *A course in convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002. [Cited on page 8]

[7] Victor V. Batyrev. On the classification of toric fano 4-folds. *J. Math. Sci. (New York)*, 94(1):10211050, 1999. [Cited on page 24]

[8] Margaret M. Bayer. Equidecomposable and weakly neighborly polytopes. *Israel J. Math.*, 81(3):301–320, 1993. [Cited on page 83]

[9] Winfried Bruns and Joseph Gubeladze. *Polytopes, Rings, and K-Theory.* Monographs in Mathematics. Springer-Verlag, 2009. XIV, 461 p. 52 illus. [Cited on page 2]

[10] Winfried Bruns and Joseph Gubeladze. *Polytopes, Rings, and K-Theory.* Springer Monographs in Mathematics part 1. Springer., New York, NY, 2009. [Cited on page 36]

[11] Jesús A. De Loera, Rambau Jörg, and Francisco Santos. *Triangulations Structures for Algorithms and Applications.* Algorithms and Computation in Mathematics, Vol. 25. Springer., Boston, MA, 2010. [Cited on pages 87, 89, 90, and 95]

[12] Mike Develin. The moduli space of n tropically collinear points in $\mathbb{R}^d$. *Collect. Math.*, 56(1):1–19, 2005. [Cited on page 45]

[13] Mike Develin and Bernd Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27 (electronic), 2004. [Cited on pages 45, 53, and 54]

[14] Mike Develin and Josephine Yu. Tropical polytopes and cellular resolutions. *Experiment. Math.*, 16(3):277–291, 2007. [Cited on page 45]

[15] Robert M. Erdahl and Sergej S. Ryshkov. On lattice dicing. *Eur. J. Comb.*, 15(5):459–481, 1994. [Cited on page 18]

[16] Günter Ewald. *Combinatorial Convexity and Algebraic Geometry*, volume 168 of *GTM*. Springer–Verlag, 1996. [Cited on page 4]

[17] R.T. Firla and G.M. Ziegler. Hilbert bases, unimodular triangulations, and binary covers of rational polyhedral cones. *Discrete Comput. Geom.*, 21(2):205–216, 1999. [Cited on page 33]

[18] R.T. Firla and G.M. Ziegler. Hilbert bases, unimodular triangulations, and binary covers of rational polyhedral cones. *Discrete Comput. Geom.*, 21(2):205–216, 1999. [Cited on page 41]

[19] William Fulton. *Introduction to Toric Varieties*, volume 131 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1993. The William H. Roever Lectures in Geometry. [Cited on page 4]

[20] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, resultants and multidimensional determinants*. Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, 2008. Reprint of the 1994 edition. [Cited on page 83]

[21] Christian Haase, Andreas Paffenholx, Lindsay C Piechnik, and Francisco Santos. Existence of unimodular triangulations — positive results. 2011. [Cited on pages 12, 14, 17, 18, 20, 21, 22, 25, 27, 30, 31, 35, 37, 39, 40, and 91]

[22] M.M. Kapranov, Bernd Sturmfels, and Andrei V. Zelevinsky. Chow polytopes and general resultants. *Duke Math. J.*, 67(1):189–218, 1992. [Cited on page 28]

[23] George R. Kempf, Finn F. Knudsen, David Mumford, and Bernard Saint-Donat. *Toroidal Embeddings I*, volume 339 of *Lecture Notes in Mathematics*. Springer–Verlag, 1973. [Cited on pages 34 and 36]

[24] Carl W. Lee. Subdivisions and triangulations of polytopes. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 271–290. CRC–Press, New York, 1997. [Cited on pages 9 and 90]

[25] Grigory Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *J. Amer. Math. Soc.*, 18(2):313–377, 2005. [Cited on page 45]

[26] Grigory Mikhalkin. Tropical geometry and its applications. In *International Congress of Mathematicians. Vol. II*, pages 827–852. Eur. Math. Soc., Zürich, 2006. [Cited on pages 44 and 45]

[27] David R. Morrison and Glenn Stevens. Terminal quotient singularities in dimension 3 and 4. *Proc. A.M.S.*, 90:15–20, 1984. [Cited on page 38]

[28] Mikkel Øbro. An algorithm for the classification of smooth fano polytopes. preprint, arXiv:0704.0049, 2007. [Cited on page 24]

[29] Tadao Oda. *Convex Bodies and Algebraic Geometry. An Introduction to the Theory of Toric Varieties*, volume 15 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer–Verlag, 1988. [Cited on page 4]

[30] Suho Oh and Hwanchul Yoo. Triangulations of $\delta_{n-1} \times \delta_{d-1}$ and tropical oriented matroids. 2010. [Cited on page 92]

[31] Hidefumi Ohsugi and Takayuki Hibi. Convex polytopes all of whose reverse lexicographic initial ideals are squarefree. *Proc. Am. Math. Soc.*, 129(9):2541–2546, 2001. [Cited on page 15]

[32] Lior Pachter and Bernd Sturmfels. Tropical geometry of statistical models. *Proc. Natl. Acad. Sci. USA*, 101(46):16132–16137 (electronic), 2004. [Cited on page 45]

[33] Julian Pfeifle. Dissections, `hom`-complexes and the Cayley trick. *J. Combin. Theory Ser. A*, 114(3):483–504, 2007. [Cited on page 83]

[34] Alexander Postnikov. Permutohedra, associahedra, and beyond. *Int. Math. Res. Not. IMRN*, (6):1026–1106, 2009. [Cited on page 83]

[35] John E. Reeve. On the volume of lattice polyhedra. *Proc. London Math. Soc.*, 7:378–395, 1957. [Cited on page 9]

[36] Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald. First steps in tropical geometry. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 2005. [Cited on pages 44 and 45]

[37] Francisco Santos. The Cayley trick and triangulations of products of simplices. In *Integer points in polyhedra—geometry, number theory, algebra, optimization*, volume 374 of *Contemp. Math.*, pages 151–177. Amer. Math. Soc., Providence, RI, 2005. [Cited on page 96]

[38] Herbert E. Scarf. Integral polyhedra in three space. *Math. Oper. Res.*, 10:403–438, 1985. [Cited on page 38]

[39] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley, Chichester, NY, 1986. [Cited on page 32]

[40] David Speyer and Bernd Sturmfels. The tropical Grassmannian. *Adv. Geom.*, 4(3):389–411, 2004. [Cited on page 45]

[41] David Speyer and Bernd Sturmfels. Tropical mathematics. 2004. [Cited on pages 44 and 45]

[42] Richard P. Stanley. Decompositions of rational convex polytopes. *Ann. Discrete Math.*, 6:333–342, 1980. [Cited on page 14]

[43] Bernd Sturmfels. *Gröbner bases and convex polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1996. [Cited on pages 6, 7, and 83]

[44] Bernd Sturmfels. *Gröbner Bases and Convex Polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1996. [Cited on pages 26, 28, and 29]

[45] Seth Sullivant. Compressed polytopes and statistical disclosure limitation. Preprint `arXiv:math.CO/0412535`, 2004. [Cited on page 15]

[46] George K. White. Lattice tetrahedra. *Canadian J. Math.*, 16:389–396, 1964. [Cited on page 38]

[47] Debbie S. Yuster. *Triangulations, tropical geometry, and applications*. ProQuest LLC, Ann Arbor, MI, 2007. Thesis (Ph.D.)–Columbia University. [Cited on page 45]

[48] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. [Cited on pages 8 and 53]