TROPICAL ALGEBRAIC GEOMETRY
THE STRUCTURE THEOREM

ISABEL LONGBOTTOM

Abstract. We develop the language and conceptual framework to state the Structure Theorem for Tropical Algebraic Varieties. This includes a background in field valuations and polyhedral geometry. Focusing on examples to illustrate the geometric intuition behind various algebraic definitions, we outline all aspects of the Structure Theorem except the precise details of the weightings needed. We conclude with a broad discussion of parts of the proof of the Structure Theorem.

1. Introduction

In algebraic geometry, the main objects of study are algebraic varieties. These are the vanishing sets in the affine space k^n of polynomial ideals in the ring $k[x_1, \ldots, x_n]$ over some field k, usually algebraically closed. In tropical algebraic geometry, however, we study tropical algebraic varieties. We still work over an algebraically closed field k, but we now consider subvarieties of the algebraic torus $(k^\times)^n$ and ideals in the Laurent polynomial ring $k[x_1^{\pm 1}, x_2^{\pm 1}, \ldots, x_n^{\pm 1}]$. Any ideal I in this Laurent polynomial ring defines a vanishing set $V(I) \subset (k^\times)^n$, and a tropical algebraic variety is the tropicalisation of such a set, which is a collection of polyhedra in $\mathbb{R}^n$. To define the tropicalization, we will use a new kind of arithmetic, in which multiplication is replaced by addition and addition is replaced by minima.

To define a tropical algebraic variety, we take a function which is the minimum of finitely many affine pieces, and then look at the set of points in the domain for which the function fails to be smooth. This is the tropical hyperplane corresponding to some Laurent polynomial, which depends on the piecewise-affine function we started with.

Our focus is the Structure Theorem for Tropical Algebraic Varieties. To give an intuitive understanding of the statement of this theorem, consider the following. Every tropical algebraic variety is a polyhedral complex (informally, a collection of polyhedra, which are what you would expect except that we allow them to be unbounded in this context), but the converse is not true. To study the converse direction, one wishes to find a criterion to determine whether a given polyhedral complex is a tropical variety. The Structure Theorem gives a partial solution, in that it describes the properties of tropical varieties corresponding to irreducible subvarieties of the algebraic torus.

Slightly more formally, the Structure Theorem states that the tropicalisation of an irreducible subvariety of the torus has the same dimension as the subvariety while satisfying balancing and connectedness conditions.

Moreover, the coefficients of the affine equations describing the polyhedra comprising the tropical variety lie in a specific subgroup of $\mathbb{R}^n$.

We will need a strong base of theory from polyhedral geometry to state the Structure Theorem, although since we are focusing on those definitions relevant to our main goal, this background will lack breadth. Where the reader desires more detail on an aspect of Section 4 in particular, [3] is recommended as an excellent reference.

We will proceed mainly via definitions and examples. The material summarised here is largely based on [2], and many of the notational choices come from this source.

2. Basic Definitions: Tropical Arithmetic and Valuations

The *tropical semiring* $(\mathbb{R} \cup \{\infty\}, \oplus, \odot)$ has elements consisting of the real numbers and one additional value ∞, representing positive infinity. The arithmetic operations in the tropical semiring are

$$x \odot y := x + y \text{ and } x \oplus y := \min(x, y).$$

These operations satisfy all the usual ring axioms, including distributivity of $\odot$ over $\oplus$, except that $\oplus$ does not have inverses. Therefore these operations make the set $\mathbb{R} \cup \{\infty\}$ into a (commutative) semiring. We note that because we are treating ∞ as positive infinity, we have

$$x \oplus \infty = x \text{ and } x \odot \infty = \infty$$

for every x. The identity element for $\odot$ is 0 and the identity element for $\oplus$ is ∞. We use exponent notation to denote repeated tropical multiplication, so $a^3 = a \odot a \odot a$. This is sometimes also denoted $a^{\odot 3}$.

This gives us a new notion of arithmetic on $\mathbb{R}$, but we really want a new notion of arithmetic on k. To this end, we define a *valuation* from k to the tropical semiring.

Definition 2.1. A valuation on a field k is a map $\mathrm{val} : k \to \mathbb{R} \cup \{\infty\}$ satisfying the following for all $a, b \in k$:

(1) $\mathrm{val}(ab) = \mathrm{val}(a) + \mathrm{val}(b) = \mathrm{val}(a) \odot \mathrm{val}(b)$,
(2) $\mathrm{val}(a + b) \geq \min\{\mathrm{val}(a), \mathrm{val}(b)\} = \mathrm{val}(a) \oplus \mathrm{val}(b)$, and
(3) $\mathrm{val}(a) = \infty$ if and only if $a = 0$.

We often work with the restriction of a valuation to the multiplicative group $k^\times$. Then its image is an additive subgroup of $\mathbb{R}$, called the *value group* and denoted Γ_{val}. A field with valuation is called a *valued field.*

The function that is constantly zero on $k \setminus \{0\}$ and that evaluates to ∞ at 0 is a valuation on any field k. To avoid such trivialities, we will assume from now on that $1 \in \Gamma_{\mathrm{val}}$.

Lemma 2.2. *Let k be algebraically closed with valuation, and $1 \in \Gamma_{\mathrm{val}}$. Then Γ_{val} is dense in $\mathbb{R}$.*

Proof. Since k is algebraically closed, $a^{1/n} \in k^\times$ for every $a \in k^\times$. Then by property (1),

$$\mathrm{val}(a^{1/n}) = \frac{1}{n}\,\mathrm{val}(a).$$

Since $1 \in \Gamma_{\text{val}}$, this means that $\mathbb{Q} \subseteq \Gamma_{\text{val}}$, and because $\mathbb{Q}$ is dense in $\mathbb{R}$, Γ_{val} must also be dense in $\mathbb{R}$. $\square$

Example 2.3. (Pascal's triangle in tropical arithmetic) The entries in Pascal's triangle come from the coefficients of the monomials in the expansion of $(x \oplus y)^{\odot n}$. First, note that in the tropical semiring

$$a \oplus a = \min\{a, a\} = a$$

for any a. Therefore

$$\begin{aligned}(x \oplus y)^{\odot n} &= x^{\odot n} \oplus x^{\odot n-1} \odot y \oplus \ldots \oplus y^{\odot n} \\ &= \bigoplus_{i=0}^{n} x^{\odot i} \odot y^{\odot (n-i)}\end{aligned}$$

Since 0 is the identity for $\odot$, all the coefficients in this expansion are 0 and so Pascal's triangle in the tropical semiring contains only zeroes.

In fact, we can simplify further: note that

$$x^{\odot i} \odot y^{\odot j} = ix + jy \geq \min\{(i+j)x, (i+j)y\} = x^{\odot (i+j)} \oplus y^{\odot (i+j)}$$

and so the cross terms in the expansion of $(x \oplus y)^{\odot n}$ are unnecessary. Therefore

$$(x \oplus y)^{\odot n} = x^{\odot n} \oplus y^{\odot n}.$$

Hence we could instead write Pascal's triangle with all the external entries (the first and last entries in each row) as 0, and the internal entries equal to ∞. We would usually use the version of Pascal's triangle not containing ∞ since then the entries lie in the value group for any choice of valuation.

Lemma 2.4. *If $a, b \in k$ and $\text{val}(a) \neq \text{val}(b)$ then*

$$\text{val}(a + b) = \min\{\text{val}(a), \text{val}(b)\} = \text{val}(a) \oplus \text{val}(b).$$

This gives us a condition under which valuations distribute across addition.

Proof. We have the following by property (1). Since $1^2 = 1$, $\text{val}(1) = 0$. Then $(-1)^2 = 1$, so $\text{val}(-1) = 0$ also. Hence $\text{val}(-b) = \text{val}(b)$ for every $b \in k$.

Now assume without loss of generality that $\text{val}(b) > \text{val}(a)$. Then (2) gives

$$\text{val}(a) \geq \min\{\text{val}(a+b), \text{val}(-b)\} = \min\{\text{val}(a+b), \text{val}(b)\} = \text{val}(a+b).$$

The last equality comes from the fact that we know $\text{val}(b) > \text{val}(a)$, so it cannot be true that $\min\{\text{val}(a+b), \text{val}(b)\} = \text{val}(b)$ since this would imply $\text{val}(a) \geq \text{val}(b)$. Property (2) also gives

$$\text{val}(a+b) \geq \min\{\text{val}(a), \text{val}(b)\} = \text{val}(a)$$

and so $\text{val}(a+b) = \text{val}(a) = \min\{\text{val}(a), \text{val}(b)\}$ as required. $\square$

The *Laurent polynomial ring in n variables* $k[x_1^{\pm}, \ldots, x_n^{\pm}]$ has a monomial basis consisting of elements of the form $\prod_{i=1}^{n} x_i^{a_i}$ where each $a_i \in \mathbb{Z}$. A Laurent polynomial in this ring can be evaluated at a point $p \in (k^{\times})^n$, the algebraic torus of rank n.

Definition 2.5. Let k be a field with valuation, and $f \in k[x]$ a univariate polynomial. To define the *tropicalisation* of f, denoted $\mathrm{trop}(f)$, we take the valuation of each of the coefficients of f and interpret multiplication and addition of terms in the polynomial f as the corresponding operations over the tropical semiring.

If $f \in k[x_1^{\pm}, \dots, x_n^{\pm}]$ is instead a Laurent polynomial in n variables, then $\mathrm{trop}(f)$ is again given by taking valuations of the coefficients of f and interpreting arithmetic over the tropical semiring. In this setting $\mathrm{trop}(f)$ is a function defined on $\mathbb{R}^n$ as a subset of n-fold product of the tropical semiring. Instead of trying to make sense of negative exponents and ∞, we exclude points with at least one coordinate equal to ∞. Since $\Gamma_{\mathrm{val}}^n \subseteq \mathbb{R}^n$, we are not losing anything by doing so. More explicitly, if f is given by

$$f = \sum_{\mathbf{w} \in \mathbb{Z}^n} c_{\mathbf{w}} x^{\mathbf{w}}$$

then its tropicalisation is

$$\mathrm{trop}(f)(\mathbf{u}) = \min_{\mathbf{w} \in \mathbb{Z}^n} \left\{ \mathrm{val}(c_{\mathbf{w}}) + \sum_{i=1}^{n} u_i w_i \right\} = \min_{\mathbf{w} \in \mathbb{Z}^n} \{ \mathrm{val}(c_{\mathbf{w}}) + \mathbf{u} \cdot \mathbf{w} \}$$

defined for $\mathbf{u} \in \mathbb{R}^n$. Note that if f happens to be a polynomial (has no negative exponents) then $\mathrm{trop}(f)$ is well-defined on the n-fold product of the tropical semiring.

Usually we would like k to be an algebraically closed field of characteristic zero, but the definition still makes sense for an arbitrary field.

Example 2.6. Take $k = \mathbb{Q}$, and let p be a prime. Consider the *p-adic valuation*, which is defined for $a \in \mathbb{Z}$ by

$$\mathrm{val}(a) = \max\{n \in \mathbb{N} \ : \ p^n \mid a\}$$

and extended to $\mathbb{Q}$ by setting $\mathrm{val}(a/b) = \mathrm{val}(a) - \mathrm{val}(b)$. This valuation measures divisibility by p. Then $\mathrm{val}(q) \in \mathbb{Z}$ for every $q \in \mathbb{Q}^{\times}$, and we take $\mathrm{val}(0) = \infty$. This valuation has value group $\mathbb{Z}$, and for a polynomial $f \in \mathbb{Q}[x]$,

$$f = a_n x^n + a_{n-1} x^{n-1} + \dots + a_0$$

we get

$$\begin{aligned} \mathrm{trop}(f) &= \mathrm{val}(a_n) \odot x^{\odot n} \oplus \mathrm{val}(a_{n-1}) \odot x^{\odot (n-1)} \oplus \dots \oplus \mathrm{val}(a_0) \\ &= \min\{\mathrm{val}(a_n) + nx, \mathrm{val}(a_{n-1}) + (n-1)x, \dots, \mathrm{val}(a_0)\} \end{aligned}$$

In general, the tropicalisation of a polynomial in several variables is piecewise affine, written as the minimum of a collection of affine functions corresponding to the monomials of the original polynomial. The affine pieces are the *tropical monomials* of $\mathrm{trop}(f)$.

For an explicit example, consider $f(x) = 10x^3 + \frac{1}{3}x^2 - 6x + \frac{32}{3}$. Then with respect to the 2-adic valuation,

$$\mathrm{trop}(f) = 1 \odot x^{\odot 3} \oplus 0 \odot x^{\odot 2} \oplus 1 \odot x \oplus 5.$$

Evaluating $\mathrm{trop}(f)$ at some point u in the tropical semiring, we have

$$\mathrm{trop}(f)(u) = \min\{1 + 3x, 2x, 1 + x, 5\}$$

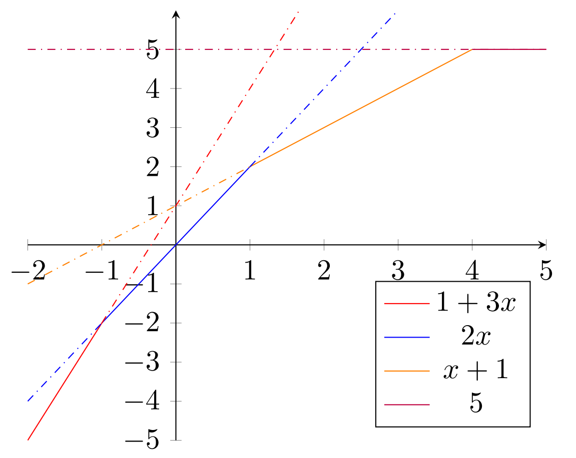

A plot of $\text{trop}(f)$ and its various affine pieces

Definition 2.7. A *tropical polynomial* is a function defined on $\mathbb{R}^n$, thought of as a subset of the n-fold product of the tropical semiring, which is the minimum of finitely many affine pieces. Such functions arise naturally as the tropicalisation of a polynomial over some valued field. We think of such polynomials as functions $\mathbb{R}^n \to \mathbb{R}$.

Every tropical polynomial is continuous, piecewise affine, and concave.

Definition 2.8. For a tropical polynomial F in n variables, $\mathbf{u} \in \mathbb{R}^n$ is a *tropical zero* of F if the minimum of the tropical monomials of F is attained at least twice at $\mathbf{u}$. With this notion of a tropical zero, we define

$$V(F) = \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \text{ is a tropical zero of } F\}$$

for a tropical polynomial F.

Returning to example 2.6, we see that $\text{trop}(f)$ has tropical zeroes $\{-1, 1, 4\}$. These are given by intersections of the tropical monomials which lie in the image of $\text{trop}(f)$. Note not every intersection of two tropical monomials corresponds to a zero.

Lemma 2.9. *Tropicalisation is compatible with polynomial evaluation. That is, if* $\mathbf{u} = \text{val}(\mathbf{c})$ *is the component-wise valuation of some* $\mathbf{c} \in (k^\times)^n$ *and* $\mathbf{u}$ *is not a tropical zero of* $\text{trop}(f)$*, then* $\text{val}(f(\mathbf{c})) = \text{trop}(f)(\mathbf{u})$*. Tropicalisation is also compatible with multiplication, that is for* $f, g \in k[x_1^\pm, \ldots, x_n^\pm]$*,*

$$\text{trop}(fg) = \text{trop}(f) \odot \text{trop}(g).$$

Proof. Since $\mathbf{u}$ is not a tropical zero, Lemma 2.4 gives the first equality. Fix $\mathbf{u} \in \Gamma_{\text{val}}^n$ not a tropical zero of $\text{trop}(f), \text{trop}(g)$ or $\text{trop}(fg)$, and choose $\mathbf{c} \in (k^\times)^n$ with $\text{val}(\mathbf{c}) = \mathbf{u}$. Then

$$\text{trop}(fg)(\mathbf{u}) = \text{val}((fg)(\mathbf{c})) = \text{val}(f(\mathbf{c})) \odot \text{val}(g(\mathbf{c})) = \text{trop}(f)(\mathbf{u}) \odot \text{trop}(g)(\mathbf{u})$$

Excluding the roots of the three tropical polynomials leaves a dense subset of Γ_{val}^n on which $\text{trop}(fg) = \text{trop}(f) \odot \text{trop}(g)$. Since Γ_{val} is dense in $\mathbb{R}$, Γ_{val}^n is also dense in $\mathbb{R}^n$, and a dense subset of Γ_{val}^n is itself dense in $\mathbb{R}^n$. Then because tropical polynomials are continuous, $\text{trop}(fg) = \text{trop}(f) \odot \text{trop}(g)$ on $\mathbb{R}^n$. □

3. The Tropical Algebraic Variety

First we consider tropical hypersurfaces.

Definition 3.1. Let $f \in k[x_1^{\pm}, \dots, x_n^{\pm}]$ be a Laurent polynomial, and let $X = V(f)$ be the hypersurface defined by f in $(k^{\times})^n$. Then the *tropical hypersurface* corresponding to f is

$$\mathrm{trop}(X) = \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \text{ is a tropical zero of } \mathrm{trop}(f)\} = V(\mathrm{trop}(f)) \subset \mathbb{R}^n.$$

This is precisely the locus in $\mathbb{R}^n$ where the piecewise affine function $\mathrm{trop}(f)$ fails to be smooth.

A tropical variety is the tropicalisation of a classical variety over a field with valuation. In our context we will start with a Laurent polynomial ideal and its subvariety of the algebraic torus and define the corresponding tropical variety in $\mathbb{R}^n$.

Definition 3.2. Suppose $I \subseteq k[x_1^{\pm}, \dots, x_n^{\pm}]$ is an ideal in the Laurent polynomial ring, with $X = V(I)$ its variety in the algebraic torus. The *tropical variety* corresponding to I is the set of common zeroes of tropicalisations of polynomials in I, that is

$$\mathrm{trop}(V(I)) = \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \text{ is a tropical zero of } \mathrm{trop}(f) \ \forall\ f \in I\}.$$

This set is also called the *tropicalisation* of X, and can be written

$$\mathrm{trop}(X) = \bigcap_{f \in I} \mathrm{trop}(V(f)) = \bigcap_{f \in I} V(\mathrm{trop}(f))$$

as the intersection of the tropical hypersurfaces defined by polynomials in I.

Remark 3.3. It can be shown that the set $\mathrm{trop}(X)$ depends only on $\sqrt{I}$, so $\mathrm{trop}(X)$ does not depend on a choice of ideal I with $V(I) = X$. In fact, the Fundamental Theorem of Tropical Algebraic Geometry tells us that $\mathrm{trop}(X)$ is the closure of the set of coordinate-wise valuations of points in X. This is why the notation for a tropical variety omits the ideal I.

Then a *tropical variety* in $\mathbb{R}^n$ is any set of the form $\mathrm{trop}(X)$ with X some subvariety of the algebraic torus $(k^{\times})^n$, k a valued field.

Remark 3.4. When defining an affine variety, it is common to take

$$V(I) = \bigcap_{i=1}^{m} V(f_i)$$

where $f_1, \dots, f_m$ is some generating set for I. In the context of tropical varieties, we might be tempted to do something similar and define

$$\mathrm{trop}(X) = \bigcap_{i=1}^{m} \mathrm{trop}(V(f_i))$$

for a generating set $f_1, \dots, f_m$ of the Laurent polynomial ideal $I(X)$. In general, this does not agree with Definition 3.2. Starting with a generating set, one must usually pass to a larger subset of I and then consider the intersection of the tropical hypersurfaces corresponding to those polynomials before the intersection is equal to $\operatorname{trop}(X)$. This is essentially because tropicalisation does not commute with taking intersections.

Significantly, this also means that not every set of the form

$$\bigcap_{i=1}^{m} \operatorname{trop}(V(f_i)) = \bigcap_{i=1}^{m} V(\operatorname{trop}(f_i))$$

is a tropical variety. Such sets are called *tropical prevarieties*. With the definition of a polyhedral complex in section 4, we see that every tropical variety is the support of some polyhedral complex but not every polyhedral complex is a tropical variety.

Example 3.5. We consider a tropical hyperplane defined by a quadratic. Consider the general tropical quadratic in 2 variables

$$p(x, y) = a \odot x^{\odot 2} \oplus b \odot x \odot y \oplus c \odot y^{\odot 2} \oplus d \odot y \oplus e \oplus f \odot x.$$

To avoid degeneracies, suppose the coefficients satisfy the following inequalities

$$e + b > f + d, f + c > d + b, 4d > c, d + a > f + b, 4f > a.$$

Then the function $p : \mathbb{R}^2 \to \mathbb{R}$ is the lower envelope of 6 planes in $\mathbb{R}^3$. The corresponding tropical quadratic is the projection onto $\mathbb{R}^2$ of the minimal intersections of these planes. It consists of 3 line segments, 6 rays and 4 vertices. By 'avoid degeneracies', all we mean here is that there are 6 distinct 2-dimensional faces, and so on. One could imagine translating one of the affine spaces defining a face in the positive z direction until its intersection with the hyperplane was empty, and thus decreasing the number of faces.

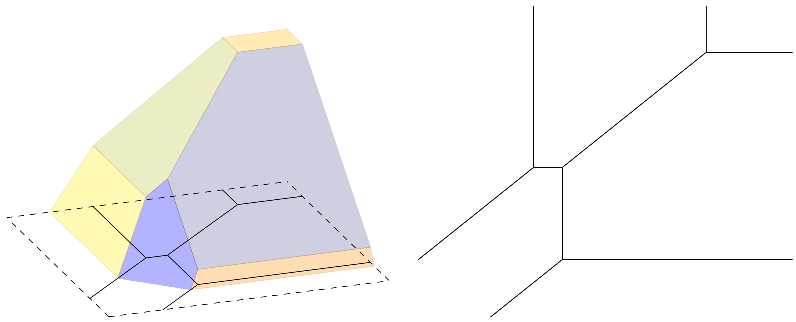

(A) A plot of $p(x, y)$ over the tropical hyperplane it defines

(B) The tropical variety

4. Polyhedra, Polyhedral Fans and Polyhedral Complexes

The next section will discuss some of the notions from polyhedral geometry that we will need to state the Structure Theorem. We will focus on examples, and proofs will generally be omitted. A reader desiring further details of any of these concepts should see [3].

Definition 4.1. A set $S \subseteq \mathbb{R}^n$ is called *convex* if for every $\mathbf{u}, \mathbf{v} \in S$ and $0 \leq \lambda \leq 1$ we have

$$\lambda\mathbf{u} + (1 - \lambda)\mathbf{v} \in S.$$

That is, the line segment joining $\mathbf{u}$ and $\mathbf{v}$ is contained in S.

Definition 4.2. A *polyhedron* P is the intersection of finitely many closed half-spaces in $\mathbb{R}^n$, that is

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}.$$

for a $d \times n$ matrix A and a vector $\mathbf{b} \in \mathbb{R}^d$. P consists of all points satisfying a system of linear inequalities. Each such linear inequality $\mathbf{a}_i \cdot \mathbf{x} \leq b_i$ for a row vector $\mathbf{a}_i$ in A defines a half-space, and P is the intersection of these.

A bounded polyhedron is called a *polytope*, and can be written as the convex hull of a finite set $U \subset \mathbb{R}^n$,

$$conv(U) = \left\{ \sum_{i=1}^{r} \lambda_i \mathbf{u}_i \;\middle|\; 0 \leq \lambda_i \leq 1, \sum_{i=1}^{r} \lambda_i = 1 \right\}$$

Note that $conv(U)$ is the smallest convex subset of $\mathbb{R}^n$ containing U.

Definition 4.3. Given a polyhedron $P \subseteq \mathbb{R}^n$, the *faces* of P are defined by linear functionals $w : \mathbb{R}^n \to \mathbb{R}$. We can represent such a functional as an inner product $w(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x}$ for a unique $\mathbf{w} \in \mathbb{R}^n$, so we can identify linear functionals with points in $\mathbb{R}^n$. Then the face of P corresponding to w is

$$\text{face}_w(P) = \{\mathbf{x} \in P \mid \mathbf{w} \cdot \mathbf{x} \leq \mathbf{w} \cdot \mathbf{y} \;\forall\; \mathbf{y} \in P\}.$$

This is the set of points in P where the value of the functional is minimised. Such a set is called a face of P if it is nonempty.

A face of P that is not contained in any larger proper face is called a *facet*.

Note that a face of a polyhedron is itself a polyhedron. It is defined by the same linear inequalities as P, plus one additional inequality,

$$\mathbf{w} \cdot \mathbf{x} \leq c$$

where c is the constant value taken by $\mathbf{w} \cdot \mathbf{x}$ on $\text{face}_w(P)$. This last linear inequality is always an equality on $\text{face}_w(P)$, and there are no points in P satisfying $\mathbf{w} \cdot \mathbf{x} < c$ because of the minimality condition.

Definition 4.4. The *affine span* of a polyhedron P is the smallest affine space containing P and is a translation of a linear subspace of $\mathbb{R}^n$. The dimension of P is defined to be the dimension of this linear subspace.

With this notion of dimension, we note that any facet of P has dimension one less than the dimension of P, and if a face F_1 of P is contained in some other face F_2 of P, $F_1 \subsetneq F_2$, then $\dim(F_1) < \dim(F_2)$. The facets of P are precisely those faces of dimension one less than the dimension of P.

Example 4.5. Consider the polyhedron $P \subseteq \mathbb{R}^2$ defined by

$$\begin{bmatrix} -4 & 1 \\ 4 & 1 \\ 0 & -1 \end{bmatrix} \mathbf{x} \leq \begin{bmatrix} 4 \\ 4 \\ 0 \end{bmatrix}$$

This is a 2-dimensional polytope which is the convex hull of the points $(0, 4), (1, 0)$ and $(-1, 0) \in \mathbb{R}^2$. These 3 points are the 0-dimensional faces of P, while the line segments joining them are the 1-dimensional faces, and also the facets of P. P is a 2-dimensional face of itself defined by the zero functional, giving a total of 7 faces.

We now turn our attention to an important special case of polyhedra.

Definition 4.6. A *polyhedral cone* $C \subseteq \mathbb{R}^n$ is a polyhedron defined by a bounding vector $\mathbf{b} = 0 \in \mathbb{R}^d$. That is, a set of the form

$$C = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq 0\}$$

for a $d \times n$ matrix A. The faces of a cone (defined in the same way as for any other polyhedron) have an alternative description as

$$\text{face}_w(C) = \{\mathbf{x} \in C \mid A'\mathbf{x} = 0\}$$

where A' is a $d' \times n$ submatrix of A depending on w. We are effectively choosing a subset of the linear inequalities defining C and forcing them to be equalities.

A cone is an intersection of linear half-spaces which include the origin, and the affine span of a cone is a linear subspace of $\mathbb{R}^n$. The faces of a cone are themselves cones. One-dimensional cones are rays from the origin.

Next we want to consider collections of polyhedra.

Definition 4.7. A *polyhedral complex* is a collection Σ of polyhedra, satisfying the following two conditions:

(i) if a polyhedron P is in Σ then so are all its faces;
(ii) for polyhedra $P, Q \in \Sigma$, their intersection $P \cap Q$ is either empty, or a face of both P and Q.

The polyhedra in a polyhedral complex are called *cells*. The *facets* of Σ are those cells which are not faces of any larger cell. The facets of these polyhedra are called *ridges* of Σ.

The set of facets of Σ defines Σ uniquely, since Σ precisely contains its facets plus all the faces of its cells.

A polyhedral complex whose facets all have the same dimension d is *pure of dimension d*.

An important special case of a polyhedral complex is a *polyhedral fan*. This is a complex whose cells are all cones.

Example 4.8. Below are some examples and non-examples of polyhedral complexes in $\mathbb{R}^2$. Only some of these are fans.

In (A) the two shaded grey regions are 2-dimensional facets, and the ray which does not border either of these regions is a 1-dimensional facet. (B) is not a polyhedral complex because the intersection of the bottom rectangular region and either of the top two regions is not a face of the bottom region. In (C), the 5 shaded grey regions are all 2-dimensional facets while the ray that does not border any of them is a 1-dimensional facet. (D) has two 2-dimensional facets, shaded grey, and four 1-dimensional ridges, the 4 rays depicted.

Definition 4.9. The *support* of a polyhedral complex Σ is its image in $\mathbb{R}^n$, which is the union of its cells. That is,

$$\text{supp}(\Sigma) = \bigcup_{P \in \Sigma} P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x} \in P \text{ for some } P \in \Sigma\}.$$

Given two polyhedral complexes Σ_1, Σ_2, we can define their intersection by taking the nonempty polyhedra $P \cap Q$ with $P \in \Sigma_1$, $Q \in \Sigma_2$.

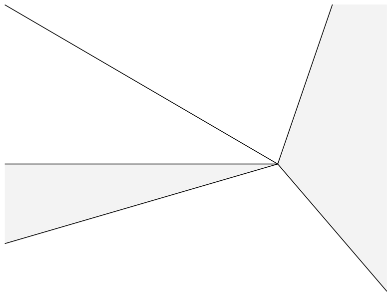

(A) A fan with 3 facets, not all of the same dimension

(B) Not a polyhedral complex

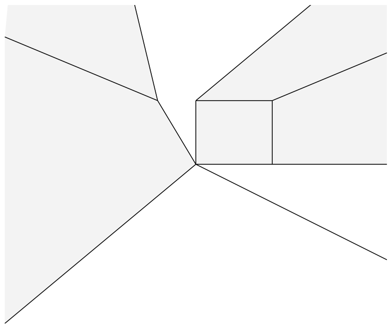

(C) A non-fan polyhedral complex with 6 facets

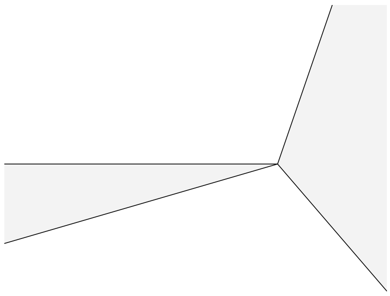

(D) A pure dimension 2 fan

This intersects their supports, so

$$\mathrm{supp}(\Sigma_1 \cap \Sigma_2) = \mathrm{supp}(\Sigma_1) \cap \mathrm{supp}(\Sigma_2).$$

Example 4.10. We can think of an octahedron in $\mathbb{R}^3$ as a polyhedral complex in two different ways. If we want the octahedron to be hollow, then the 8 classical faces of the octahedron will be the facets of the corresponding polyhedral complex, while the 12 classical edges are the ridges. The polyhedral complex will have a total of $8 + 12 + 6 = 26$ cells.

Alternatively, we could build a polyhedral complex representing a solid octahedron. This complex has one 3-dimensional facet, which contains all the points in the octahedron, and 8 2-dimensional ridges which were the facets in our last construction. It has a total of 27 cells. Because this polyhedral complex has only one facet, it is really a polyhedron. In fact this is a polytope because it is bounded.

We now define a useful notion of connectedness for polyhedral complexes.

Definition 4.11. Suppose Σ is a pure d-dimensional polyhedral complex in $\mathbb{R}^n$, some $d \leq n$. We say that Σ is *connected through codimension one* if there is a path along facets and ridges between any two facets of Σ. That is, for any two d-dimensional polyhedra $P, P' \in \Sigma$, there is a chain

$$P = P_0, P_1, P_2, \ldots, P_s = P'$$

with each $P_i \in \Sigma$ a facet, and where for each $0 \leq i \leq s-1$, $P_i \cap P_{i+1}$ is a ridge of Σ (and hence a facet of both P_i and P_{i+1}). This is a path along facets and ridges of Σ connecting P and P'.

Example 4.12. Any polyhedral complex which is connected through codimension one is path-connected in the usual topological sense. For a pure 1-dimensional complex, these notions are equivalent because any path-connected complex is connected through dimension zero.

None of the three polyhedral complexes shown above is connected through codimension one, but the example below is.

A fan, connected through codimension 1

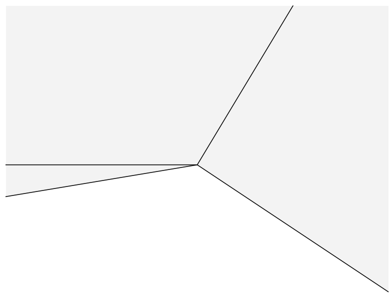

There is one further basic concept we will need regarding polyhedral complexes.

Definition 4.13. Let Γ be an additive subgroup of $\mathbb{R}$. A Γ-*rational polyhedron* is one whose defining matrix A has entries in $\mathbb{Q}$, and whose bounding vector $\mathbf{b} \in \Gamma^d$ has entries in Γ. That is,

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}$$

with $A \in \mathbb{Q}^{d \times n}$ and $\mathbf{b} \in \Gamma^d$.

A polyhedral complex is Γ-rational if all its constituent polyhedra are, or equivalently if all its facets are Γ-rational.

If $\Gamma = \mathbb{Q}$, we shorten $\mathbb{Q}$-*rational* to just *rational.*

We will, of course, be interested in the case where $\Gamma = \Gamma_{\mathrm{val}}$ is the value group.

Remark 4.14. The bounding vector of a cone is by definition the zero vector, so a cone is $\{0\}$-rational if and only if it is Γ-rational for every additive subgroup $\Gamma \subseteq \mathbb{R}$. In particular, a cone is rational if and only if it is Γ-rational for at least one such Γ.

The same applies to a fan, since all the cells of a fan are cones.

5. Weights and Balancing

Next we will meet and investigate weighted versions of pure-dimensional polyhedral complexes, and define what it means for a weighted complex to be balanced. A good place to start is with a balanced one-dimensional fan.

Definition 5.1. Let Σ be a pure d-dimensional polyhedral complex in $\mathbb{R}^n$. We give Σ the structure of a *weighted complex* by assigning a positive integer weight $m(P) \in \mathbb{N}$ to each facet $P \in \Sigma$. Recall that the facets are precisely the d-dimensional cells since Σ is pure-dimensional.

Definition 5.2. Suppose $\Sigma \in \mathbb{R}^n$ is a one-dimensional weighted rational fan consisting of s rays from the origin, denoted $P_1, \ldots, P_s$. Since each ray is rational, there are integer lattice points lying on each ray. Let $\mathbf{u}_i$ be the first such nonzero lattice point on the ray $P_i \in \Sigma$. Then the fan is *balanced* if the zero-tension condition holds, that is

$$\sum_{i=1}^{s} m(P_i) \cdot \mathbf{u}_i = 0.$$

Next we want to extend this definition to a pure d-dimensional weighted rational fan for $d > 1$, and then extend further to a general pure dimensional polyhedral complex. Before we can do this we need one more definition.

Definition 5.3. The *Minkowski sum* of two polyhedra is their pointwise vector sum, that is

$$P + Q = \{\mathbf{x} + \mathbf{y} \mid \mathbf{x} \in P \text{ and } \mathbf{y} \in Q\}$$

for polyhedra $P, Q \subseteq \mathbb{R}^n$. The Minkowski sum of two polyhedra is a polyhedron and the Minkowski sum of two cones is a cone.

We consider weighted rational fans in $\mathbb{R}^n$ of pure dimension d. Let Σ be such a fan, with weightings $m(P)$ for the cones $P \in \Sigma$ of dimension d. Fix a cone $Q \in \Sigma$ with $\dim(Q) = d - 1$. We can define a 1-dimensional fan associated to Q, and this fan will inherit weightings from the weightings of Σ. It has rays corresponding to facets $P \in \Sigma$ which have Q as a face.

Let L be the linear span of Q, so $\dim(L) = d - 1$. We consider the set of integer lattice points contained in L, $L_{\mathbb{Z}} = L \cap \mathbb{Z}^n$. Since Q is a rational cone, the linear equation defining L has rational coefficients, and so $L_{\mathbb{Z}}$ is a free abelian group of the same rank as L. Since L is a subspace of $\mathbb{R}^n$, $N_Q = \mathbb{Z}^n / L_{\mathbb{Z}}$ has a free abelian component of rank $n - d + 1$, and possibly also some additional torsion. Therefore $N_Q \otimes \mathbb{R} \cong \mathbb{R}^{n-d+1}$ since tensoring with $\mathbb{R}$ gets rid of the torsion components and we get one copy of $\mathbb{R}$ for each copy of $\mathbb{Z}$ in N_Q.

For each facet $P \in \Sigma$ that has Q as a face, the set $(P + L)/L$ can be thought of as a one-dimensional cone in $N_Q \otimes \mathbb{R}$ in the following way. The linear span of $(P + L)/L$ is 1-dimensional, so it is spanned by a single basis vector. Since P is closed under linear combinations with nonnegative coefficients, $(P + L)/L$ is also, and so the set $(P + L)/L$ is a ray.

Since P, Q are both rational, there are integer lattice points lying on this ray. Let $\mathbf{u}_P$ be the minimal nonzero such lattice point. We can do this for each facet P containing Q, and then consider

$$\sum_{P \supsetneq Q, \dim(P) = d} m(P) \cdot \mathbf{u}_P.$$

If this sum is 0 — that is, the 1-dimensional rational fan at Q we constructed is balanced — then we say that Σ is *balanced at* Q.

The fan Σ is *balanced* if it is balanced at all its ridges, or equivalently if it is balanced at all its $(d-1)$-dimensional cones.

More informally, what we are doing in the above construction is finding the line through the affine span of P that is perpendicular to Q, and taking $\mathbf{u}_P$ to be the displacement from Q of the integer lattice point on this line which lies on the same side of Q as P and is closest to Q. This is sometimes called an *inward normal vector* for Q in P, since it is normal to Q and directed into P. There is a unique line perpendicular to Q in the affine span of P because $\dim(P) = \dim(Q) + 1$.

Example 5.4. Consider the following two rational fans in $\mathbb{R}^2$. Given a ridge and its adjacent facets, we derive a 1-dimensional fan around that ridge consisting of inward normal vectors (vectors perpendicular to the ridge, pointing into a given adjacent facet). For the fan to be balanced, each of these derived 1-dimensional fans must be balanced with the inherited weightings. Hence (A) cannot be balanced, since the 1-dimensional fan we get at either of the lower 1-cells has only a single ray.

The 1-dimensional fan derived from the second example at each ridge consists of two rays pointing in precisely opposite directions, represented by a lattice point and its negation. Therefore each 1-dimensional fan is balanced if and only if the two adjacent 2-cells are weighted equally. Thus (B) will be balanced when all its 2-cells have the same weight.

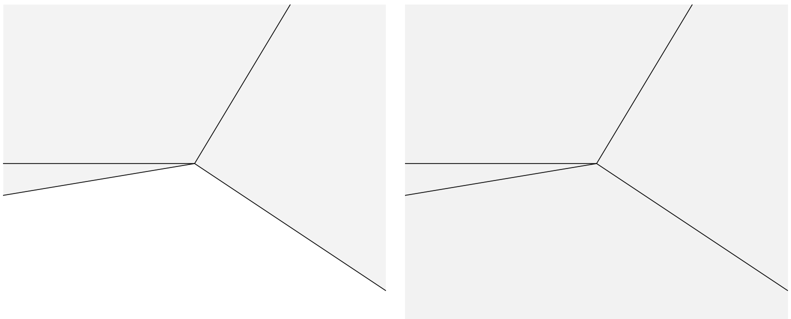

(A) Cannot be balanced

(B) Balanced as long as the 2-cells are all weighted the same

To extend this definition to apply to arbitrary polyhedral complexes, we use a similar construction to get a fan at Q for every ridge $Q \in \Sigma$ of a pure-dimensional polyhedral complex. This time the fan we derive need not be 1-dimensional.

Definition 5.5. Let Σ be a polyhedral complex in $\mathbb{R}^n$. For a cell $Q \in \Sigma$, the *star of* Q is a fan in $\mathbb{R}^n$ whose cones are indexed by the cells in Σ which contain Q as a face. For each such cell P with $Q \subsetneq P$, the cone corresponding to P in this fan is

$$\overline{P} = \{\lambda(\mathbf{x} - \mathbf{y}) \mid \lambda \geq 0, \mathbf{x} \in P, \mathbf{y} \in Q\}.$$

This is a cone because it is the positive hull of all points of the form $\mathbf{x} - \mathbf{y}$ for $\mathbf{x} \in P, \mathbf{y} \in Q$.

Then the star of Q in Σ is

$$\mathrm{star}_\Sigma(Q) = \{\overline{P} \mid P \in \Sigma, P \cap Q = Q\}$$

This is generally not a pure-dimensional fan.

At first glance this may not look like it defines a polyhedral complex, because it is not clear that $\mathrm{star}_\Sigma(Q)$ contains all the faces of its cells. It is a polyhedral complex, basically because given a cone $\overline{P} \in \mathrm{star}_\Sigma(Q)$, each of its faces is given by $\overline{P'}$ for some face P' of P which also contains Q.

Example 5.6. The cone $\overline{P}$ has the same dimension as P because the linear subspace of $\mathbb{R}^n$ spanned by the cone is a translation of the affine span of P.

In the special case where Q is a facet of Σ, the only cell of Σ containing Q is Q itself, and so $\mathrm{star}_\Sigma(Q)$ has one facet of the same dimension as Q. Hence $\mathrm{star}_\Sigma(Q)$ is essentially a polyhedral cone. Technically, $\mathrm{star}_\Sigma(Q)$ is a set containing this polyhedral cone and all its faces, but morally we can identify the fan with its single facet.

If Q is a ridge of Σ, then the only cells of Σ containing Q are Q itself, and some facets of Σ. If Σ happens to be pure-dimensional, these facets all have the same dimension, and so the cones of $\mathrm{star}_\Sigma(Q)$ are all the same dimension, except the cone corresponding to Q itself. This cone is not a facet of $\mathrm{star}_\Sigma(Q)$ because it is a face of all the others. Therefore $\mathrm{star}_\Sigma(Q)$ is also pure-dimensional.

All the stars of a Γ-rational polyhedral complex are also Γ-rational.

Definition 5.7. Let Σ be a weighted Γ_{val}-rational polyhedral complex that is pure of dimension d, and $Q \in \Sigma$ be any ridge. The facets of $\mathrm{star}_\Sigma(Q)$ are indexed by the facets of Σ which contain Q as a face, so $\mathrm{star}_\Sigma(Q)$ inherits a weighting function. The complex Σ is *balanced* if for every ridge Q, $\mathrm{star}_\Sigma(Q)$ is balanced under the inherited weights.

Remark 5.8. We only defined what it means for a rational pure-dimensional fan to be balanced, so for this definition to make sense, we need $\mathrm{star}_\Sigma(Q)$ to be rational for every ridge Q. We know that $\mathrm{star}_\Sigma(Q)$ is Γ_{val}-rational because Σ is, and then because this is a fan, it must also be rational.

Example 5.9. To get some geometric intuition for $\mathrm{star}_\Sigma(Q)$, consider the following. Let $P \in \Sigma$ have Q as a face, so that $\overline{P}$ is a cone in $\mathrm{star}_\Sigma(Q)$. Then to form the cone $\overline{P}$, we take the union of all the sets obtained by translating P so that one of the points in Q lies at the origin. Then $\overline{P}$ consists of all positive scalar multiples of points in this union.

In this sense, $\overline{P}$ is the collection of rays parametrised by the gradients of displacement vectors from Q to P.

As an explicit example, consider the complex Σ on the next page with 2 facets, P_1 and P_2, and 6 ridges. The fans $\mathrm{star}_\Sigma(Q_i)$ for $i = 1, 2, 3$ are also depicted.

In the following description, the qualifiers 'pure-dimensional' and 'weighted' have been omitted because they apply everywhere that makes sense.

To define what it means for a Γ-rational complex to be balanced, we first defined what it means for a 1-dimensional rational fan to be balanced, using the zero-tension condition. Given a d-dimensional rational fan, we can derive a 1-dimensional rational fan at each ridge by taking the inward normal vectors with respect to adjacent facets, and we already know what it means for this derived fan to be balanced. The d-dimensional fan is balanced if the derived 1-dimensional fan at each ridge is balanced. Then, given a Γ-rational complex Σ, we can derive a d-dimensional rational fan at each

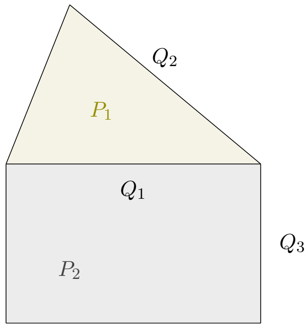

(A) The polyhedral complex Σ, with some faces labelled

(B) The star at Q_3 in Σ, with one 2-dimensional facet corresponding to P_2

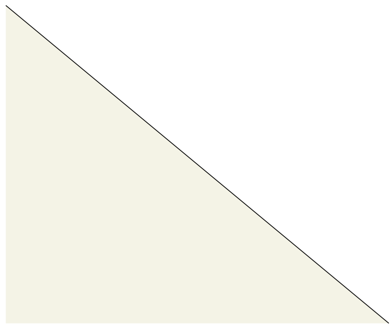

(C) The star at Q_2 in Σ, with one 2-dimensional facet corresponding to P_1

(D) The star at Q_1 in Σ, with two 2-dimensional facets

ridge by taking the star. The complex Σ is balanced if the star at each ridge is balanced.

To check whether a Γ-rational complex is balanced, we first compute the star at each ridge. For each star, we then compute the derived 1-dimensional fan at each ridge *of the star*. The star itself is pure of the same dimension as Σ, so its ridges have the same dimension as the ridges of Σ. If all the 1-dimensional rational fans derived in this way are balanced, then Σ is balanced.

6. The Structure Theorem

At this juncture we can state the precise form of the structure theorem. The remainder of this section will be dedicated to giving an overview of parts of the proof and discussing the various properties in the statement of the structure theorem. We note for motivation that every tropical algebraic variety is the support of a polyhedral complex, but the converse is not true. We would like to have a general criterion which can be used to determine whether for a given polyhedral complex, there is some tropical variety equal to its support. The structure theorem gives a partial solution, in that it describes the properties of a polyhedral complex whose support is the tropicalisation of an irreducible subvariety of the torus.

Theorem 6.1. *(The Structure Theorem) Let X be an irreducible subvariety in $(k^{\times})^n$ of dimension d, over an algebraically closed valued field. Then* trop(X) *is the support of a balanced* Γ_{val}*-rational polyhedral complex that is pure of dimension* d*. This complex is moreover connected through codimension one.*

Morally, this theorem tells us which polyhedral complexes are tropical varieties. The reader is encouraged to verify these properties on Example 3.5.

We note that whether a polyhedral complex is balanced depends on the weighting function, and the precise nature of the weighting which balances this complex is not given in the theorem. This is for several reasons. Perhaps most importantly, not every polyhedral complex has a weighting which balances it, so the theorem asserts the existence of a balanced weighting. We will not discuss the details of this weighting function or prove that the complex is balanced by it.

Example 6.2. Consider a 1-dimensional fan in $\mathbb{R}^2$ with rays represented by $(1, 0)$ and $(0, 1)$. There is no positive integer weighting on this complex that balances it because the integer lattice points are linearly independent.

Lemma 6.3. *For an irreducible variety* $X \subset (k^\times)^n$, trop(X) *is the support of a* Γ_{val}*-rational polyhedral complex.*

Proof. We will use the fact that any tropical variety can be written as the intersection of finitely many tropical hypersurfaces. The proof of this uses tropical bases; for details see Section 2.6 of [2].

A tropical hypersurface is the support of a Γ_{val}-rational polyhedral complex because for a single Laurent polynomial f, trop(f) is the minimum of finitely many affine pieces, with coefficients in the value group. Taking the intersection of two polyhedral complexes intersects their supports, so the intersection of all the polyhedral complexes corresponding to these hypersurfaces has support equal to trop(X).

The intersection of two Γ_{val}-rational polyhedra is Γ_{val}-rational, so the intersection of these tropical hypersurfaces gives the desired polyhedral complex. □

For a fixed subvariety X of the torus, denote a polyhedral complex obtained as in the proof of this lemma by Σ_X.

We will not discuss the proof of connectedness through codimension one in much depth. It can be shown by induction on the dimension d. The result for the $d = 1$ case is nontrivial. When $d = 1$ the polyhedral complex is a graph in $\mathbb{R}^n$, so it is connected through codimension one if and only if it is connected; one therefore need only show that trop(X) is connected as a set, so this result does not depend in any way on Σ_X.

Remark 6.4. This connectedness result is largely important for computation of tropical varieties. Given a subvariety X of the torus, we can define a graph whose vertices are the facets (d-dimensional cells) of trop(X), and two vertices are connected by an edge when their corresponding facets share a common ridge. This graph is connected precisely when trop(X) is connected through codimension one, since a path between two vertices in the graph corresponds to a facet-ridge path between the corresponding facets in trop(X). There are computational methods which, starting with one vertex in this graph, allow adjacent vertices to be identified. The graph being connected makes this method effective, since we do not require multiple starting vertices in different connected components.

So far we have not addressed the dimension of Σ_X. The proof that Σ_X is pure of dimension d relies on the fact that for any polyhedral complex structure on $\text{trop}(X)$, the star of a cell is itself a tropical variety. It also uses the fact that any subvariety X of the torus with $\text{trop}(X)$ finite must consist of finitely many points in the torus. See Section 3.3 of [2] for details.

Remark 6.5. The structure theorem tells us that every tropical variety $\text{trop}(X)$ for irreducible X is the support of a balanced weighted Γ_{val}-rational pure-dimensional polyhedral complex. It is natural to wonder whether the converse holds — given such a polyhedral complex, is its support always a tropical variety? In general, this is not the case. However, this is a complete classification for hypersurfaces.

Theorem 6.6. *Let Σ be a balanced weighted Γ_{val}-rational polyhedral complex in $\mathbb{R}^n$ that is pure of dimension $n-1$. Then there is a tropical polynomial F whose coefficients lie in Γ_{val} with $V(F) = \Sigma$. By choosing a Laurent polynomial $f \in k[x_1^{\pm}, \ldots, x_n^{\pm}]$ s.t.* $\text{trop}(f) = F$*, we see that* $\Sigma = \text{trop}(V(f))$ *and so* $\text{supp}(\Sigma)$ *is a tropical variety.*

7. References

[1] M. Michalek and B. Sturmfels, *Invitation to Nonlinear Algebra*. unpublished, https://personal-homepages.mis.mpg.de/michalek/NonLinearAlgebra.pdf, 2020.

[2] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*. Graduate Studies in Mathematics, American Mathematical Society, 2015.

[3] G. M. Ziegler, *Lectures on polytopes*, vol. 152. Springer Science & Business Media, 2012.