The structure of Tropical Varieties

N.Davis

November 8, 2005

Contents

1 Introduction

The main result proved in this paper is on the structure of tropical varieties. Tropical varieties are an interesting kind of subset of $\mathbb{R}^n$ which can be defined in many different ways, and have a variety of applications in different fields of mathematics. In our second section we give a first definition of these sets. The third section is devoted to developing a little Gröbner basis theory, and in the fourth section we use these ideas to show that two quite different definitions also define the same sets. The proof about the structure of tropical varieties is the content of the fifth section.

2 Tropical Varieties

We shall be approaching tropical varieties from the perspective of tropical geometry, an active and relatively new field of research. The purpose of this section is to give a very brief introduction to tropical geometry, and explain how the need for tropical varieties arises when one wishes to study it. Much of this section is based on the article *First Steps in Tropical Geometry* [7], by Jürgen Richter–Gerbert, Bernd Sturmels and Thorsten Theobald, to which the reader should refer for a fuller introduction, including explanatory diagrams.

2.1 The Tropical Semiring

Tropical algebraic geometry is the geometry of the *tropical semiring* $(\mathbb{R} \cup \infty, \oplus, \odot)$. The underlying set $\mathbb{R} \cup \infty$ is the set of real numbers augmented by $+\infty$. The arithmetic operation $\oplus$, or *tropical addition*, is the standard minimum function on real numbers, while $\odot$, or *tropical multiplication*, is ordinary addition. It is easy to check that what we have here is indeed a semiring. The minimum function is associative and distributive, ordinary addition is associative, and we also have distributivity. The element ∞ acts as an additive unit, and 0 acts as a multiplicative unit. Note that we don't have a ring, because of a lack of additive inverses.

Using the operations of this semiring, we can construct new definitions of both monomials and polynomials. A *tropical monomial* (in n variables) has the form $c \odot x_1^{a_1} \odot \cdots \odot x_n^{a_n}$, where $c \in \mathbb{R}$. The powers of the variables are also computed tropically; for instance, x_1^3 denotes $x_1 \odot x_1 \odot x_1$. Thus a tropical monomial represents a classical linear function $\mathbb{R}^n \to \mathbb{R}$, where $(x_1, \ldots, x_n)$ maps to $a_1 x_1 + \cdots + a_n x_n + c$. A *tropical polynomial* is defined as a finite tropical sum of tropical monomials,

$$F(x_1, \ldots, x_n) = c_1 \odot x_1^{a_{11}} \odot \cdots \odot x_n^{a_{1n}} \oplus \cdots \oplus c_r \odot x_1^{a_{r1}} \odot \cdots \odot x_n^{a_{rn}}$$

Since each tropical monomial appearing in F represents a classical linear function, and tropical addition is in fact the standard minimum function, F

will represent a piecewise–linear concave function which is the minimum of r linear functions.

For every tropical polynomial F, there will be some closed set of points $x = (x_1, \ldots, x_n)$ in $\mathbb{R}^n$ at which F is non–linear. These are the points at which the minimum is attained by at least two of the tropical monomials appearing in F. We define this set as the *tropical hypersurface* $\mathscr{T}(F)$ of the tropical polynomial F.

2.2 Tropical Linear Spaces

In the study of tropical geometry, a definition of a tropical linear space is naturally required. There are two obvious candidates for this definition:

- A tropical linear space L is a subset of $\mathbb{R}^n$ which consists of all solutions $(x_1, x_2, \ldots, x_n)$ to a finite system of tropical linear equations

$$a_1 \odot x_1 \oplus \cdots \oplus a_n \odot x_n = b_1 \odot x_1 \oplus \cdots \oplus b_n \odot x_n$$

- A tropical linear space L in $\mathbb{R}^n$ consists of all tropical linear combinations $\lambda \odot a \oplus \mu \odot b \oplus \cdots \oplus \nu \odot c$ of a fixed finite subset $\{a, b, \ldots, c\} \subset \mathbb{R}^n$.

The authors of *First Steps in Tropical Geometry* show that both of these possible definitions are in fact unsatisfactory, in the first case because the sets defined are too small, and in the second case because they are too large. This is where tropical varieties enter the picture. They are used in the correct definition of a tropical linear space.

2.3 Tropical Varieties

The definition of a tropical variety is not too complex, but does make use of a map which is an example of a *valuation*. It is therefore convenient to explain a small amount of valuation theory at this point, most of which we will not use until later.

2.3.1 Some valuation theory

There are two basic types of valuations. We will be interested in non-Archimedean valuations.

Definition 2.1 (non–Archimedean Valuation). *Let K be a field. A non–Archimedean valuation is a map $\nu : K \to \mathbb{R} \cup \infty$ with the following properties:*

1. $\nu(x) = \infty \Leftrightarrow x = 0$

2. $\nu(xy) = \nu(x) + \nu(y)$

3. $\nu(x+y) \geq \inf(\nu(x), \nu(y))$

There is another general fact about non–Archimedean valuations which we shall find particularly useful later on.

Proposition 2.1. *The subset of K defined by $R_K = \{x \in K : \nu(x) \geq 0\}$ is a local ring, with unique maximal ideal $M_K = \{x \in K : \nu(x) > 0\}$*

Proof. First we verify that that the subset R_K as we have defined it is a ring. Property 2 gives us closure under multiplication, and also implies that $\nu(1) = 0$, which in turn tells us that $\nu(-1) = 0$ since

$$\nu(-1) + \nu(-1) = \nu((-1)(-1)) = \nu(1) = 0.$$

Thus $\nu(x) = \nu(-x)$ for all $x \in K$, guaranteeing that each element in R_K has an additive inverse. Property 3 ensures closure under addition.

Property 2 tells us that our set M_K is precisely the set of non–units in R_K, which means that it contains all proper ideals of R_K. We need only show that M_K is an ideal, and our proof will be complete. Take an element $x \in M_K$, and an element $r \in R_K$. It is clear that $rx \in M_K$ since

$$\nu(rx) = \nu(r) + \nu(x) \geq \nu(x) > 0.$$

Now suppose that x and y lie in M_K. Then

$$\nu(x+y) \geq \inf(\nu(x), \nu(y)) > 0,$$

so $x + y \in M_K$. So M_K is in fact an ideal, and hence the unique maximal ideal of the local ring R_K. □

2.3.2 Definition of a tropical variety

As noted earlier, there is one valuation in particular that is central in our definition of a tropical variety. The field K on which we shall make this valuation is the algebraic closure of the rational function field, i.e. $K = \overline{\mathbb{C}[t]}$. Any non–zero algebraic function $p(t)$ in this field can be locally expressed as a *Puiseux series*

$$p(t) = c_1 t^{q_1} + c_2 t^{q_2} + c_3 t^{q_3} + \cdots$$

where $c_1, c_2, \ldots$ are non–zero complex numbers and $q_1 < q_2 < \cdots$ are rational numbers with bounded denominator. The map that we are interested in is $\deg : K^* \to \mathbb{Q}$, which sends $p(t) \in K$ to q_1, the smallest exponent of t which occurs in it. If we define $\deg(0) = \infty$, and consider $\mathbb{Q}$ as a subset of $\mathbb{R}$, we can check that we do indeed have a valuation $K \to \mathbb{R}$:

$\deg(0) = \infty$, and the value ∞ is clearly not attained by any other element of K.

$$\begin{aligned}
\deg(xy) &= \deg((c_{11}t^{q_{11}} + c_{12}t^{q_{12}} + \cdots)(c_{21}t^{q_{21}} + c_{22}t^{q_{22}} + \cdots)) \\
&= \deg(c_{11}c_{21}t^{q_{11}+q_{21}} + c_{11}c_{22}t^{q_{11}+q_{22}} + c_{12}c_{21}t^{q_{12}+q_{21}} + c_{12}c_{22}t^{q_{12}+q_{22}} + \cdots) \\
&= q_{11} + q_{21} \\
&= \deg(x) + \deg(y)
\end{aligned}$$

$$\begin{aligned}
\deg(x+y) &= \deg(c_{11}t^{q_{11}} + c_{12}t^{q_{12}} + \cdots + c_{21}t^{q_{21}} + c_{22}t^{q_{22}} + \cdots) \\
&\geq \inf(q_{11}, q_{21}) \\
&= \inf(\deg(x), \deg(y))
\end{aligned}$$

We also note now that this valuation is dense in $\mathbb{R}$, since it takes all rational values; for any $p, q \in \mathbb{Z}$, we have $t^{p/q} \in K$ since it is a root of $x^q - t^p$, and clearly $\deg(t^{p/q}) = p/q$. We can now define a tropical variety. We note that it is not critical for the definition that *deg* is a valuation. We will be using that fact later on.

Definition 2.2 (Tropical Variety). *Let I be any ideal in the Laurent polynomial ring $K[x_1^{\pm 1}, \ldots, x_n^{\pm n}]$ and consider its affine variety $V(I) \subset (K^*)^n$ over the algebraically closed field K. The image of $V(I)$ under the map*

$$(p_1(t), \ldots, p_n(t)) \mapsto (\deg(p_1(t)), \ldots, \deg(p_n(t)))$$

is a subset of $\mathbb{Q}^n$. We take the topological closure of this set. The resulting subset of $\mathbb{R}^n$ is the tropical variety $\mathscr{T}(I)$.

Here we have introduced the notation K^* to mean $K \setminus \{0\}$. We now have our definition of a tropical variety, and are only a few short steps away from the definition of a tropical linear space. A polynomial f is *homogeneous* if all monomials $x_1^{i_1} \cdots x_n^{i_n}$ appearing in it have the same total degree $d = i_1 + \cdots + i_n$. An ideal I is homogeneous if the generating polynomials which define it are all homogeneous. Such a homogeneous ideal I defines a variety in projective space $\mathbb{P}_K^{n-1}$ minus the the coordinate hyperplanes $x_i = 0$. This works because if f is homogeneous, $f(\lambda a_0, \lambda a_1, \ldots, \lambda a_n) = \lambda^d f(a_0, a_1, \ldots, a_n)$ for all $\lambda \in K$ and so scalar multiplication fixes the variety in $\mathbb{R}^n$. The *tropical projective variety* of a homogeneous ideal I is then the image of $\mathscr{T}(I)$ in the space $\mathbb{R}^n/\mathbb{R}(1, 1, \ldots, 1) = \mathbb{TP}^{n-1}$. It is this definition that leads to the definition of a *tropical linear space* as a subset of $\mathbb{TP}^{n-1}$ of the form $\mathscr{T}(I)$ where I is generated by linear forms $p_1(t) \cdot x_1 + p_2(t) \cdot x_2 + \cdots + p_n(t) \cdot x_n$, whose coefficients $p_i(t)$ are algebraic functions in one complex variable t.

We will not be delving any deeper into tropical geometry. It is tropical varieties themselves that are the focus of what follows.

3 Some Gröbner basis theory

In the previous section, we defined tropical varieties. It turns out that there are a number of equivalent ways to define these sets, and that they are known by a number of different names: logarithmic limit sets, Bergman fans, Bieri–Groves sets, and non–Archimedean amoebas are examples. In section three, we will show three definitions of tropical varieties to be equivalent, and in section four, we will demonstrate how these new perspectives can lead to a simplification of proofs in the area. In both sections, we will be using Gröbner basis techniques. This section provides the necessary theory and definitions. Much of the material is taken from *Gröbner bases and convex polytopes.* [9]

3.1 The basics

Let k be any field, and $k[x] = k[x_1, \ldots, x_n]$ the polynomial ring on n variables. We will denote the monomials in $k[x]$ by $x^a = x_1^{a_1} x_2^{a_2} \ldots x_n^{a_n}$. Note that they can be identified with the lattice points $a = (a_1, \ldots, a_n)$ in $\mathscr{A} = \mathbb{N}^n$. A total order $\prec$ on $\mathbb{N}^n$ is thus a total order on the monomials of $k[x]$. Such a total order $\prec$ is called a *term order* if it satisfies the following conditions for all monomials $x^a, x^b, x^c \in k[x]$:

1. if $x^a \neq 1$ then $1 \prec x^a$
2. if $x^a \prec x^b$ then $x^a x^c \prec x^b x^c$

Given a term order $\prec$ on $k[x] = k[x_1, \ldots, x_n]$, and a non–zero polynomial $f \in k[x]$, f has a unique *initial monomial*, denoted by $\text{in}_\prec(f)$, which is the greatest term of f under the ordering $\prec$. This definition leads to a division algorithm for polynomials in $k[x]$.

Theorem 3.1 (Division Algorithm). *Suppose $f \in k[x]$ and $g_1, g_2, \ldots, g_s$ are non–zero elements of $k[x]$. Let $\prec$ be a term order on $k[x]$. Then there exist $u_1, \ldots, u_s, h$ so that*

$$f = u_1 g_1 + \cdots + u_s g_s + h$$

where no term of h is divisible by any $\text{in}_\prec(g_i)$ and for each i, $\text{in}_\prec(u_i g_i) \leq \text{in}_\prec(f)$.

When f is written in this form, we say that f is *reduced* to h via the g_i. We won't prove the theorem here; it is a standard result which can be found in any text book which deals with Gröbner bases, for example [3] p.63.

If I is an ideal in $k[x]$, then its *initial ideal* is the monomial ideal

$$\text{in}_\prec(I) := \langle \text{in}_\prec(f) : f \in I \rangle$$

The set of monomials which do not lie in $\mathrm{in}_\prec(I)$ are called the *standard monomials*.

We now define what is obviously the central object in Gröbner basis theory.

Definition 3.1 (Gröbner basis). *A finite subset $\mathcal{G} \subset I$ is a Gröbner basis for I with respect to $\prec$ if the following equivalent statements are true:*

- $\mathrm{in}_\prec(I)$ *is generated by* $\{\mathrm{in}_\prec(g) : g \in \mathcal{G}\}$.
- *A polynomial $f \in k[x]$ reduces to* 0 *via $\mathcal{G}$ if and only if $f \in I$.*

We can now prove the following useful result.

Proposition 3.1. *The images of the standard monomials form a k–vector space basis for the residue ring $k[x]/I$.*

Proof. By the Hilbert basis theorem, we need only finitely many elements g_i in a Gröbner basis $\mathcal{G}$. The division algorithm then tells us that we can write $f = u_1 g_1 + \cdots + u_s g_s + h$ for any $f \in k[x]$, and that the terms in h must be standard monomials, since they are not divisible by the $\mathrm{in}_\prec(g_i)$, which generate $\mathrm{in}_\prec(I)$. By the definition of a Gröbner basis, h will be equal to 0 if and only if $f \in I$. So the standard monomials span $k[x]/I$. They are also linearly independent, since if some linear combination of standard monomials were in I, one of them would necessarily be in $\mathrm{in}_\prec(I)$. □

There are some further refinements which can be made to the definition of a Gröbner basis. If no monomial in a Gröbner basis $\mathcal{G}$ is redundant, then $\mathcal{G}$ is *minimal*. If the coefficient of $\mathrm{in}_\prec(g)$ is 1 for all $g \in \mathcal{G}$, and for any two distinct elements $g, g' \in \mathcal{G}$, no term of g' is divisible by $\mathrm{in}_\prec(g)$, then $\mathcal{G}$ is *reduced*. The reduced Gröbner basis of an ideal is unique. It can be shown ([9] p.1) that every ideal $I \subset k[x]$ has only finitely many distinct initial ideals, and so there is some finite subset $\mathcal{U} \subset I$ which is a Gröbner basis for I with respect to all term orders $\prec$ simultaneously. We call a set $\mathcal{U}$ fulfilling this requirement a *universal Gröbner basis* for I.

For our purposes, we need to generalise these definitions a little. Specifically, we are interested in representing term orders by *weight vectors*. The idea is quite simple. Since we are denoting the monomials in $k[x]$ by $x^a = x_1^{a_1} x_2^{a_2} \ldots x_n^{a_n}$, a polynomial f in $k[x]$ will have the form $f = \sum c_i \cdot x^{a_i}$, where i sums over the different terms. Now, for a fixed vector $w = (w_1, \ldots, w_n) \in \mathbb{R}^n$, the *initial form* $\mathrm{in}_w(f)$ is defined as the sum of the terms $c_i \cdot x^{a_i}$ such that the inner product $w \cdot a_i$, or *weight*, is maximal. We define the *initial ideal* $\mathrm{in}_w(I)$ of an ideal I in the obvious way:

$$\mathrm{in}_w(I) := \langle \mathrm{in}_w(f) : f \in I \rangle$$

Note that while this ideal is monomial (i.e. the generators are all monomials) for sufficiently generic choices of w, it needn't be for all w. There may be polynomials in I which contain two or more terms of equal maximal weight. The special values of w for which this happens are not a problem for us. In fact, they are precisely the values in which we will be interested.

For $w \geq 0$ and $\prec$ an arbitrary term order, we can also define a new term order $\prec_w$ in the following way. For $a, b \in \mathscr{A}$ we have

$$x^a \prec_w x^b \quad :\Longleftrightarrow \quad w \cdot a < w \cdot b \text{ or } (w \cdot a = w \cdot b \text{ and } x^a \prec x^b)$$

That is, we use w to order monomials as far as possible, and then apply $\prec$ when necessary. We shall need this notation later, along with the following propositions.

Proposition 3.2. *For every ideal $I \subset k[x]$,* $\quad \mathrm{in}_\prec(\mathrm{in}_w(I)) = \mathrm{in}_{\prec_w}(I)$

Proof. From the definition of the term order $\prec_w$, we can see that for a given polynomial f, we have $\mathrm{in}_\prec(\mathrm{in}_w(f)) = \mathrm{in}_{\prec_w}(f)$. This implies that $\mathrm{in}_\prec(\mathrm{in}_w(I))$ and $\mathrm{in}_{\prec_w}(I)$ are generated by the same set of monomials, and hence they are equal. □

Proposition 3.3. *If $w \geq 0$ and $\mathcal{G}$ is a Gröbner basis for I with respect to $\prec_w$, then $\{\mathrm{in}_w(g) : g \in \mathcal{G}\}$ is a Gröbner basis for $\mathrm{in}_w(I)$ with respect to $\prec$.*

Proof. This follows immediately from Proposition 3.2. □

Proposition 3.4. *Let $I \subset k[x]$ be an ideal, $w', w \in \mathbb{R}^n$ and $\epsilon > 0$ sufficiently small. Then*

$$\mathrm{in}_w(\mathrm{in}_{w'}(I)) = \mathrm{in}_{w'+\epsilon \cdot w}(I)$$

Proof. Let $\mathcal{G}$ be the reduced Gröbner basis of I with respect to the term order $\prec_{w'+\epsilon \cdot w}$. For each $g \in \mathcal{G}$ we have $\mathrm{in}_w(\mathrm{in}_{w'}(g)) = \mathrm{in}_{w'+\epsilon \cdot w}(g)$. By Proposition 3.3, the set

$$\{\mathrm{in}_{w+\epsilon \cdot w'}(g) : g \in \mathcal{G}\} = \{\mathrm{in}_w(\mathrm{in}_{w'}(g)) : g \in \mathcal{G}\}$$

is a Gröbner basis for $\mathrm{in}_{w'+\epsilon \cdot w}(I)$ with respect to $\prec$. This tells us that $\mathrm{in}_{w'+\epsilon \cdot w}(I) \subseteq \mathrm{in}_w(\mathrm{in}_{w'}(I))$ since the elements in a Gröbner basis of the ideal on the left are clearly contained in the ideal on the right. If the containment were proper, then it would remain proper after we took the initial monomial ideals of each side with respect to $\prec$. This is impossible by Proposition 3.1, since it would require the ring $k[x]/I$ to have two different dimensions as a k–vector space. □

3.2 The Gröbner fan

In this subsection we shall define the Gröbner fan of an ideal. This is an object which will be particularly useful in section four. First we need a little polyhedral geometry.

3.2.1 Polyhedral geometry

A *polyhedron* is a finite intersection of closed half spaces in $\mathbb{R}^n$. A polyhedron P can be written as

$$P = \{x \in \mathbb{R}^n : A \cdot x \leq b\},$$

where A is a matrix with n columns, and a and b are vectors in $\mathbb{R}^n$. In the special case where $b = 0$, P is a *(polyhedral) cone*, and can be written in form

$$P = \{\lambda_1 u_1 + \cdots + \lambda_n u_n : \lambda_i \in \mathbb{R}_+, u_i \in \mathbb{R}^n\}$$

for some vectors $u_1, \ldots, u_n \in \mathbb{R}^n$. For a polyhedron P in $\mathbb{R}^n$, and a vector $w \in \mathbb{R}^n$, we define

$$\text{face}_w(P) := \{u \in P : w \cdot u \geq w \cdot v \text{ for all } v \in P\}$$

Every subset F of P which has this form is called a *face* of P. If $P \subset \mathbb{R}^n$ is a polyhderon, and F is a face of P, then the *normal cone* of F at P is

$$\mathcal{N}_P(F) = \{w \in \mathbb{R}^n : \text{face}_w(P) = F\}$$

With every polynomial $f = \sum_{i=1}^m c_i \cdot x^{a_i}$ in $k[x]$ we can associate a bounded polyhedron, or *polytope*, which is the convex hull of the points a_i. We call this polytope the *Newton polytope*:

$$\text{New}(f) := \mathit{conv}\{a_i : i = 1, \ldots, m\}$$

The Newton polytope satisfies the equation $\text{New}(f \cdot g) = \text{New}(f) + \text{New}(g)$. We also have the following relationship between faces and initial forms: $\text{face}_w(\text{New}(f)) = \text{New}(\text{in}_w(f))$.

A *(polyhedral) complex* Δ is a finite collection of polyhedra in $\mathbb{R}^n$ such that

1. if $P \in \Delta$ and F is a face of P, then $F \in \Delta$.
2. if $P_1, P_2 \in \Delta$, then $P_1 \cap P_2$ is a face of P_1 and P_2.

The *support* of a complex Δ is defined as the union of all polyhedra in Δ. A complex Δ which consists entirely of cones is called a *fan.*

3.2.2 The Gröbner fan

Now that we have a little polyhedral geometry, we can proceed with our definition of the Gröbner fan, and give a proof that it is indeed a fan in the sense we have just defined. It will be comprised of equivalence classes of weight vectors.

Fix an ideal $I \in k[x]$. Two weight vectors $w, w' \in \mathbb{R}^n$ are called *equivalent (with respect to I)* if and only if $\text{in}_w(I) = \text{in}_{w'}(I)$.

Proposition 3.5. *Each equivalence class of weight vectors is a relatively open convex polyhedral cone.*

Proof. Let $C[w]$ denote the equivalence class containing w. We fix an arbitrary term order $\prec$ as a "tie breaker". Let $\mathcal{G}$ be a reduced Gröbner basis of I with respect to $\prec_w$. We claim the following formula holds:

$$C[w] = \{w' \in \mathbb{R}^n : \mathrm{in}_{w'}(g) = \mathrm{in}_w(g) \text{ for all } g \in \mathcal{G}\}$$

The right hand side of this formula is defined by the equations $w' \cdot a = w' \cdot b$ and the inequalities $w' \cdot a > w' \cdot c$ where x^a and x^b run over the terms in $\mathrm{in}_w(g)$ and x^c runs over the terms of g which do not appear in $\mathrm{in}_w(g)$. Thus, if the formula holds, $C[w]$ is expressible as an intersection of hyperplanes and open half–spaces, and the proposition is true.

We first prove the inclusion "$\supseteq$". If w' lies in the set on the right hand side, then

$$\mathrm{in}_w(I) = \langle \mathrm{in}_{w'}(g) : g \in \mathcal{G}\rangle \subseteq \mathrm{in}_{w'}(I).$$

The equality on the left follows immediately from Proposition 3.3. We need a little more work to prove equality on the right, since $\mathcal{G}$ is defined with respect to the term order $\prec_w$, not $\prec_{w'}$, but we have an inclusion since $\mathcal{G} \subseteq I$. If the inclusion were proper, then the following inclusion of initial ideals would also be proper:

$$\mathrm{in}_{\prec_w}(I) = \mathrm{in}_{\prec}(\mathrm{in}_w(I)) \subset \mathrm{in}_{\prec}(\mathrm{in}_{w'}(I)) = \mathrm{in}_{\prec_{w'}}(I)$$

This follows from the fact that two ideals are equal if and only if their reduced Gröbner bases are equal, and uses Proposition 3.2. As we noted in the proof of Proposition 3.4, a strict inclusion of initial monomial ideals is impossible by Proposition 3.1, so we conclude that $\mathrm{in}_w(I) = \mathrm{in}_{w'}(I)$.

Now to prove the inclusion "$\subseteq$". Suppose $w' \in C[w]$. Again using Proposition 3.2, we can see that the set $\mathrm{in}_w(\mathcal{G}) = \{\mathrm{in}_w(g) : g \in \mathcal{G}\}$ is the reduced Gröbner basis of $\mathrm{in}_w(I) = \mathrm{in}_{w'}(I)$ with respect to $\prec$. Fix $g \in \mathcal{G}$. Then $\mathrm{in}_{w'}(g) \in \mathrm{in}_{w'}(I)$ reduces to zero with respect to $\mathrm{in}_w(\mathcal{G})$ using the term order $\prec$. Now, the fact that $\mathcal{G}$ is a *reduced* Gröbner basis of $\mathrm{in}_w(I)$ with respect to $\prec$ means that $\mathrm{in}_{\prec_w}(g)$ is the only monomial of g which can lie in $\mathrm{in}_{\prec_w}(I)$. So, the fact that $\mathrm{in}_{w'}(g)$ reduces to zero means that the monomial $m := \mathrm{in}_{\prec_w}(g)$ must appear in $\mathrm{in}_{w'}(g)$. Thus we can write $\mathrm{in}_w(g) = m + h$ and also $\mathrm{in}_{w'}(g) = m + h'$, where h and h' are both polynomials consisting of monomial terms not lying in $\mathcal{G}$. Then, when we look explicitly at the reduction of $\mathrm{in}_{w'}(g)$ via $\mathrm{in}_w(\mathcal{G})$, we arrive at the polynomial $h' - h$ after one step. This polynomial must lie in the ideal $\mathrm{in}_w(I)$, but by construction none of its terms lie in $\mathrm{in}_{\prec_w}(I) = \mathrm{in}_{\prec}(\mathrm{in}_w(I))$. This means that $h' - h$ can only equal 0, so that we have $\mathrm{in}_{w'}(g) = \mathrm{in}_w(g)$. This completes the proof. □

We are now ready for our definition of the Gröbner fan.

Definition 3.2 (Gröbner fan). *The Gröbner fan $GF(I)$ is the set of closed cones $\overline{C[w]}$ for all $w \in \mathbb{R}^n$.*

As we mentioned earlier, this definition requires some justification.

Proposition 3.6. *The Gröbner fan $GF(I)$ is a fan.*

Proof. For the proof, we need to reformulate our cones $C[w]$. We write

$$C[w] = \mathcal{N}_Q(\text{face}_w(Q)), \text{ where } Q := \text{New}(\prod_{g\in\mathcal{G}} g) = \sum_{g\in\mathcal{G}} \text{New}(g)$$

We can quite easily convince ourselves that this is an equivalent definition.

$$\begin{aligned}
C[w] &= \mathcal{N}_Q(\text{face}_w(\sum_{g\in\mathcal{G}} \text{New}(g))) \\
&= \{w' \in \mathbb{R}^n : \text{face}_{w'}(\sum_{g\in\mathcal{G}} \text{New}(g)) = \text{face}_w(\sum_{g\in\mathcal{G}} \text{New}(g))\} \\
&= \{w' \in \mathbb{R}^n : \sum_{g\in\mathcal{G}}(\text{face}_{w'}(\text{New}(g))) = \sum_{g\in\mathcal{G}}(\text{face}_w(\text{New}(g)))\} \\
&= \{w' \in \mathbb{R}^n : \sum_{g\in\mathcal{G}}(\text{New}(\text{in}_{w'}(g))) = \sum_{g\in\mathcal{G}}(\text{New}(\text{in}_w(g)))\} \\
&= \{w' \in \mathbb{R}^n : \text{in}_{w'}(g) = \text{in}_w(g) \text{ for all } g \in \mathcal{G}\}
\end{aligned}$$

Now we can verify that the Gröbner fan satisfies the two axioms of a complex. Let w' be any vector in the closure $\overline{C[w]}$ of an equivalence class $C[w]$. Then by Proposition 3.4,

$$\text{in}_w(\text{in}_{w'}(I)) = \text{in}_{w'+\epsilon\cdot w}(I) = \text{in}_w(I),$$

so $\text{in}_w(I)$ is an initial ideal of $\text{in}_{w'}(I)$. Hence there exists a term order $\prec$ such that $\text{in}_{\prec_w}(I) = \text{in}_{\prec_{w'}}(I)$. Let $\mathcal{G}$ be the reduced Gröbner basis of I with respect to $\prec_w$, and let Q be the polytope $\text{New}(\prod_{g\in\mathcal{G}} g) = \sum_{g\in\mathcal{G}} \text{New}(g)$ as above. Since $\mathcal{G}$ is the reduced Gröbner basis with respect to $\prec_{w'}$ as well as $\prec_w$, the equivalence classes $C[w]$ and $C[w']$ are normal cones of the *same* polytope Q, with

$$C[w] = \mathcal{N}_Q(\text{face}_w(Q)) \text{ and } C[w'] = \mathcal{N}_Q(\text{face}_{w'}(Q))$$

In fact, because $\text{in}_w(I)$ is an initial ideal of $\text{in}_{w'}(I)$, $\text{face}_w(Q)$ is actually a face of $\text{face}_{w'}(Q)$. This means that $\overline{C[w']}$ is a face of $\overline{C[w]}$. So $GF(I)$ satisfies the first axiom of a complex, since if we choose w' in some face F of a closure of an equivalence class $C[w]$, then $F = \overline{C[w']}$ is a face of $\overline{C[w]}$. The second also follows. Let $w, w' \in \mathbb{R}^n$ and consider the closed convex cone $P := \overline{C[w]} \cap \overline{C[w']}$. We have proved that for each $w'' \in P$ the cone $\overline{C[w'']}$ is a face of both $\overline{C[w]}$ and $\overline{C[w']}$. Hence P is a finite union of such common faces. The only way that such a union of faces can be convex is if one contains all of the others, so P is itself a common face of $\overline{C[w]}$ and $\overline{C[w']}$. This completes the proof that $GF(I)$ is a fan. □

4 Three equivalent definitions of a Tropical Variety

We have now developed enough Gröbner basis theory and polyhedral geometry to examine the proof of a theorem, found in *The tropical Grassmanian* [8], which is an example of how different definitions of tropical varieties can be shown to be equivalent.

4.1 Statement of equivalence theorem

First, we adapt the definitions of Gröbner basis theory to our specific needs. The field we are dealing with is $K = \overline{\mathbb{C}[t]}$. Recall that we have defined a valuation $\deg : K \to \mathbb{R}$, and that there will be a corresponding local ring R_K contained in K. Polynomials in $K[x]$ will have form $f = \sum c_i \cdot x^{a_i}$, with $c_i \in K^*$, but we will define the *w–weight* of a term $c_a \cdot x_1^{a_1} \dots x_n^{a_n}$ as $\deg(c_a) + a_1 w_1 + \dots + a_n w_n$, rather than $c_a + a_1 w_1 + \dots + a_n w_n$. The *initial form* $\mathrm{in}_w(f)$ of a polynomial f will also be defined slightly differently. Set $\tilde{f}(x_1, \dots, x_n) = f(t^{w_1} x_1, \dots, t^{w_n} x_n)$. Let ι be the smallest weight of any term of f, so that $t^{-\iota}\tilde{f}$ is a non–zero element in $R_K[x]$. Then $\mathrm{in}_w(f)$ is defined as $t^{-\iota}\tilde{f} \bmod M_K$. Since in this case $R_K/M_K = \mathbb{C}$, we have $\mathrm{in}_w(f) \in \mathbb{C}[x]$. So, given any ideal $I \subset K[x]$ its *initial ideal* will be defined

$$\mathrm{in}_w(I) = \langle \mathrm{in}_w(f) : f \in I \rangle \subset \mathbb{C}[x]$$

Note that although we are now using the terms of smallest weight in the initial ideals, instead of those of maximal weight, this makes no real difference to our theory.

We need one more definition before we state the thereom. We define the image of f under the map

$$\sum_{a \in \mathscr{A}} c_a x_1^{a_1} \dots x_n^{a_n} \mapsto \sum_{a \in \mathscr{A}} \deg(c_a) x_1^{a_1} \dots x_n^{a_n}.$$

as the tropical polynomial $\mathrm{trop}(f)$. We are now ready to proceed.

Theorem 4.1. *For an ideal $I \subset K[x]$ the following subsets of $\mathbb{R}^n$ coincide:*

1. *The closure of the set* $\{(\deg(u_1), \dots, \deg(u_n)) : (u_1, \dots, u_n) \in V(I)\}$
2. *The intersection of the tropical hypersurfaces* $\mathscr{T}(\mathrm{trop}(f))$ *where* $f \in I$
3. *The set of all vectors $w \in \mathbb{R}^n$ such that $\mathrm{in}_w(I)$ contains no monomial.*

We are now defining the tropical variety of an ideal $I \subset K[x]$, rather than one in $K[x^{\pm 1}]$, but the concepts are very similar. We can see the Laurent polynomial ring $K[x^{\pm 1}]$ as the localisation of $K[x]$ at the maximal ideal corresponding to 0. We will not explore this point in any detail here.

One of the inclusions in the proof of this theorem is quite technical, so it will be convenient to prove some lemmas before launching into it.

4.2 Technical Lemmas

Lemma 4.1. *For an ideal $I \subset K[x]$, $\text{in}_w(I)$ contains no monomial if and only if $\text{in}_w(f)$ contains no monomial for all $f \in I$.*

Proof. It is clear that if $\text{in}_w(I)$ contains no monomial then $\text{in}_w(f)$ contains no monomial for all $f \in I$. The other direction is not quite as obvious. We need to show that all elements of $\text{in}_w(I)$ have the form $\text{in}_w(f)$ for some f. Consider the polynomial $h = p(x) \cdot \text{in}_w(f)$, where

$$p(x) = a_0 + a_1 x + \cdots + a_n x^n \in \mathbb{C}[x].$$

We have

$$\begin{aligned} h &= a_0 \text{in}_w(f) + a_1 x \text{in}_w(f) + \cdots + a_n x^n \text{in}_w(f) \\ &= \text{in}_w(a_0 f) + \text{in}_w(a_1 x f) + \cdots + \text{in}_w(a_n x^n f) \\ &= \text{in}_w(f_0) + \text{in}_w(f_1) + \cdots + \text{in}_w(f_n) \end{aligned}$$

where $f_i \in I$ for all i.
Also, for any $f_1, f_2 \in I$, we have

$$\text{in}_w(f_1) + \text{in}_w(f_2) = \text{in}_w(f_1 + t^\eta f_2) = \text{in}_w(f_3)$$

for some $f_3 \in I$, where η is the weight by which the term of least weight in f_2 is less than the term of least weight in f_1. □

Lemma 4.2. *Let R_1 and R_2 be rings, and φ a homomorphism from R_1 to R_2. Given a prime ideal P of R_2, the preimage of the intersection $\varphi(R_1) \cap (P)$ is a prime ideal $Q = \varphi^{-1}(\varphi(R_1) \cap P)$ of R_1.*

Proof.

- For $a, b \in Q$, $\varphi(a) = A$ and $\varphi(b) = B$ lie in $\varphi(R_1) \cap P$. Since P is an ideal, $A + B \in P$, and since φ is a homomorphism, $A + B = \varphi(a + b)$, so $a, b \in Q$ implies $a + b \in Q$.

- For $a \in Q$, $r \in R_1$, we have $\varphi(ar) = \varphi(a)\varphi(r)$. Since $\varphi(a) \in P$, and P is an ideal, this means $\varphi(ar) \in P$, so $ar \in Q$.

- Let $ab \in Q$. Then $\varphi(ab) = \varphi(a)\varphi(b) \in P$, and since P is prime, either $\varphi(a)$ or $\varphi(b)$ lies in P. Hence a or b lies in Q.

□

Lemma 4.3. *Let K be a field, R_K the local ring corresponding to a valuation ν of K into $\mathbb{R}$, and M_K the corresponding maximal ideal of R_K. Then M_K is the only proper non–zero prime ideal of R_K.*

Proof. Assume that there is some other proper prime ideal P. Since M_K is the only maximal ideal of the ring, P is strictly contained in M_K. Choose an element a that is in M_K but not P, and some non–zero element b of P. Since a has a positive valuation, and neither a nor b can have an infinite valuation, there is some positive integer n such that $n\nu(a) > \nu(b)$. Then $\nu(a^n) > \nu(b)$ and so $\nu(a^n b^{-1}) > 0$. (Remembering this is allowed because our elements come from a field K). We now have $a^n b^{-1} \in R_K$, which implies that $a^n = a^n b^{-1} b \in P$, as P is a prime ideal of R_K. It then follows that a must be in P, so we have derived a contradiction. □

Lemma 4.4. *Let R be a ring, S a multiplicatively closed subset of R, and I an ideal of R. Then there is a bijective correspondence between prime ideals of $S^{-1}R/S^{-1}(I)$ and prime ideals of R/I which do not meet S.*

Proof. By the fourth ring isomorphism theorem, there is a bijection between the ideals of $S^{-1}R/S^{-1}(I)$ and the ideals of $S^{-1}R$ which contain $S^{-1}(I)$. It is also well known that there is a bijection between the prime ideals of $S^{-1}R$ and the prime ideals of R which do not meet S. (See [5] p.401, for instance.) Together, these facts give us the required bijective correspondence between the prime ideals of R/I which do not meet S, and the prime ideals of $S^{-1}R/S^{-1}(I)$. □

4.3 Proof of equivalence theorem

Our aim in this subsection is to prove Theorem 4.1.
Proof. It is not too hard to see that conditions (2) and (3) are equivalent. Fix some $f \in I$. We can see $\text{trop}(f)$ as the minimum of the terms $\deg(c_a) + a_1x_1 + \cdots + a_nx_n$ where a takes all values from f, so that $\mathscr{T}(\text{trop}(f))$ consists of points x at which $\deg(c_a) + a_1x_1 + \cdots + a_nx_n$ is minimal for at least two values of a from f. Now, it is clear from the definition of $\text{in}_w(f)$ that it will be a monomial precisely when there is a single term in f of minimal weight. Looking at it the other way, this means $\text{in}_w(f)$ will not be a monomial for any w where the minimum of all weights $\deg(c_a) + a_1w_1 + \cdots + a_nw_n$ is achieved by at least two values of a in f. These conditions are thus equivalent, and so, taking note of Lemma 4.1, conditions (2) and (3) are equivalent.

Now we show that (2) contains (1). We note that (2) is closed, being an intersection of closed sets in $\mathbb{R}^n$. This means that it is enough to show that an arbitrary point $w = (\deg(u_1), \ldots, \deg(u_n))$ in (1) lies in (2). Now, we know by definition that for any $f \in I$, we have $f(u_1, \ldots, u_n) = 0$. That is, $\sum_{a\in\mathscr{A}} c_a u_1^{a_1} \ldots u_n^{a_n} = 0$. If we write

$$c_\alpha u_1^{\alpha_1} \ldots u_n^{\alpha_n} = \sum_{a\in\mathscr{A}\setminus\alpha} -c_a u_1^{a_1} \ldots u_n^{a_n}$$

where we have chosen α such that the valuation of the monomial $c_\alpha u_1^{\alpha_1} \ldots u_n^{\alpha_n}$ is minimal in f, the valuation of both sides must be equal. Since the valuation of a sum is greater than or equal to the minimum valuation of its parts, this means that there must be some monomial $c_\beta u_1^{\beta_1} \dots u_n^{\beta_n}$ in the sum on the right with valuation less than or equal to that of $c_\alpha u_1^{\alpha_1} \dots u_n^{\alpha_n}$. It cannot be less, due to our choice of $c_\alpha u_1^{\alpha_1} \dots u_n^{\alpha_n}$, so there must be at least two terms in f with minimal degree. We showed earlier that

$$\deg(c_a u_1^{a_1} \dots u_n^{a_n}) = \deg(c_a) + a_1 \deg(u_1) + \cdots + a_n \deg(u_n),$$

so this means that $\deg(c_a) + a_1\deg(u_1) + \cdots + a_n\deg(u_n)$ is minimal for at least two values of a from f, and so $w = (\deg(u_1), \dots, \deg(u_n))$ clearly does lie in (2).

We now prove the trickiest inclusion; that (3) is contained in (1). Choose a vector w in (3) such that $w = (\deg(v_1), \dots, \deg(v_n))$ for some $v \in (K^*)^n$. As we noted when we defined deg, its image is dense in $\mathbb{R}$, so we can choose such a vector w arbitrarily close to *any* element of (3). Because of this fact, and the fact that the set (1) is closed, it suffices to prove that $w = (\deg(u_1), \dots, \deg(u_n))$ for some $u \in V(I)$. By making the change of coordinates $x_i = x_i \cdot v_i^{-1}$, we may assume that $w = (0, 0, \dots, 0)$.

We choose a non–zero point $\overline{u} \in V(\text{in}_w(I)) \subset (k^*)^n$. Because k is algebraically closed, Hilbert's Nullstellensatz ensures that this is possible, since if $V(\text{in}_w(I))$ were empty, we could use it to show that $1 \in \text{in}_w(I)$, and in the case that $V(\text{in}_w(I)) = \{0\}$, we could use it to show that $\text{in}_w(I)$ contains monomials; in fact it would contain some power of *any* given monomial, since they all have a zero at the origin. Again because k is algebraically closed, there will be a maximal ideal $\overline{m}$ in $k[x]$ corresponding to this point $\overline{u}$. Let S be the set of polynomials f in $R_K[x]$ whose reduction modulo M_K is not in $\overline{m}$. Then S is a multiplicative set in $R_K[x]$ disjoint from I.
Consider the induced map

$$\varphi : R_K \to S^{-1}R_K[x]/S^{-1}(I \cap R_K[x])$$

We claim that φ is injective. We prove this by contradiction. Assume φ is not injective, and choose some $c \in R_K[x] \setminus \{0\}$ with $\varphi(c) = 0$. This implies the existence of some $f \in S$ such that $cf \in I$. Since c^{-1} exists in K, it follows that $f \in S \cap I$, a contradiction.

Now consider the prime ideals of the ring $S^{-1}R_K[x]/S^{-1}(I \cap R_K[x])$. Combining Lemma 4.2 and Lemma 4.3, we see that each must meet $\varphi(R_K)$ either in $\varphi(R_K)$, in $\varphi(M_K)$ or in $\{0\}$. If they all met either $\varphi(R_K)$ or $\varphi(M_K)$, then since $\varphi(M_K) \subset \varphi(R_K)$, the subset $\varphi(M_K)$ would lie in their intersection, which is of course the nilradical of $S^{-1}R_K[x]/S^{-1}(I \cap R_K[x])$. This is impossible, since φ is injective and the elements of M_K are not nilpotent. So, there exists a prime ideal P of the ring $S^{-1}R_K[x]/S^{-1}(I \cap R_K[x])$ such that $P \cap \varphi(R_K) = \{0\}$.

By Lemma 4.4, the existence of this prime ideal P guarantees we have a proper prime ideal P_1 of $R_K[x]/(I \cap R_K[x]) = (R_K[x] + I)/I$ corresponding to

P, which does not meet S. The image of P_1 in the ring $K[x]/I$ will generate a proper ideal P_2 of that ring, since if there were some element $p \in P_1$ with $k \cdot p$ a unit in $K[x]/I$ for some $k \in K$, then of course $k^{-1} \cdot k \cdot p = p$ would also be a unit in $K[x]/I$ and hence in $(R_K[x] + I)/I$. This ideal P_2 will be contained in some maximal ideal M of $K[x]/I$. If we think of M as a maximal ideal of $K[x]$ which contains I, then since K is algebraically closed, M will correspond to a point in K^n. This point must lie in $V(I)$, since if all polynomials in M evaluate to 0 there, certainly all polynomials of $I \subset M$ must. So M has the form $\langle x_1 - u_1, \ldots, x_n - u_n \rangle$ for some $u \in V(I) \subset (K^*)^n$. Note also that $M \cap R_K[x]$ does not meet S, because there is an injection from the integral domain $(R_K[x] + P_2)/P_2$ to the ring $S^{-1}(R_K[x] + P_2)/S^{-1}(P_2)$, guaranteeing that M corresponds to a proper ideal on the right, and so cannot contain any element of S.

We claim now that $u_i \in R_K$ and furthermore, $u_i \equiv \overline{u_i} \pmod{M_K}$. This will imply that $\deg(u_1) = \deg(u_2) = \cdots = \deg(u_n) = 0$, putting w in (1) and so completing the proof.

Take some $x_i - u_i \in M$. First suppose that $u_i \notin R_K$. Then necessarily, $\deg(u_i) < 0$. In this case $\deg(u_i^{-1}) > 0$, which means $u_i^{-1} \in M_K$. Consider the polynomial $x_i u_i^{-1} - 1 \in M \cap R_K[x]$. The image of this element of M when we mod out by M_K is clearly 1, which is not equal to 0 when evaluated at $\overline{u}$, and so we have an element of $M \cap S$, a contradiction. So we know that $u_i \in R_K$, and hence $x_i - u_i \in M \cap R_K[x]$. Now suppose that $u_i \not\equiv \overline{u_i} \pmod{M_K}$. This means that $x_i - u_i \in S$, since when we evaluate $x_i - u_i \pmod{M_K}$ at $\overline{u_i}$, the result will not be 0. So once again we have an element of $M \cap S$, a contradiction. Thus (3) is contained in (1) and all three definitions are equivalent. $\square$

Speyer and Sturmfels note as a corollary to this theorem that a tropical variety $\mathscr{T}(I)$ is the intersection of the tropical hypersurfaces $\mathscr{T}(\mathrm{trop}(f))$, where the functions f are those in a universal Gröbner basis $\mathcal{U}$. This is a useful fact, because there exist algorithms for computing universal Gröbner bases. ([9] Chapter 3.)

5 The Bieri–Groves Theorem

Our knowledge of Gröbner basis theory and tropical varieties is now sufficient to examine Bernd Sturmfels' proof of the Bieri–Groves Theorem. As its name suggests, the theorem was first proven by Bieri and Groves [2], but the techniques they used were quite different. It is a result about the structure of tropical varieties.

5.1 Preliminaries

We have defined the tropical variety of an ideal $I \in K[x]$ where $K = \overline{\mathbb{C}(t)}$. In the case where the auxiliary variable t does not appear in I, our definition can be simplified. I can be seen as an ideal in $\mathbb{C}[x]$, so that the corresponding variety lies in $(\mathbb{C}^*)^n$. The tropical variety is then equal to the set

$$\tilde{B}(X) = \{w \in \mathbb{R}^n : \mathrm{in}_w(I) \text{ does not contain a monomial}\}$$

This set $\tilde{B}(X)$ is known as the *Bergman fan* of the variety X, where of course X is the variety corresponding to the ideal I. It is in relation to the Bergman fan that Sturmfels, in his book *Solving Systems of Polynomial Equations* [10] proves the Bieri–Groves Theorem.

In his proof, Sturmfels deals with homogeneous ideals and projective varieties. There is however no loss of generality. We can *homogenize* an ideal I, and still recover the initial ideal $\mathrm{in}_w(I)$. By homogenization of a polynomial, we mean the following:

Definition 5.1 (Homogenization). *Let $g = \sum c_i \cdot x^{a_i}$ be a polynomial in $k[x_1, \ldots, x_n]$. Let d be the maximum over i of $a_1 + \cdots + a_n$. Then the homogenization g^d of g is the polynomial in $k[x_0, x_1, \ldots, x_n]$ computed using the formula $g^d = x_0^d \cdot g(\frac{x_1}{x_0}, \cdots, \frac{x_n}{x_0})$.*

Homogenizing an ideal is simply a matter of homogenizing each of the polynomials which generate it. If J is the ideal in $k[x_0, x_1, \ldots, x_n]$ obtained by homogenizing I, then we can recover I by substituting 1 for x_0 in each of the generating polynomials, and once again considering the result as an ideal in $k[x_1, \ldots, x_n]$. We can also recover the initial ideal. For any $w \in \mathbb{R}^n$, the initial ideal $\mathrm{in}_w(I)$ is found from J using the vector $(0, w) \in \mathbb{R}^{n+1}$. We compute the initial ideal $\mathrm{in}_{(0,w)}(J)$ and then replace x_0 with 1. This gives us an alternative definition of the Bergman fan.

$$\tilde{B}(X) = \{w \in \mathbb{R}^n : \mathrm{in}_{(0,w)}(J) \text{ contains no monomial in } S\}$$

Once we have homogenized our ideals, we will see that we can apply the following lemma, which will be crucial to our proof of the theorem.

Lemma 5.1. *Let L be a homogeneous ideal in the polynomial ring S, containing no monomials and $X(L)$ its zero set in the algebraic torus $(\mathbb{C}^*)^n$. Then the following are equivalent:*

1. *Every proper initial ideal of L contains a monomial.*
2. *There exists a subtorus T of $(\mathbb{C}^*)^n$ such that $X(L)$ consists of finitely many T–orbits.*
3. *The Bergman fan $\tilde{B}(X(L))$ is a linear subspace of $\mathbb{R}^n$.*

Proof. We aim to simplify the statement of our lemma before proving it. Let $\mathcal{L}$ denote the linear subspace of $\mathbb{R}^n$ consisting of all vectors w such that $\text{in}_w(L) = L$. Then a non–zero vector $(w_1, \ldots, w_n)$ lies in $\mathcal{L}$ if and only if the one–parameter subgroup $\{t^{w_1}, \ldots, t^{w_n} : t \in \mathbb{C}^*\}$ fixes L; that is, if and only if we have $\tilde{f}(x_1, \ldots, x_n) = f(t^{w_1}x_1, \ldots, t^{w_n}x_n) \in L$ for all $f \in L$. It is not too hard to see why this is the case. If we consider some monomial x^a, then the effect of $(t^{w_1}, \ldots, t^{w_n})$ on x^a is a multiplication by $t^{w \cdot a}$. If f is a polynomial in $\text{in}_w(L)$, the weight $w \cdot a$ will be equal for each term, so multiplying f by $(t^{w_1}, \ldots, t^{w_n})$ will amount to multiplying by $t^{w \cdot a} \in \mathbb{C}$. Thus $(t^{w_1}, \ldots, t^{w_n})$ fixes $\text{in}_w(L)$, and hence L in the case where $\text{in}_w(L) = L$. Now suppose that $(t^{w_1}, \ldots, t^{w_n})$ fixes L. This implies that for every $f \in L$, the polynomials which are formed by each group of terms in f with equal weight are also in L. This is because we can isolate groups of terms with equal weight and then, after a suitable scalar multiplication, cancel them to form new polynomials in the ideal L. The best way to see this is probably by looking at an easy example. Imagine that we know $f = x_1^2 + x_1x_2 + x_2^2$ is in $L \subset k[x_1, x_2]$, and that (t^2, t^1) fixes L. Then we have $t^4x_1^2 + t^3x_1x_2 + t^2x_2^2 \in L$. It follows that

$$\begin{aligned} &t^2(x_1^2 + x_1x_2 + x_2^2) - (t^4x_1^2 + t^3x_1x_2 + t^2x_2^2) \\ &= (t^2 - t^4)x_1^2 + (t^2 - t^3)x_1x_2 \end{aligned}$$

is in L. Multiplying by (t^2, t^1) again, we then arrive at the conclusion that

$$\begin{aligned} &t^3((t^2 - t^4)x_1^2 + (t^2 - t^3)x_1x_2) - (t^4(t^2 - t^4)x_1^2 + t^3(t^2 - t^3)x_1x_2) \\ &= (t^3 - t^4)(t^2 - t^4)x_1^2 \end{aligned}$$

is also in L. It is clear that we can always isolate terms of equal weight in this way. We see immediately that this implies $\text{in}_w(L) \subseteq L$. Furthermore, since polynomials may have only finitely many terms, if we repeatedly cancel out the terms of minimal weight in this way, the process must terminate somewhere, so it also follows that any given $f \in L$ can be written as a finite k–linear combination of terms from $\text{in}_w(L)$. Thus $L \subseteq \text{in}_w(L)$ and so $L = \text{in}_w(L)$.

The subtorus $T \subseteq (\mathbb{C}^*)^n$ which is generated by the one–parameter subgroups $\{(t^{w_1}, \ldots, t^{w_n}) : t \in \mathbb{C}^*\}$ with $w \in \mathcal{L}$ will have the same dimension as $\mathcal{L}$, and it fixes the variety $X(L)$ since it fixes L. Now we are able to simplify our lemma, making it much simpler to prove. We replace $(\mathbb{C}^*)^n$ by its quotient $(\mathbb{C}^*)^n/T$, and $\mathbb{R}^n$ by its quotient $\mathbb{R}^n/\mathcal{L}$, and our lemma reduces to the following easier assertion.

For a homogeneous ideal L which contains no monomial the following are equivalent:

1. *For any non–zero vector w, the initial ideal* $\text{in}_w(L)$ *contains a monomial.*
2. *$X(L)$ is finite.*
3. *$\tilde{B}(X(L)) = \{0\}$.*

Here, the equivalence of the first and third conditions is immediate from the definition of the Bergman fan. The equivalence of the second and third conditions is harder to prove. We shall handle it seperately, as our next Lemma. □

In the proof of this next Lemma, we will make use of yet another definition of a tropical variety, proven by Bergman [1] to be equivalent.

Definition 5.2. *The tropical variety of an ideal I in $k[x^{\pm 1}]$ is the set of n–tuples $(-\nu(x_1), \ldots, -\nu(x_n))$ where ν runs over all real–valued valuations on $k[x^{\pm 1}]/I$ satisfying $\sum \nu(x_j)^2 = 1$.*

The real–valued valuations that Bergman refers to in this theorem are the restrictions to the ring $k[x^{\pm 1}]/I$ of valuations on the field of fractions of the integral domain $k[x^{\pm 1}]/P$, where P is a prime ideal containing I. In addition to the normal axioms of non–Archimedean valuations we have given, they satisfy $\nu(a) = 0$ for $a \in k \setminus \{0\}$. See Bergman's paper for more details.

Lemma 5.2. *A variety $X(L)$ is finite if and only if its Bergman fan $\tilde{B}(X(L)) = \{0\}$.*

Proof. We first present a proof that the variety $X(L)$ is finite if and only if the coordinate ring $\mathbb{C}[X(L)]$ which is associated to it is in fact a finite dimensional $\mathbb{C}$–vector space. ([3] p.232) Assume that $\mathbb{C}[X(L)] = \mathbb{C}[x_1, \ldots, x_n]/L$ is finite dimensional. To show that $X(L)$ is finite, it suffices to show that for each i there can be only finitely many distinct ith coordinates for the points of $X(L)$. Fix i and consider the classes $[x_i^j]$ in $\mathbb{C}[x_1, \ldots, x_n]/L$, where $j = 0, 1, 2, \ldots$. Since $\mathbb{C}[x_1, \ldots, x_n]/L$ is finite dimensional, the $[x_i^j]$ must be linearly dependent in $\mathbb{C}[x_1, \ldots, x_n]/L$. So for some m there exist constants $c_j \in \mathbb{C}$, not all zero, such that

$$\sum_{j=0}^{m} c_j [x_i^j] = [\sum_{j=0}^{m} c_j x_i^j] = [0].$$

This implies that $\sum_{j=0}^{m} c_j x_i^j \in L$. This non–zero polynomial can have only finitely many roots in $\mathbb{C}$, and so there are only finitely many possible values for the ith coordinate of a point in $X(L)$. This shows that $X(L)$ is finite.

Now suppose that $X(L)$ is finite. If $X(L)$ is empty, then $1 \in L$ so the result is obvious. If $X(L)$ is nonempty, then for a fixed i, let $a_j, j = 1, \dots, k$ be the distinct complex numbers occuring as ith coordinates of points in $X(L)$. Form the one variable polynomial $f(x_i) = \prod_{j=1}^{k}(x_i - a_j)$. This polynomial vanishes at every point of $X(L)$, so by Hilbert's Nullstellensatz, some power of it $f^m(x_i)$ must lie in L. If we let $\mathcal{G}$ be a Gröbner basis for L with respect to some $\prec$, this means that $\text{in}_\prec(f^m) = x_i^{km} \in \text{in}_\prec(L)$. So for each i we have some m_i such that $x_i^{m_i} \in \text{in}_\prec(L)$. This implies that there are only finitely many standard monomials in $k[x_i, \dots, x_n]$, since all monomials $x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_n^{\alpha_n}$ where $\alpha_i \geq m_i$ for some i lie in $\text{in}_\prec(L)$. We know (Lemma 3.1) that the standard monomials form a $\mathbb{C}$–vector space basis for $\mathbb{C}[x_i, \dots, x_n]/L$, so this vector space must be finite dimensional.

Our goal now is to show that $\mathbb{C}[X(L)]$ is a finite dimensional vector space if and only if $\tilde{B}(X(L)) = \{0\}$. Assume that $\mathbb{C}[X(L)]$ is a finite dimensional $\mathbb{C}$–vector space. Then (the image of) each variable x_i in $\mathbb{C}[X(L)]$ is algebraic over $\mathbb{C}$. So, we have

$$c_m x_i^m + c_{m+1} x_i^{m+1} + \cdots + c_{n-1} x_i^{n-1} + c_n x_i^n = 0$$

for some m and n such that c_m and c_n are non–zero. If we write $-c_m x_i^m = c_{m+1} x_i^{m+1} + \cdots + c_{n-1} x_i^{n-1} + c_n x_i^n$, and take a valuation ν on both sides, we see that $\nu(-c_m x_i^m) = m \cdot \nu(x_i) \geq \inf((m+1) \cdot \nu(x_i), \dots, n\nu(x_i))$. Therefore $\nu(x_i) \leq 0$. Now we write $c_m x_i^m + c_{m+1} x_i^{m+1} + \cdots + c_{n-1} x_i^{n-1} = -c_n x_i^n$ and again take the valuation of both sides. We see that $n\nu(x_i) \geq \inf(m\nu(x_i), \dots, (n-1)\nu(x_i))$. Thus the only possible value for $\nu(x_i)$ is 0. This holds for all variables, and all valuations, so we can see from Bergman's definition of a tropical variety that $\tilde{B}(X(L)) = \{0\}$.

Now we assume that $\mathbb{C}[X(L)] = \mathbb{C}[x_1, \dots, x_n]/L$ is an infinite dimensional $\mathbb{C}$–vector space and aim to show that $\tilde{B}(X(L))$ must contain some non–zero vector. Since $\mathbb{C}[x_1, \dots, x_n]/L$ is infinite dimensional, there is some prime P containing L such that $\mathbb{C}[x_1, \dots, x_n]/P$ is infinite dimensional. This is a consequence of the fact that $\mathbb{C}[x_1, \dots, x_n]$ is Noetherian. There exists a finite chain of $\mathbb{C}[x_1, \dots, x_n]/L$–modules

$$\{0\} = M_0 \leq M_1 \leq \cdots \leq M_k = \mathbb{C}[x_1, \dots, x_n]/L$$

such that $M_{i+1}/M_i \cong (\mathbb{C}[x_1, \dots, x_n]/L)/P_i$ for each i, with P_i a prime ideal of $\mathbb{C}[x_1, \dots, x_n]/L$ [4]. Since $\mathbb{C}[x_1, \dots, x_n]/L$ is infinite dimensional, at least one of these quotients $(\mathbb{C}[x_1, \dots, x_n]/L)/P_i$ must have infinite dimension. The prime ideal P for which this is the case will clearly correspond to a prime ideal of $\mathbb{C}[x_1, \dots, x_n]$ containing I such that $\mathbb{C}[x_1, \dots, x_n]/P$ is an infinite dimensional k–vector space.

Looking now at $\mathbb{C}[x_1, \ldots, x_n]/P$, which is an integral domain, we see there must be some x_i whose image in $\mathbb{C}[x_1, \ldots, x_n]/P$ is not algebraic. We construct a valuation ν_i on the image of $\mathbb{C}[x_i]$ by defining $\nu_i(\sum c_a x_i^a) = \min(a)$, since a takes only finitely many values from the natural numbers for a given polynomial. It is easy to check that this map satisfies the axioms of a valuation. By definition, $\nu_i(f+g) = \inf(\nu_i(f), \nu_i(g))$, and it is also clear that $\nu_i(fg) = \nu_i(f) + \nu_i(g)$. So we have a non–zero valuation on the image of $\mathbb{C}[x_i]$ in $\mathbb{C}[x_1, \ldots, x_n]/P$. We claim that we can extend this non–zero valuation to a valuation on $\mathbb{C}[x_1, \ldots, x_n]/I$ into the reals, which proves that $\tilde{B}(X(L)) \neq \{0\}$. To substantiate this claim, we quote Theorems 9.11 and 9.12 from [5]. The first of these theorems says that any valuation on a field can be extended to a valuation on an extension field, although the *value group* or image of the valuation may need to be extended. In the case that our extension is transcendental, this is not a problem, because we can assign valuations within the reals to any transcendental element. If our field extension is finite, then we can use the second theorem, which states that a valuation into the reals (non–Archimedean absolute value) can be extended to a valuation into the reals on any finite dimensional extension field. $\square$

There is one further piece of information needed for the proof of our main theorem. We shall claim that given a prime ideal P, which defines an irreducible subvariety of dimension d, the variety defined by $\text{in}_w(P)$, while not necessarily irreducible, must be a union of irreducible varieties of dimension d. We quote without proof Theorem 2 of [6]:

Theorem 5.1. *Let P be a prime ideal in $k[x_1, \ldots, x_n]$ and $w \in \mathbb{N}^n$. Then the initial complex $\Delta_w(P)$ is pure of dimension dim$(P) - 1$ and strongly connected.*

Clearly we cannot use this result without further explanation. First of all, we need to define the objects in the theorem. A subset X of $\{x_1, \ldots, x_n\}$ is *independent modulo I*, if and only if $k[X] \cap I = \{0\}$. The initial complex $\Delta(I)$ is defined as the set

$$\{X \subseteq \{x_1, \ldots, x_n\} : X \text{ is independent modulo } I\}$$

and $\Delta_w(I)$ is defined as $\Delta(\text{in}_w(I))$. The *dimension* of a complex $\Delta(I)$ is the maximum dimension of a face in $\Delta(I)$. A complex is *pure* if all of its maximal faces have the same dimension.

The following properties are shown to follow easily from the definition of $\Delta_w(I)$:

- If $I \subseteq J$ then $\Delta_w(I) \supseteq \Delta_w(J)$.
- $\Delta_w(\text{Rad}(I)) = \Delta_w(I)$.
- $\Delta_w(I \cap J) = \Delta_w(I) \cup \Delta_w(J)$.

If $I = Q_1 \cap \cdots \cap Q_r$ is a primary decomposition of the ideal I, then we can see that

$$\Delta_w(I) = \Delta_w(Q_1) \cup \cdots \cup \Delta_w(Q_r) = \Delta_w(P_1) \cup \cdots \cup \Delta_w(P_r)$$

where the P_i are the associated primes corresponding to each Q_i. (Since each P_i is the radical of the corresponding Q_i). So given that the theorem claims that the complex $\Delta_w P$ is pure of dimension $\dim(P) - 1$, we can see from our last equation that the complex of each of the associated primes of $\mathrm{in}_w(P)$ must have the same dimension $\dim(P)-1$. Kalkbrener and Sturmfels make the further claim that the Krull dimension of an ideal I is in fact the maximum cardinality of a face in $\Delta(I)$, so it follows that the varieties of the associated prime ideals all have dimension $d - 1 + 1 = d$. Since the variety of an ideal is the union of the irreducible varieties of its associated primes, this means that our claim is true.

5.2 The Theorem

Theorem 5.2 (Bieri–Groves theorem). *The Bergman fan $\tilde{B}(X)$ of a d–dimensional irreducible subvariety X of $(\mathbb{C}^*)^n$ is a finite union of rational d–dimensional convex polyhedral cones.*

Proof. It is clear from the definition of the Bergman fan that it is the support of a subfan of the Gröbner fan of the ideal I, and hence a finite union of rational polyhedral cones in $\mathbb{R}^n$. What we need to prove is that the maximal cones occuring in $\tilde{B}(X)$ all have the same dimension d.

If C is a cone in the Gröbner fan, we shall write $\mathrm{in}_C(J)$ for $\mathrm{in}_w(J)$ where w is any vector in the relative interior of C. If X_C is the zero set of the initial ideal $\mathrm{in}_C(J)$ in $(\mathbb{C}^*)^n$, then we have

$$\tilde{B}(X_C) = \mathrm{star}_C\tilde{B}(X) + \mathbb{R} \cdot C,$$

where $\mathrm{star}_C\tilde{B}(X)$ is the fan consisting of all cones in $\tilde{B}(X)$ which contain C. This equality is not immediately obvious.

First we prove the "$\supseteq$" direction. It is clear that $C \subseteq \tilde{B}(X)$. Now note that if some cone D of the Bergman fan contains C, then C is a face of D. Then for any $w' \in C$ and $w \in D$, by Proposition 3.4,

$$\mathrm{in}_w(\mathrm{in}_{w'}(J)) = \mathrm{in}_w(\mathrm{in}_C(J)) = \mathrm{in}_{w'+\epsilon \cdot w}(J) = \mathrm{in}_D(J).$$

Thus since $\mathrm{in}_D(J)$ contains no monomials (D is a cone in the Bergman fan), the ideal $\mathrm{in}_w(\mathrm{in}_C(J))$ contains no monomials for all $w \in D$, and so $D \subseteq \tilde{B}(X_C)$.

Now for the "$\subseteq$" direction. Imagine $\mathrm{in}_w(\mathrm{in}_C(J))$ has no monomials. This means that $\mathrm{in}_w(\mathrm{in}_{w'}(J))$ has no monomials for all $w' \in C$. But

$\mathrm{in}_w(\mathrm{in}_{w'}(J)) = \mathrm{in}_{w'}(\mathrm{in}_w(J))$, so if D is the cone of the Gröbner fan containing w, then D lies in the Bergman fan. Using Lemma 3.4 twice, we must also have $\mathrm{in}_{w+\epsilon\cdot w'}(J) = \mathrm{in}_{w'+\epsilon\cdot w}(J)$, and *this* cone of equivalent vectors lies in the Bergman fan and has both C and D as faces. (Not necessarily properly). So we also have the reverse inclusion $\tilde{B}(X_C) \subseteq \mathrm{star}_C\tilde{B}(X)+\mathbb{R}\cdot C$.

Now let C be a cone of the Gröbner fan of J which is maximal with respect to containment in $\tilde{B}(X)$. Let $L = \mathrm{in}_C(J)$. The projective variety defined by L is not necessarily irreducible, but its components all have the same dimension d as the irreducible variety defined by the prime ideal J. (See Theorem 5.1 and the surrounding discussion.) That is, we have $\dim(X(L)) = \dim(X) = d$.

The maximality of C implies that any proper initial ideal of L will contain a monomial, since if $\mathrm{in}_w(\mathrm{in}_C(J))$ were proper and contained no monomial, then C would be a face of, and hence contained in, the cone of the Gröbner fan D with $\mathrm{in}_{w'+\epsilon\cdot w}(J) = \mathrm{in}_D(J)$, where of course w' is a representative of C. This means that we can apply Lemma 5.1. Considering the proof of the Lemma, we see that $\tilde{B}(X(L))$ has the same dimension as $X(L)$, which we have seen is d. Now since C is a maximal cone in the Gröbner fan of J, $\mathrm{star}_C\tilde{B}(X)$ is empty, so that $\tilde{B}(X(L)) = \tilde{B}(X_C) = \mathbb{R}\cdot C$. We conclude that C has dimension d. Thus all cones of $\tilde{B}(X)$ which are maximal with respect to containment have dimension d, and the proof is complete. □

References

[1] G. Bergman. The Logarithmic Limit–Set of an Algebraic Variety. *Transactions of the American Mathematical Society*, Vol. 157 (Jun., 1971), 459–469.

[2] R. Bieri and J. Groves. The geometry of the set of characters induced by valuations. *Journal für die reine und angewandte Mathematik* **347** (1984), 168–195.

[3] D. Cox, J. Little, D. O'Shea. *Ideals, Varieties, and Algorithms.* Springer Verlag. 1992.

[4] J. Groves. Notes for 620–422 Commutative Algebra. (2005)

[5] N. Jacobson. *Basic Algebra II.* W. H. Freeman and Company. 1980.

[6] M. Kalkbrener and B. Sturmfels. Initial complexes of prime ideals, *Advances in Mathematics* **116** (1995), 365–376.

[7] J. Richter–Gebert, B. Sturmfels, T. Theobald. First Steps in Tropical Geometry. `arXiv:math.AG/0306366`

[8] D. Speyer and B. Sturmfels. The tropical Grassmannian. *Advances in Geometry* **4** (2004), 389–411.

[9] B. Sturmfels. *Gröbner bases and convex polytopes*, volume 8 of *University Lecture Series.* Amer. Math. Soc. 1996.

[10] B. Sturmfels. *Solving systems of polynomial equations*, volume 97 of *CBMS Regional Conference Series in Mathematics.* Amer. Math. Soc. 2002.