Introduction to Tropical Geometry

Diane Maclagan and Bernd Sturmfels

November 4, 2009

Contents

Pr	reface		7												
1	Intr	oduction	9												
	1.1	Arithmetic)												
	1.2	Dynamic Programming	5												
	1.3	Plane Curves)												
	1.4	Amoebas and their Tentacles	7												
	1.5	Implicitization	2												
	1.6	Group Theory	ô												
	1.7	Curve Counting	3												
	1.8	Compactifications	3												
	1.9	Exercises	3												
2	Bui	ding Blocks 57	7												
	2.1	Fields	7												
	2.2	Algebraic Varieties	ô												
	2.3	Polyhedral Geometry	1												
	2.4	Gröbner Bases	ô												
	2.5	Tropical Bases	ô												
	2.6	Exercises)												
3	Tropical Varieties														
	3.1	Tropical Hypersurfaces	3												
	3.2	The Fundamental Theorem	9												
	3.3	The Structure Theorem	õ												
	3.4	Multiplicities and Balancing	9												
	3.5	Connectivity	2												
	3.6	Exercises	5												

CONTENIS	CO	NT	EN	TS
----------	----	----	----	----

4	The	e Tropical Rain Forest 117
	4.1	Plane Curves
	4.2	Surfaces
	4.3	Linear Spaces
	4.4	Grassmannians
	4.5	Complete Intersections
	4.6	Exercises
5	Lin	ear Algebra 127
	5.1	Eigenvalues and Eigenvectors
	5.2	Matroids
	5.3	Tropical Convexity
	5.4	The Rank of a Matrix
	5.5	Dressians
	5.6	Exercises
6	Tor	ic Connections 185
	6.1	Toric Background
	6.2	Subvarieties of Toric Varieties
	6.3	Tropical Compactifications
	6.4	Geometric Tropicalization
	6.5	Degenerations
	6.6	Intersection theory
	6.7	Exercises
7	Elir	nination and Implicitization 189
	7.1	Tropical Maps
	7.2	Projections and Tropical Bases
	7.3	Discriminants and Resultants
	7.4	Mixed Fiber Polytopes
	7.5	Parametrized Surfaces
	7.6	Hadamard Products
	7.7	Exercises
8	Rea	lizability 193
	8.1	Hypersurfaces
	8.2	Matroids
	8.3	Curves

CONTENTS

	$\begin{array}{c} 8.4\\ 8.5\end{array}$	Prevarieties	94 94
9	Fur	ther Topics 19	95
	9.1	Berkovich Spaces	95
	9.2	Abstract Tropical Intersection Theory	95
	9.3	Tropical Curves and Riemann-Roch	95
	9.4	Tropical Moduli Spaces	95

CONTENTS

6

Preface

This is a very preliminary and incomplete draft of the forthcoming textbook on tropical geometry by Diane Maclagan and Bernd Sturmfels. Besides the visible gaps, there are certainly still plenty of mistakes. We will greatly appreciate all of your comments and questions. Please send us an e-mail to D.Maclagan@warwick.ac.uk with a copy to bernd@math.berkeley.edu. In your e-mail, please include the date you see on the title page of this version.

Tropical geometry is a rapidly expanding field, and only a small selection of topics can be covered in an introductory book. In this text we focus on the study of tropical varieties that arise from classical algebraic varieties. Our approach is primarily algebraic and combinatorial. The study of tropical geometry as an intrinsic geometry in its own right will appear in the book by Grisha Mikhalkin [Mik]. We also give less coverage to those topics that have already received a book exposition elsewhere, such as the use of tropical methods in enumerative and real algebraic geometry [IMS07], and applications of the min-plus semiring in the sciences and engineering [BCOQ92, PS05].

This book is intended to be suitable for a class on tropical geometry for beginning graduate students in mathematics. We have attempted to make the first part of the book (Chapters 1–5) accessible to readers with a minimal background in algebraic geometry, say, at the level of the undergraduate text book *Ideals, Varieties, and Algorithms* by Cox, Little, and O'Shea [CLO07].

Later chapters will demand more mathematical maturity and expertise. For instance, Chapter 6 relates tropical geometry to toric geometry, and it will help to have acquaintance with toric varieties, for instance, from having read Fulton's book [Ful93]. Likewise, readers of Chapter 7 will benefit from having had plenty of hands-on experience with Gröbner bases and resultants.

A one-semester course could be based on the first four chapters, plus selected topics from the later chapters. Covering the entire book would require a two semester course, or an exceptionally well-prepared group of students.

CONTENTS

Chapter 1

Introduction

In tropical algebra, the sum of two numbers is their minimum and the product of two number is their product. This algebraic structure is known as the *tropical semiring* or as the min-plus algebra. With minimum replaced by maximum we get the isomorphic max-plus algebra. The adjective "tropical" was coined by French mathematicians, notably Jean-Eric Pin [Pin98], to honor their Brazilian colleague Imre Simon [Sim88], who pioneered the use of min-plus algebra in optimization theory. There is no deeper meaning in the adjective "tropical". It simply stands for the French view of Brazil.

The origins of algebraic geometry lie in the study of zero sets of systems of multivariate polynomials. These objects are algebraic varieties, and they include familiar examples such as plane curves and surfaces in threedimensional space. It makes perfect sense to define polynomials and rational functions over the tropical semiring. The functions they define are piecewiselinear. Also, algebraic varieties can be defined in the tropical setting. They are now subsets of \mathbb{R}^n that are composed of convex polyhedra. Thus, tropical algebraic geometry is a piecewise-linear version of algebraic geometry.

This introductory chapter offers an invitation to tropical mathematics. We present the basic concepts concerning the tropical semiring, we discuss some of the historical origins of tropical geometry, and we show by way of elementary examples how tropical methods can be used to solve problems in algebra, geometry and combinatorics. Proofs, precise definitions, and the general theory will be postponed to later chapters. Our primary objective here is to show the reader that the tropical approach is both useful and fun.

1.1 Arithmetic

Our basic object of study is the tropical semiring $(\mathbb{R} \cup \{\infty\}, \oplus, \odot)$. As a set this is just the real numbers \mathbb{R} , together with an extra element ∞ which represents infinity. In this semiring, the basic arithmetic operations of addition and multiplication of real numbers are redefined as follows:

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

In words, the *tropical sum* of two numbers is their minimum, and the *tropical product* of two numbers is their usual sum. Here are some examples of how to do arithmetic in this exotic number system. The tropical sum of 4 and 9 is 4. The tropical product of 4 and 9 equals 13. We write this as follows:

$$4 \oplus 9 = 4$$
 and $4 \odot 9 = 13$.

Many of the familiar axioms of arithmetic remain valid in tropical mathematics. For instance, both addition and multiplication are *commutative*:

$$x \oplus y = y \oplus x$$
 and $x \odot y = y \odot x$.

These two arithmetic operations are also associative, and the times operator \odot takes precedence when plus \oplus and times \odot occur in the same expression.

The distributive law holds for tropical addition and multiplication:

$$x \odot (y \oplus z) = x \odot y \oplus x \odot z.$$

Here is a numerical example to show distributivity:

$$\begin{array}{rcl} 3 \odot (7 \oplus 11) & = & 3 \odot 7 & = & 10, \\ 3 \odot 7 & \oplus & 3 \odot 11 & = & 10 \oplus & 14 & = & 10. \end{array}$$

Both arithmetic operations have a neutral element. Infinity is the *neutral* element for addition and zero is the *neutral* element for multiplication:

$$x \oplus \infty = x$$
 and $x \odot 0 = 0$.

We also note the following identities involving the two neutral elements:

$$x \odot \infty = \infty$$
 and $x \oplus 0 = \begin{cases} 0 & \text{if } x \ge 0, \\ x & \text{if } x < 0. \end{cases}$

1.1. ARITHMETIC

Elementary school students tend to prefer tropical arithmetic because the multiplication table is easier to memorize, and even long division becomes easy. Here is a tropical *addition table* and a tropical *multiplication table*:

\oplus	1	2	3	4	5	6	7	(\cdot	1	2	3	4	5	6	7
1	1	1	1	1	1	1	1		1	2	3	4	5	6	7	8
2	1	2	2	2	2	2	2		2	3	4	5	6	7	8	9
3	1	2	3	3	3	3	3	:	3	4	5	6	7	8	9	10
4	1	2	3	4	4	4	4	2	4	5	6	7	8	9	10	11
5	1	2	3	4	5	5	5	ł	5	6	$\overline{7}$	8	9	10	11	12
6	1	2	3	4	5	6	6		6	7	8	9	10	11	12	13
7	1	2	3	4	5	6	7	,	7	8	9	10	11	12	13	14

An essential feature of tropical arithmetic is that there is no subtraction. There is no real number x which we can call "13 minus 4" because the equation $4 \oplus x = 13$ has no solution x at all. Tropical division is defined to be classical subtraction, so $(\mathbb{R} \cup \{\infty\}, \oplus, \otimes)$ satisfies all ring (and indeed field) axioms except the existence of additive inverse. Such objects are called semirings, whence the name tropical semiring.

It is extremely important to remember that "0" is the multiplicatively neutral element. For instance, the tropical *Pascal's triangle* looks like this:

The rows of Pascal's triangle are the coefficients appearing in the *Binomial Theorem*. For instance, the third row in the triangle represents the identity

$$\begin{aligned} (x \oplus y)^3 &= & (x \oplus y) \odot (x \oplus y) \odot (x \oplus y) \\ &= & 0 \odot x^3 \oplus 0 \odot x^2 y \oplus 0 \odot x y^2 \oplus 0 \odot y^3. \end{aligned}$$

Of course, the zero coefficients can be dropped in this identity:

$$(x \oplus y)^3 \quad = \quad x^3 \ \oplus \ x^2y \ \oplus \ xy^2 \ \oplus \ y^3.$$

Moreover, the Freshman's Dream holds for all powers in tropical arithmetic:

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

The validity of the three displayed identities is easily verified by noting that the following equations hold in classical arithmetic for all $x, y \in \mathbb{R}$:

$$3 \cdot \min\{x, y\} = \min\{3x, 2x + y, x + 2y, 3y\} = \min\{3x, 3y\}.$$

The linear algebra operations of adding and multiplying vectors and matrices make perfect sense over the tropical semiring. For instance, the tropical scalar product in \mathbb{R}^3 of a row vector with a column vector is the scalar

$$(u_1, u_2, u_3) \odot (v_1, v_2, v_3)^{\mathrm{T}} = u_1 \odot v_1 \oplus u_2 \odot v_2 \oplus u_3 \odot v_3 = \min\{u_1 + v_1, u_2 + v_2, u_3 + v_3\}.$$

Here is the product of a column vector and a row vector of length three:

$$\begin{array}{l} (u_1, u_2, u_3)^T \odot (v_1, v_2, v_3) \\ = \begin{pmatrix} u_1 \odot v_1 & u_1 \odot v_2 & u_1 \odot v_3 \\ u_2 \odot v_1 & u_2 \odot v_2 & u_2 \odot v_3 \\ u_3 \odot v_1 & u_3 \odot v_2 & u_3 \odot v_3 \end{pmatrix} = \begin{pmatrix} u_1 + v_1 & u_1 + v_2 & u_1 + v_3 \\ u_2 + v_1 & u_2 + v_2 & u_2 + v_3 \\ u_3 + v_1 & u_3 + v_2 & u_3 + v_3 \end{pmatrix}$$

Any matrix which can be expressed as such a product has *tropical rank one*.

Here are a few more examples of arithmetic with vectors and matrices:

$$2 \odot (3, -7, 6) = (5, -5, 8), \quad (\infty, 0, 1) \odot (0, 1, \infty)^{T} = 1,$$

$$\begin{pmatrix} 3 & 3 \\ 0 & 7 \end{pmatrix} \oplus \begin{pmatrix} 4 & 1 \\ 5 & 2 \end{pmatrix} = \begin{pmatrix} 3 & 1 \\ 0 & 2 \end{pmatrix} \text{ and } \begin{pmatrix} 3 & 3 \\ 0 & 7 \end{pmatrix} \odot \begin{pmatrix} 4 & 1 \\ 5 & 2 \end{pmatrix} = \begin{pmatrix} 7 & 4 \\ 4 & 1 \end{pmatrix}$$

If we are given a $d \times n$ -matrix A, then we might be interested in computing its image, i.e. the set $\{A \odot x : x \in \mathbb{R}^n\}$, and in solving the linear systems $A \odot x = b$ for various right hand sides b. We will discuss the relevant geometry in Chapter 5. For an introduction to solving tropical linear systems, and to engineering applications thereof, we recommend the book on *Synchronization* and Linearity by Baccelli, Cohen, Olsder and Quadrat [BCOQ92].

Students of computer science and discrete mathematics may encounter tropical matrix multiplication in algorithms for finding shortest paths in graphs and networks. The general framework for such algorithms is known as *dynamic programming*. We shall explore this connection in the next section.

Let x_1, x_2, \ldots, x_n be variables which represent elements in the tropical semiring $(\mathbb{R} \cup \{\infty\}, \oplus, \odot)$. A monomial is any product of these variables,

where repetition is allowed. Throughout this book, we generally allow negative integer exponents. By commutativity, we can sort the product and write monomials in the usual notation, with the variables raised to exponents:

$$x_2 \odot x_1 \odot x_3 \odot x_1 \odot x_4 \odot x_2 \odot x_3 \odot x_2 \quad = \quad x_1^2 x_2^3 x_3^2 x_4.$$

A monomial represents a function from \mathbb{R}^n to \mathbb{R} . When evaluating this function in classical arithmetic, what we get is a linear function:

$$x_2 + x_1 + x_3 + x_1 + x_4 + x_2 + x_3 + x_2 = 2x_1 + 3x_2 + 2x_3 + x_4$$

Remark 1.1.1. Every linear function with integer coefficients arises in this way, so tropical monomials are precisely the linear functions with integer coefficients.

A *tropical polynomial* is a finite linear combination of tropical monomials:

$$p(x_1,\ldots,x_n) = a \odot x_1^{i_1} x_2^{i_2} \cdots x_n^{i_n} \oplus b \odot x_1^{j_1} x_2^{j_2} \cdots x_n^{j_n} \oplus \cdots$$

Here the coefficients a, b, \ldots are real numbers and the exponents i_1, j_1, \ldots are integers. Every tropical polynomial represents a function $\mathbb{R}^n \to \mathbb{R}$. When evaluating this function in classical arithmetic, what we get is the minimum of a finite collection of linear functions, namely

$$p(x_1, \dots, x_n) = \min(a + i_1 x_1 + \dots + i_n x_n, b + j_1 x_1 + \dots + j_n x_n, \dots)$$

This function $p: \mathbb{R}^n \to \mathbb{R}$ has the following three important properties:

- p is continuous,
- p is piecewise-linear, where the number of pieces is finite, and
- p is concave, i.e., $p(\frac{x+y}{2}) \ge \frac{1}{2}(p(x) + p(y))$ for all $x, y \in \mathbb{R}^n$.

Every function which satisfies these three properties can be represented as the minimum of a finite set of linear functions; see Exercise 1.10.1. We conclude:

Lemma 1.1.2. The tropical polynomials in n variables x_1, \ldots, x_n are precisely the piecewise-linear concave functions on \mathbb{R}^n with integer coefficients.

It is instructive to examine tropical polynomials and the functions they define even for polynomials of one variable. For instance, consider the general cubic polynomial in one variable x:

$$p(x) = a \odot x^3 \oplus b \odot x^2 \oplus c \odot x \oplus d.$$
(1.1)

To graph this function we draw four lines in the (x, y) plane: y = 3x + a, y = 2x + b, y = x + c and the horizontal line y = d. The value of p(x) is the smallest y-value such that (x, y) is on one of these four lines, i.e., the graph of p(x) is the lower envelope of the lines. All four lines actually contribute if

$$b-a \leq c-b \leq d-c. \tag{1.2}$$

These three values of x are the breakpoints where p(x) fails to be linear, and the cubic has a corresponding factorization into three linear factors:

$$p(x) = a \odot (x \oplus (b-a)) \odot (x \oplus (c-b)) \odot (x \oplus (d-c))$$

The three breakpoints (1.2) of the graph are called the *roots* of the cubic polynomial p(x). The graph and its breakpoints are shown in Figure 1.1.

Every tropical polynomial function can be written uniquely as a tropical product of tropical linear functions (i.e., the *Fundamental Theorem of Algebra* holds tropically). In this statement we must underline the word "function". Distinct polynomials can represent the same function $p: \mathbb{R}^n \to \mathbb{R}$. We are not claiming that every polynomial factors into linear functions. What we are claiming is that every polynomial can be replaced by an equivalent polynomial, representing the same function, that can be factored into linear factors. Here is an example of a quadratic polynomial function and its unique factorization into linear polynomial functions:

$$x^2 \oplus 17 \odot x \oplus 2 = x^2 \oplus 1 \odot x \oplus 2 = (x \oplus 1)^2.$$

Unique factorization of tropical polynomials holds in one variable, but it no longer holds in two or more variables. What follows is a simple example of a bivariate polynomial that has two distinct irreducible factorizations:

$$\begin{array}{l} (x \oplus 0) \odot (y \oplus 0) \odot (x \odot y \oplus 0) \\ = (x \odot y \oplus x \oplus 0) \odot (x \odot y \oplus y \oplus 0). \end{array}$$

Here is a geometric way of interpreting this identity. The Newton polygon of a polynomial f(x, y) is the convex hull of all points (i, j) such that $x^i y^j$ appears in f(x, y). The Newton polygon of the polynomial above is a hexagon. It is expressed as the sum of two triangles and as the sum of three line segments.



Figure 1.1: The graph of a cubic polynomial and its roots

1.2 Dynamic Programming

To see why tropical arithmetic might be relevant for computer science, let us consider the problem of finding shortest paths in a weighted directed graph. We fix a directed graph G with n nodes that are labeled by $1, 2, \ldots, n$. Every directed edge (i, j) in G has an associated length d_{ij} which is a non-negative real number. If (i, j) is not an edge of G then we set $d_{ij} = +\infty$.

We represent the weighted directed graph G by its $n \times n$ adjacency matrix $D_G = (d_{ij})$ whose off-diagonal entries are the edge lengths d_{ij} . The diagonal entries of D_G are zero, i.e., $d_{ii} = 0$ for all i. The matrix D_G need not be a symmetric matrix, i.e., it may well happen that $d_{ij} \neq d_{ji}$ for some i, j. However, if G is an undirected graph with edge lengths, then we represent G as a directed graph with two directed edges (i, j) and (j, i) for each undirected edge $\{i, j\}$. In that special case, D_G is a symmetric matrix, and we can think of $d_{ij} = d_{ji}$ as the distance between node i and node j. For a general directed graph G, the adjacency matrix D_G will not be symmetric.

Consider the $n \times n$ -matrix with entries in $\mathbb{R}_{\geq 0} \cup \{\infty\}$ that results from

tropically multiplying the given adjacency matrix D_G with itself n-1 times:

$$D_G^{\odot n-1} = D_G \odot D_G \odot \cdots \odot D_G.$$
(1.3)

Proposition 1.2.1. Let G be a weighted directed graph on n nodes with $n \times n$ adjacency matrix D_G . Then the entry of the matrix $D_G^{\odot n-1}$ in row i and column j equals the length of a shortest path from node i to node j in G.

Proof. Let $d_{ij}^{(r)}$ denote the minimum length of any path from node i to node j which uses at most r edges in G. We have $d_{ij}^{(1)} = d_{ij}$ for any two nodes i and j. Since the edge weights d_{ij} were assumed to be non-negative, a shortest path from node i to node j visits each node of G at most once. In particular, any such shortest path in the directed graph G uses at most n-1 directed edges. Hence the length of a shortest path from i to j equals $d_{ij}^{(n-1)}$.

For $r \ge 2$ we have the following recursive formula for the lengths of these shortest paths:

$$d_{ij}^{(r)} = \min \{ d_{ik}^{(r-1)} + d_{kj} : k = 1, 2, \dots, n \}.$$
(1.4)

Using tropical arithmetic, this formula can be rewritten as follows:

$$\begin{aligned} d_{ij}^{(r)} &= d_{i1}^{(r-1)} \odot d_{1j} \oplus d_{i2}^{(r-1)} \odot d_{2j} \oplus \cdots \oplus d_{in}^{(r-1)} \odot d_{nj} \\ &= (d_{i1}^{(r-1)}, d_{i2}^{(r-1)}, \dots, d_{in}^{(r-1)}) \odot (d_{1j}, d_{2j}, \dots, d_{nj})^T. \end{aligned}$$

From this it follows, by induction on r, that $d_{ij}^{(r)}$ coincides with the entry in row i and column j of the $n \times n$ matrix $D_G^{\odot r}$. Indeed, the right hand side of the recursive formula is the tropical product of row i of $D_G^{\odot r-1}$ and column jof D_G , which is the (i, j) entry of $D_G^{\odot r}$. In particular, $d_{ij}^{(n-1)}$ coincides with the entry in row i and column j of $D_G^{\odot n-1}$. This proves the claim. \Box

The iterative evaluation of the formula (1.4) is known as the *Floyd–Warshall algorithm* for finding shortest paths in a weighted digraph. This algorithm and its running time are featured in most standard undergraduate text books on Discrete Mathematics, and it also has a nice Wikipedia page.

For us, running the algorithm means performing the matrix multiplication

$$D_G^{\odot r} = D_G^{\odot r-1} \odot D_G$$
 for $r = 2, \dots, n-1$.

Example 1.2.2. Let G be the complete bidirected graph on n=4 nodes with

$$D_G = \begin{pmatrix} 0 & 1 & 3 & 7 \\ 2 & 0 & 1 & 3 \\ 4 & 5 & 0 & 1 \\ 6 & 3 & 1 & 0 \end{pmatrix}.$$

The first and second tropical power of this matrix are found to be

$$D_G^{\odot 2} = \begin{pmatrix} 0 & 1 & 2 & 4 \\ 2 & 0 & 1 & 2 \\ 4 & 4 & 0 & 1 \\ 5 & 3 & 1 & 0 \end{pmatrix} \quad \text{and} \quad D_G^{\odot 3} = \begin{pmatrix} 0 & 1 & 2 & 3 \\ 2 & 0 & 1 & 2 \\ 4 & 4 & 0 & 1 \\ 5 & 3 & 1 & 0 \end{pmatrix}.$$

The entries in $D_G^{\odot 3}$ are the lengths of the shortest paths in the graph G.

The tropical computation mirrors the following matrix computation in ordinary arithmetic. Let ϵ denote an indeterminate that represents a very small positive real number, and let $A_G(\epsilon)$ be the $n \times n$ matrix whose entry in row *i* and column *j* is the monomial $\epsilon^{d_{ij}}$. In our example we have

$$A_G(\epsilon) = \begin{pmatrix} 1 & \epsilon^1 & \epsilon^3 & \epsilon^7 \\ \epsilon^2 & 1 & \epsilon^1 & \epsilon^3 \\ \epsilon^4 & \epsilon^5 & 1 & \epsilon^1 \\ \epsilon^6 & \epsilon^3 & \epsilon^1 & 1 \end{pmatrix}.$$

Now, we compute the third power of this matrix in ordinary arithmetic:

$$A_G(\epsilon)^3 = \begin{pmatrix} 1+3\epsilon^3+\cdots & 3\epsilon+\epsilon^4+\cdots & 3\epsilon^2+3\epsilon^3+\cdots & \epsilon^3+6\epsilon^4+\cdots\\ 3\epsilon^2+4\epsilon^5+\cdots & 1+3\epsilon^3+\cdots & 3\epsilon+\epsilon^3+\cdots & 3\epsilon^2+3\epsilon^3+\cdots\\ 3\epsilon^4+2\epsilon^6+\cdots & 3\epsilon^4+6\epsilon^5+\cdots & 1+3\epsilon^2+\cdots & 3\epsilon+\epsilon^3+\cdots\\ 6\epsilon^5+3\epsilon^6+\cdots & 3\epsilon^3+\epsilon^5+\cdots & 3\epsilon+\epsilon^3+\cdots & 1+3\epsilon^2+\cdots \end{pmatrix}.$$

The entry of the classical matrix power $A_G(\epsilon)^3$ in row *i* and column *j* is a polynomial in ϵ which represents the lengths of all paths from node *i* to node *j* using at most three edges. The lowest exponent appearing in this polynomial is the (i, j)-entry in the tropical matrix power $D_G^{\odot 3}$.

This is a general phenomenon, summarized informally as follows:

tropical =
$$\lim_{\epsilon \to 0} \log_{\epsilon} (classical(\epsilon))$$
 (1.5)

This process of passing from classical arithmetic to tropical arithmetic is referred to as *tropicalization*. Equation (1.5) is not a mathematical statement. To make this rigorous we use the algebraic notion of *valuations* which will be developed in the subsequent chapters.

We shall give two more examples of how tropical arithmetic ties in naturally with algorithms in discrete mathematics. The first example concerns the dynamic programming approach to *integer linear programming*. The integer linear programming problem can be stated as follows. Let $A = (a_{ij})$ be a $d \times n$ matrix of non-negative integers, let $w = (w_1, \ldots, w_n)$ be a row vector with real entries, and let $b = (b_1, \ldots, b_d)^T$ be a column vector with nonnegative integer entries. Our task is to find a non-negative integer column vector $u = (u_1, \ldots, u_n)$ which solves the following optimization problem:

Minimize
$$w \cdot u$$
 subject to $u \in \mathbb{N}^n$ and $Au = b$. (1.6)

Let us further assume that all columns of the matrix A sum to the same number α and that $b_1 + \cdots + b_d = m\alpha$. This assumption is convenient because it ensures that all feasible solutions $u \in \mathbb{N}^n$ of (1.6) satisfy $u_1 + \cdots + u_n = m$.

We can solve the integer programming problem (1.6) using tropical arithmetic as follows. Let x_1, \ldots, x_d be indeterminates and consider the expression

$$w_1 \odot x_1^{a_{11}} \odot x_2^{a_{21}} \odot \cdots \odot x_d^{a_{d1}} \oplus \cdots \oplus w_n \odot x_1^{a_{1n}} \odot x_2^{a_{2n}} \odot \cdots \odot x_d^{a_{dn}}.$$
(1.7)

Proposition 1.2.3. The optimal value of (1.6) is the coefficient of the monomial $x_1^{b_1} x_2^{b_2} \cdots x_d^{b_d}$ in the mth power of the tropical polynomial (1.7).

The proof of this proposition is not difficult and is similar to that of Proposition 1.2.1. The process of taking the *m*th power of the tropical polynomial (1.7) can be regarded as solving the shortest path problem in a certain graph. This is precisely the dynamic programming approach to integer linear programming. This approach furnishes a polynomial-time algorithm for integer programming under the assumption that the integers in A are bounded.

Example 1.2.4. Let d = 2, n = 5 and consider the instance of (1.6) with

$$A = \begin{pmatrix} 4 & 3 & 2 & 1 & 0 \\ 0 & 1 & 2 & 3 & 4 \end{pmatrix}, \quad b = \begin{pmatrix} 5 \\ 7 \end{pmatrix} \text{ and } w = (2, 5, 11, 7, 3).$$

Here we have $\alpha = 4$ and m = 3. The matrix A and the cost vector w are encoded by a tropical polynomial as in (1.7):

$$f = 2x_1^4 \oplus 5x_1^3x_2 \oplus 11x_1^2x_2^2 \oplus 7x_1x_2^3 \oplus 3x_2^4.$$

The third power of this polynomial, evaluated tropically, is equal to

$$\begin{split} f \odot f \odot f \ &= \ 6x_1^{12} \oplus 9x_1^{11}x_2 \oplus 12x_1^{10}x_2^2 \oplus 11x_1^9x_2^3 \oplus 7x_1^8x_2^4 \oplus 10x_1^7x_2^5 \oplus 13x_1^6x_2^6 \\ & \oplus 12x_1^5x_2^7 \oplus 8x_1^4x_2^8 \oplus 11x_1^3x_2^9 \oplus 17x_1^2x_2^{10} \oplus 13x_1x_2^{11} \oplus 9x_2^{12}. \end{split}$$

1.2. DYNAMIC PROGRAMMING

The coefficient 12 of $x_1^5 x_2^7$ in this tropical polynomial is the optimal value. An optimal solution to this integer programming problem is $u = (1, 0, 0, 1, 1)^{\mathrm{T}}$.

Our final example concerns the notion of the determinant of an $n \times n$ matrix $X = (x_{ij})$. As there is no negation in tropical arithmetic, the *tropical determinant* is the same as the *tropical permanent*, namely, the sum over the diagonal products obtained by taking all n! permutations π of $\{1, 2, \ldots, n\}$:

tropdet(X) :=
$$\bigoplus_{\pi \in S_n} x_{1\pi(1)} \odot x_{2\pi(2)} \odot \cdots \odot x_{n\pi(n)}.$$
 (1.8)

Here S_n is the symmetric group of permutations of $\{1, 2, ..., n\}$. Evaluating the tropical determinant means solving the classical assignment problem of combinatorial optimization. Consider a company which has n jobs and nworkers, and each job needs to be assigned to exactly one of the workers. Let x_{ij} be the cost of assigning job i to worker j. The company wishes to find the cheapest assignment $\pi \in S_n$. The optimal total cost is the minimum:

$$\min\{x_{1\pi(1)} + x_{2\pi(2)} + \dots + x_{n\pi(n)} : \pi \in S_n\}.$$

This number is precisely the tropical determinant of the matrix $Q = (x_{ij})$;

Proposition 1.2.5. The tropical determinant solves the assignment problem.

In the assignment problem we need to find the minimum over n! quantities, which appears to require exponentially many operations. However, there is a well-known polynomial-time algorithm for solving this problem. It was developed by Harold Kuhn in 1955 who called it the *Hungarian Assignment Method* [Kuh55]. This algorithm maintains a price for each job and an (incomplete) assignment of workers and jobs. At each iteration, an unassigned worker is chosen and a shortest augmenting path from this person to the set of jobs is chosen. The total number of arithmetic operations is $O(n^3)$.

In classical arithmetic, the evaluation of determinants and the evaluation of permanents are in different complexity classes. The determinant of an $n \times n$ matrix can be computed in $O(n^3)$ steps, namely by *Gaussian elimination*, while computing the permanent of an $n \times n$ matrix is a fundamentally harder problem. A famous theorem due to Leslie Valiant says that computing the (classical) permanent is #P-complete. In tropical arithmetic, computing the permanent is easier, thanks to the Hungarian Assignment Method. We can think of that method as a certain tropicalization of Gaussian Elimination. For an example, consider a 3×3 matrix $A(\epsilon)$ whose entries are polynomials in the indeterminate ϵ . For each entry we list the term of lowest order:

$$A(\epsilon) = \begin{pmatrix} a_{11}\epsilon^{x_{11}} + \cdots & a_{12}\epsilon^{x_{12}} + \cdots & a_{13}\epsilon^{x_{13}} + \cdots \\ a_{21}\epsilon^{x_{21}} + \cdots & a_{22}\epsilon^{x_{22}} + \cdots & a_{23}\epsilon^{x_{23}} + \cdots \\ a_{31}\epsilon^{x_{31}} + \cdots & a_{32}\epsilon^{x_{32}} + \cdots & a_{33}\epsilon^{x_{33}} + \cdots \end{pmatrix}.$$

Suppose that the a_{ij} are sufficiently general non-zero integers, so that no cancellation occurs in the lowest-order coefficient when we expand the determinant of $A(\epsilon)$. Writing X for the 3×3 matrix with entries x_{ij} , we have

$$\det(A(\epsilon)) = \alpha \cdot \epsilon^{\operatorname{tropdet}(X)} + \cdots \quad \text{for some } \alpha \in \mathbb{R} \setminus \{0\}.$$

Thus the tropical determinant of X can be computed from this expression by taking the logarithm and letting ϵ tend to zero, as suggested by (1.5).

The material in this section is extracted from Chapter 2 in the book Algebraic Statistics for Computational Biology by Lior Pachter and Bernd Sturmfels [PS05]. Algorithms in computational biology that are based on dynamic programming, such as sequence alignment and gene prediction, can be interpreted as evaluating a tropical polynomial. The book [PS05], and the paper [PS04] that preceeds it, argue that the tropical interpretation of dynamic programming algorithms makes sense in the framework of statistics.

1.3 Plane Curves

A tropical polynomial function $p : \mathbb{R}^n \to \mathbb{R}$ is given as the minimum of a finite set of linear functions. We define the *hypersurface* V(p) of p to be the set of all points $x \in \mathbb{R}^n$ at which this minimum is attained at least twice. Equivalently, a point $x \in \mathbb{R}^n$ lies in V(p) if and only if p is not linear at x.

For instance, let n = 1 and let p be the univariate cubic polynomial in (1.1). If the assumption (1.2) holds then

$$V(p) = \{ b - a, c - b, d - c \}.$$

Thus the hypersurface V(p) is the set of "roots" of the polynomial p(x).

For an example of a tropical polynomial in many variables consider the determinant function p = tropdet in (1.8). Its hypersurface V(p) consists of all $n \times n$ -matrices that are *tropically singular*. A square matrix being tropically singular means that the optimal solution to the assignment problem is

20

not unique, so among all n! ways of assigning the n workers to the n jobs, there are at least two assignments both of which minimize the total cost.

In this section we study the geometry of a polynomial in two variables:

$$p(x,y) = \bigoplus_{(i,j)} c_{ij} \odot x^i \odot y^j.$$

The corresponding tropical hypersurface V(p) is a *plane tropical curve*. The following proposition summarizes the salient features of such a tropical curve.

Proposition 1.3.1. The curve V(p) is a finite graph which is embedded in the plane \mathbb{R}^2 . It has both bounded and unbounded edges, all edge slopes are rational, and this graph satisfies a zero tension condition around each node.

The zero tension condition is the following geometric condition. Consider any node (x, y) of the graph and suppose it is the origin, i.e., (x, y) = (0, 0). Then the edges adjacent to this node lie on lines with rational slopes. On each such ray emanating from the origin consider the first non-zero lattice vector. Zero tension at (x, y) means that the sum of these vectors is zero.

Our first example is a *line* in the plane. It is defined by a polynomial:

$$p(x,y) = a \odot x \oplus b \odot y \oplus c$$
 where $a, b, c \in \mathbb{R}$.

The tropical curve V(p) consists of all points (x, y) where the function

$$p : \mathbb{R}^2 \to \mathbb{R}, \quad (x,y) \mapsto \min(a+x, b+y, c)$$

is not linear. It consists of three half-rays emanating from the point (x, y) = (c - a, c - b) into northern, eastern and southwestern direction.

Two lines in the tropical plane will always meet in one point. This is shown in Figure 1.2. When the lines are in special position, it can happen that the set-theoretic intersection is a halfray, and in that case the notion of stable intersection discussed below is used to get a unique intersection point.

Let p be any tropical polynomial in x and y and consider any term $\gamma \odot x^i \odot y^j$ appearing in p. In classical arithmetic this represents the linear function $(x, y) \mapsto \gamma + ix + jy$, and the tropical polynomial function $p : \mathbb{R}^2 \to \mathbb{R}$ is given by the minimum of these linear functions. The graph of p is concave and piecewise linear. It looks like a tent over the plane \mathbb{R}^2 . The tropical curve V(p) is the set of all points in \mathbb{R}^2 over which the graph is not smooth.



Figure 1.2: Two lines in the tropical plane meet in one point



Figure 1.3: The graph and the curve defined by a quadratic polynomial



Figure 1.4: Two subdivisions of the Newton polygon of a biquadratic curve

As an example we consider the general quadratic polynomial

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

Suppose that the coefficients $a, b, c, d, e, f \in \mathbb{R}$ satisfy the inequalities

$$2b < a + c$$
, $2d < a + f$, $2e < c + f$.

Then the graph of $p : \mathbb{R}^2 \to \mathbb{R}$ is the lower envelope of six planes in \mathbb{R}^3 . This is shown in Figure 1.3, where each linear piece of the graph is labeled by the corresponding linear function. Below this "tent" lies the tropical quadratic curve $V(p) \subset \mathbb{R}^2$. This curve has four vertices, three bounded edges and six half-rays (two northern, two eastern and two southwestern).

If p is a tropical polynomial then its curve V(p) is a planar graph dual to the graph of a regular subdivision of its Newton polygon Newt(p). Such a subdivision is a *unimodular triangulation* if each cell is a lattice triangle of unit area 1/2. In this case we call V(p) a smooth tropical curve.

The unbounded rays of a tropical curve V(p) are perpendicular to the edges of the Newton polygon. For example, if p is a biquadratic polynomial then Newt(p) is the square with vertices (0,0), (0,2), (2,0), (2,2), and V(p) has two unbounded rays for each of the four edges of the square. Figure 1.4 shows two subdivisions. The corresponding tropical curves are shown in Figure 1.5. The curve on the left is smooth, and it has genus one. The unique cycle corresponds to the interior lattice point of Newt(p). This is an example of a tropical elliptic curve. The curve on the right is not smooth.

If we draw tropical curves in the plane, then we discover that they intersect and interpolate just like algebraic curves do. In particular, we observe:



Figure 1.5: Two tropical biquadratic curves. The curve on the left is smooth.

- Two general lines meet in one points (Figure 1.2).
- Two general points lie on a unique line.
- A line and a quadric meet in two points (Figure 1.7).
- Two quadrics meet in four points (Figures 1.6 and 1.8).
- Five general points lie on a unique quadric (Exercise 1.4).

A classical result from algebraic geometry, known as *Bézout's Theorem*, holds in tropical algebraic geometry as well. In order to state this theorem, we need to introduce multiplicities. First of all, every edge of a tropical curve comes with an attached multiplicity which is a positive integer. For any point (x, y) in the relative interior of any edge, consider the terms $\gamma \odot x^i \odot y^j$ which attain the minimum. The sum of these terms is effectively a polynomial in one variable, and the number of nonzero roots of that polynomial equals the lattice length of the edge in question. Next, we consider any two lines with distinct rational slopes in \mathbb{R}^2 . If their primitive direction vectors are $(u_1, u_2) \in \mathbb{Z}^2$ and $(v_1, v_2) \in \mathbb{Z}^2$ respectively, then the intersection multiplicity of the two lines at their unique common point is the determinant $|u_1v_2-u_2v_1|$.

We now focus on tropical curves whose Newton polygons are the standard triangles, with vertices (0,0), (0,d) and (d,0). We refer to such a curve as a *curve of degree d*. A curve of degree *d* has *d* rays, possibly counting multiplicities, perpendicular to each of the three edges of its Newton triangle.



Figure 1.6: Bézout's Theorem: Two quadratic curves meet in four points



Figure 1.7: The stable intersection of a line and a conic

Suppose that C and D are two tropical curves in \mathbb{R}^2 that intersect transversally, that is, every common point lies in the relative interior of a unique edge in C and also in D. The multiplicity of that point is the product of the multiplicities of the edges times the intersection multiplicity $|u_1v_2 - u_2v_1|$.

Theorem 1.3.2 (Bézout). Consider two tropical curves C and D of degree c and d in \mathbb{R}^2 . If the two curves intersect transversally, then the number of intersection points, counted with multiplicities as above, is equal to $c \cdot d$.

Just like in classical algebraic geometry, it is possible to remove the restriction "intersect transversally" from the statement of Bézout's Theorem. In fact, the situation is even better here because of the following important phenomenon which is unfamiliar from classical geometry, namely, intersections can be continued across the entire parameter space of coefficients.

We explain this for the intersection of two curves C and D of degrees cand d in \mathbb{R}^2 . Suppose the intersection of C and D is not transverse or not even finite. Pick **any** nearby curves C_{ϵ} and D_{ϵ} such that C_{ϵ} and D_{ϵ} intersect transversely in finitely many points. Then, according to the refined count of Theorem 1.3.2, the intersection $C_{\epsilon} \cap D_{\epsilon}$ is a multiset of cardinality $c \cdot d$.

Theorem 1.3.3 (Stable Intersection Principle). The limit of the point configuration $C_{\epsilon} \cap D_{\epsilon}$ is independent of the choice of perturbations. It is a well-defined multiset of $c \cdot d$ points contained in the intersection $C \cap D$.

Here the limit is taken as ϵ tends to 0. Multiplicities add up when points collide. The limit is a finite configuration of point in \mathbb{R}^2 with multiplici-



Figure 1.8: The stable intersection of a conic with itself.

ties, where the sum of the multiplicities is cd. We call this limit the *stable* intersection of the curves C and D, and we denote this multiset of points by

 $C \cap_{\mathrm{st}} D = \lim_{\epsilon \to 0} (C_{\epsilon} \cap D_{\epsilon}).$

Hence we can strengthen the statement of Bézout's Theorem as follows.

Corollary 1.3.4. Any two curves of degrees c and d in \mathbb{R}^2 , no matter how special they might be, intersect stably in a well-defined multiset of cd points.

The Stable Intersection Principle is illustrated in Figures 1.7 and 1.8. In Figure 1.7 we see the intersection of a tropical line with a tropical conic, moving from general position to special position. In the diagram on the right, the set-theoretic intersection of the two curves is infinite, but the stable intersection is well-defined. It consists of precisely the two points A and B.

Figure 1.8 shows an even more dramatic situation. In that picture, a conic is intersected stably with itself. For any small perturbation the coefficients of the defining quadratic polynomials, we obtain four intersection points that lie near the four nodes of the original conic. This shows that the stable intersection of a conic with itself consists precisely of the four nodes.

1.4 Amoebas and their Tentacles

One of the earliest sources in tropical algebraic geometry is the 1971 paper by George Bergman [Ber71] on the *logarithmic limit of an algebraic variety*. With hindsight, the structure introduced by Bergman is the same as the tropical variety arising from a subvariety in a complex algebraic torus $(\mathbb{C}^*)^n$. The *amoeba* of such a variety is its image under taking the coordinatewise logarithm of the absolute value of any point on the variety. The term amoeba was coined by Gel'fand, Kapranov and Zelevinsky in their monograph on *Discriminants*, *Resultants*, and *Multidimensional Determinants* [GKZ08]. Bergman's logarithmic limit arises from the amoeba as the set of all tentacle directions. In this section we review material from these and related sources, and we discuss what the authors had in mind.

Let *I* be an ideal in the Laurent polynomial ring $S = \mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}, \dots, x_n^{\pm 1}]$. Its algebraic variety is the common zero set of all Laurent polynomials in *I*:

$$V(I) = \left\{ z \in (\mathbb{C}^*)^n : f(z) = 0 \text{ for all } f \in I \right\}.$$

Note that this is well-defined because $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$. The *amoeba* of the ideal I is the subset of \mathbb{R}^n defined as image of the coordinate-wise logarithm map:

$$\mathcal{A}(I) = \left\{ \left(\log(|z_1|), \log(|z_2|), \dots, \log(|z_n|) \right) \in \mathbb{R}^n : z = (z_1, \dots, z_n) \in V(I) \right\}$$

If n = 1 and I is a proper ideal in $S = \mathbb{C}[x, x^{-1}]$ then I is a principal ideal, which is generated by a single polynomial f(x) that factors over \mathbb{C} :

$$f(x) = (u_1 + iv_1 - x)(u_2 + iv_2 - x) \cdots (u_m + iv_m - x).$$

Here $u_1, v_1, \ldots, u_m, v_m \in \mathbb{R}$ are the real and imaginary parts of the various roots of f(x), and the amoeba is the following set of $\leq m$ real numbers:

$$\mathcal{A}(I) = \mathcal{A}(f) = \left\{ \log(\sqrt{u_1^2 + v_1^2}), \log(\sqrt{u_2^2 + v_2^2}), \dots, \log(\sqrt{u_m^2 + v_m^2}) \right\}.$$

The name "amoeba" begins to make more sense once one examines the case when n = 2. Suppose that $I = \langle f(x_1, x_2) \rangle$ is the principal ideal of a curve $\{f(x_1, x_2) = 0\}$ in $(\mathbb{C}^*)^2$. The amoeba $\mathcal{A}(f)$ of that curve is a closed subset of \mathbb{R}^2 whose boundary is described by analytic functions. It has finitely many tentacles that emenate towards infinity, and the directions of these tentacles are precisely the directions perpendicular to the edges of the Newton polygon Newt(f). The complement $\mathbb{R}^2 \setminus \mathcal{A}(f)$ of the amoeba is a finite union of open convex subsets of the plane \mathbb{R}^2 . Pictures of amoebas of curves and surfaces are supposed to like their biological counterparts.

We refer to the work of Passare and his collaborators [PR04, PT05] for the foundational results on amoebas of hypersurfaces in $(\mathbb{C}^*)^n$, and to the article by Theobald [The02] for methods for computing and drawing amoebas. An interesting Nullstellensatz for amoebas was established by Purbhoo [Pur08].



Figure 1.9: The amoeba of a plane curve and its spine

Example 1.4.1. Figure 1.9 shows the complex amoeba of the curve

$$f(z,w) = 1 + 5zw + w^2 - z^3 + 3z^2w - z^2w^2.$$

This picture uses the max-convention under which logarithms are negated. (Passare tells us that he prefers max-convention because his son's first name is Max and not Min). So, what is depicted in Figure 1.9 is the set

$$-\mathcal{A}(f) = \{ (-\log(|z|), -\log(|w|)) \in \mathbb{R}^2 : z, w \in \mathbb{C}^* \text{ and } f(z, w) = 0 \}.$$

Note the two bounded convex components in the complement of $\mathcal{A}(f)$. They correspond to the two interior lattice points of the Newton polygon of f. The tentacles of the amoeba converge to four rays in \mathbb{R}^2 , and the union of these rays is precisely the plane curve V(p) defined by the tropical polynomial

$$p = \operatorname{trop}(f) = 0 \oplus u \odot v \oplus v^2 \oplus u^3 \oplus u^2 \odot v \oplus u^2 \odot v^2$$

This expression is the tropicalization of f. All coefficients of p are zero because the coefficients of f are complex numbers. There are no parameters.

Inside the amoeba of Figure 1.9, we see a tropical curve V(q) defined by a specific tropical polynomial whose coefficients c_1, \ldots, c_6 are nonzero reals:

$$q = c_1 \oplus c_2 \odot u \odot v \oplus c_3 \odot v^2 \oplus c_4 \odot u^3 \oplus c_5 \odot u^2 \odot v \oplus c_6 \odot u^2 \odot v^2$$

The tropical curve V(q) is a canonical deformation retract of $\mathcal{A}(f)$. It is known as the *spine* of the amoeba. The coefficients c_i are defined below. \Box

There are three ways in which tropical varieties arise from amoebas. They are different and we associate the name of a mathematician with each of them.

The Passare Construction: Every complex hypersurface amoeba $\mathcal{A}(f)$ has a spine which is a canonical tropical hypersurface contained in $\mathcal{A}(f)$. Suppose f = f(z, w) is a polynomial in two variables. Then its Ronkin function is

$$N_f(u,v) = \frac{1}{(2\pi i)^n} \int_{\log^{-1}(u,v)} \log |f(z,w)| \frac{dz}{z} \wedge \frac{dw}{w}.$$

Passare and Rullgard [PR04] showed that this function is concave, and that it is linear on each connected component of the complement of $\mathcal{A}(f)$. Let q(u, v) be denote the minimum of these affine-linear functions, one for each component in the amoeba complement. Then q(u, v) is a tropical polynomial

1.4. AMOEBAS AND THEIR TENTACLES

function (i.e. piecewise-linear concave function) which satisfies $N_f(u, v) \leq q(u, v)$ for all $(u, v) \in \mathbb{R}^2$. Its tropical curve V(q) is the spine of $\mathcal{A}(f)$.

The Maslov Construction: Tropical varieties arise as limits of amoebas as one changes the base of the logarithm and makes it either very large or very small. This limit process is also known as Maslov dequantization, and it can be made precise as follows. Given h > 0, we redefine arithmetic as follows:

$$x \oplus_h y = h \cdot \log\left(\exp\left(\frac{x}{h}\right) + \exp\left(\frac{y}{h}\right)\right)$$
 and $x \odot_h y = x + y.$

This is what happens to the operations addition and multiplication of positive real numbers under the coordinate transformation $\mathbb{R}_+ \to \mathbb{R}, x \mapsto h \cdot \log(x)$.

We now consider a polynomial $f_h(z, w)$ whose coefficients are rational functions of the parameter h. For each h > 0, we take the amoeba $\mathcal{A}_h(f_h)$ of f_h with respect to scaled logarithm map $(z, w) \mapsto h \cdot (\log(|z|), \log(|w|))$. The limit in the Hausdorff topology of the set $\mathcal{A}_h(f_h)$ as $h \to 0+$ is a tropical hypersurface V(q). For details see [Mik04]. The coefficients of the tropical polynomials q are the orders (of poles or zeros) of the coefficients at t = 0. This process can be thought of as a sequence of amoebas converging to their spine but it is quite different from the construction using Ronkin functions.

The Bergman Construction: Our third connection between amoebas on tropical varieties arises by examining their tentacles. Here we disregard the interior structure of $\mathcal{A}(f)$, such as the bounded convex regions in the complement, and we focus only on the asymptotic limit directions. This makes sense for any subvariety of $(\mathbb{C}^*)^n$, so our input now is an ideal $I \subset S$ as above.

We denote the unit sphere by $\mathbb{S}^{n-1} = \{x \in \mathbb{R}^n : ||x|| = 1\}$. For any real number M > 0 we consider the following scaled subset of the amoeba:

$$\mathcal{A}_M(I) = \frac{1}{M} \mathcal{A}(I) \cap \mathbb{S}^{n-1}.$$

The logarithmic limit set $\mathcal{A}_{\infty}(I)$ is the set of all points v on the sphere \mathbb{S}^{n-1} such that there exists a sequence of points $v_M \in \mathcal{A}_M(I)$ converging to v, i.e.,

$$\lim_{M\to\infty} v_M = v.$$

The following result establishes the relationship to the tropical variety $\operatorname{trop}(I)$ of I. Here $\operatorname{trop}(I)$ is defined as the intersection of the tropical hypersurfaces V(p) where $p = \operatorname{trop}(f)$ is the tropicalization of any polynomial $f \in I$.

Theorem 1.4.2. The tropical variety of I coincides with the cone over the logarithmic limit set $\mathcal{A}_{\infty}(I)$, i.e., a non-zero vector $w \in \mathbb{R}^n$ lies in trop(I) if and only if the corresponding unit vector $\frac{1}{||w||}w$ lies in $\mathcal{A}_{\infty}(I)$.

In Chapter 3 we will show that V(I) is a polyhedral fan, and we shall establish various structural properties of that fan. Theorem 1.4.2 and the fan property of V(I) implies that $\mathcal{A}_{\infty}(I)$ is a spherical polyhedral complex.

It is interesting to see the motivation behind the paper [Ber71]. Bergman introduced tropical varieties in order to prove a conjecture of Zalessky concerning the multiplicative action of $GL(n,\mathbb{Z})$ on the Laurent polynomial ring S. Here, an invertible integer matrix $g = (g_{ij})$ acts on S as the ring homomorphism that maps each variable x_i to the Laurent monomial $\prod_{j=1}^n x_j^{g_{ij}}$.

If I is a proper ideal in S then we consider its stabilizer subgroup:

$$\operatorname{Stab}(I) = \left\{ g \in GL(n, \mathbb{Z}) : gI = I \right\}.$$

The following result answers Zalessky's question. It is Theorem 1 in [Ber71]:

Corollary 1.4.3. The stabilizer $\operatorname{Stab}(I) \subset GL(n, \mathbb{Z})$ of a proper ideal $I \subset S$ has a subgroup of finite index which stabilizes a nontrivial sublattice of \mathbb{Z}^n .

Proof. The tropical variety V(I) has the structure of a proper polyhedral fan in \mathbb{R}^n . Let \mathcal{U} be the finite set of linear subspaces of \mathbb{R}^n that are spanned by the maximal cones in V(I). While the fan structure on V(I) is not unique, the set \mathcal{U} of linear subspaces of \mathbb{R}^n is uniquely determined by I. The set \mathcal{U} does not change under refinement or coarsening of the fan structure on V(I).

The group $\operatorname{Stab}(I)$ acts by linear transformations on \mathbb{R}^n , and it leaves the tropical variety V(I) invariant. This implies that it acts by permutations on the finite set \mathcal{U} of subspaces in \mathbb{R}^n . Fix one particular subspace $U \in \mathcal{U}$ and let G be the subgroup of all elements $g \in \operatorname{Stab}(I)$ that fix U. Then G has finite index in $\operatorname{Stab}(I)$ and it stabilizes the sublattice $U \cap \mathbb{Z}^n$ of \mathbb{Z}^n . \Box

1.5 Implicitization

An algebraic variety can be represented either as the image of a rational map or as the zero set of some multivariate polynomials. The latter representation exists for all algebraic varieties while the former representation requires that the variety be *unirational*, which is a very special property in algebraic geometry. The transition between two representations is a basic problem in computer algebra. Implicitization is the problem of passing from the first representation to the second, that is, given a rational map Φ , one seeks to determine the prime ideal of all polynomials that vanish on the image of Φ .

In this section we examine the simplest instance, namely, we consider the case of a plane curve in \mathbb{C}^2 that is given by a rational parametrization:

$$\Phi : \mathbb{C} \to \mathbb{C}^2, t \mapsto (\phi_1(t), \phi_2(t)).$$

To make the map Φ actually well-defined, we here tacitly assume that the poles of ϕ_1 and ϕ_2 have been removed from the domain \mathbb{C} . The implicitization problem is to compute the unique (up to scaling) irreducible polynomial f(x, y) vanishing on the curve $\operatorname{Image}(\Phi) = \{(\phi_1(t), \phi_2(t)) \in \mathbb{C}^2 : t \in \mathbb{C}\}.$

Example 1.5.1. Consider the plane curve defined parametrically by

$$\Phi(t) = \left(\frac{t^3 + 4t^2 + 4t}{t^2 - 1}, \frac{t^3 - t^2 - t + 1}{t^2}\right).$$

The implicit equation of this curve equals

$$f(x,y) = x^{3}y^{2} - x^{2}y^{3} - 5x^{2}y^{2} - 2x^{2}y - 4xy^{2} - 33xy + 16y^{2} + 72y + 81.$$

Two standard methods used in computer algebra for solving implicitization problems are Gröbner bases and resultants. These methods are explained in the text book by Cox, Little and O'Shea [CLO07]. Specifically, the desired polynomial f(x, y) equals the Sylvester resultant of the numerator of $x - \phi_1(t)$ and the numerator of $y - \phi_2(t)$ with respect to variable t. For instance, the implicit equation in Example 1.5.1 is easily found as follows:

$$f(x,y) = \operatorname{resultant}_t \left(t^3 + 4t^2 + 4t - (t^2 - 1)x, t^3 - t^2 - t + 1 - t^2y \right).$$

However, for larger problems in higher dimensions, Gröbner bases and resultants often do not perform well enough, or do not give enough geometric insight. This is where the tropical approach to implicatization comes in. We shall explain the basic idea behind this approach for the case of plane curves.

The Newton polygon of the implicit equation f(x, y) is the convex hull in \mathbb{R}^2 of all points $(i, j) \in \mathbb{Z}^2$ such that $x^i y^j$ appears with non-zero coefficient in the expansion of f(x, y). We denote the Newton polygon of f by Newt(f). For example, the Newton polygon of the polynomial above equals

Newt
$$(f) = \operatorname{Conv}\left\{ \begin{pmatrix} 0\\0 \end{pmatrix}, \begin{pmatrix} 2\\1 \end{pmatrix}, \begin{pmatrix} 3\\2 \end{pmatrix}, \begin{pmatrix} 2\\3 \end{pmatrix}, \begin{pmatrix} 2\\1 \end{pmatrix} \right\}$$
 (1.9)

This pentagon has four additional lattice points in its interior, so Newt(f) contains precisely nine lattice points, one for each of the nine terms of f(x, y).

Suppose we are given the parametrization Φ and suppose that the implicit equation f(x, y) is unknown and hard to get. Let us further assume that the Newton polygon the implicit equation is known. That information reveals

$$f(x,y) = c_1 x^3 y^2 + c_2 x^2 y^3 + c_3 x^2 y^2 + c_4 x^2 y + c_5 x y^2 + c_6 x y + c_7 y^2 + c_8 y + c_9 y^2 + c_9 y^2 + c_8 y + c_9 y^2 + c_8 y^2 + c_8 y + c_9 y^2 + c_8 y^2 + c_8 y^2 + c_8 y^2 + c_8 y^2 + c_9 y^2 + c_8 y^2 + c_9 y^2 +$$

where the coefficients c_1, c_2, \ldots, c_9 are unknown parameters. At this point we can set up a linear system of equations as follows. For any choice of complex number τ , the equation $f(\phi_1(\tau), \phi_2(\tau)) = 0$ holds. This equation translates into one linear equation for the nine unknowns c_i . Eight of such linear equations will determine the coefficients uniquely (up to scaling). For instance, in our example, if we take $\tau = \pm 2, \pm 3, \pm 4, \pm 5$, then we get eight linear equations which stipulate that the column vector $(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8, c_9)^T$ lies in the kernel of the following 8×9-matrix of rational numbers

This matrix has rank 8, so its kernel is 1-dimensional. Any generator of that kernel translates into a scalar multiple of the polynomial f(x, y). From this example we see that the implicit equation f(x, y) can be recovered using (numerical) linear algebra from the Newton polytope Newt(f), but it also suggests that the matrices in the resulting systems of linear equations tend to be dense and ill-conditioned. This means that it is rather non-trivial computational problem to solve the equations when f(x, y) has thousands of terms. However, from a geometric perspective it makes sense to consider the implicitization problem solved once the Newton polytope has been found. Thus, in what follows, we consider the following alternative version of implicitization:

Tropical Implicitization Problem: Given two rational functions $\phi_1(t)$ and $\phi_2(t)$, compute the Newton polytope Newt(f) of the implicit equation f(x, y).

1.5. IMPLICITIZATION

We shall present the solution to the tropical implicitization problem for plane curves. By the Fundamental Theorem of Algebra, the two given rational functions are products of linear factors over the complex numbers \mathbb{C} :

$$\begin{aligned}
\phi_1(t) &= (t - \alpha_1)^{u_1} (t - \alpha_2)^{u_2} \cdots (t - \alpha_1)^{u_m} \\
\phi_2(t) &= (t - \alpha_1)^{v_1} (t - \alpha_2)^{v_2} \cdots (t - \alpha_1)^{v_m}
\end{aligned} \tag{1.10}$$

Here the α_i are the zeros and poles of either the two functions ϕ_1 and ϕ_2 . It may occur that u_i is zero while v_j is non-zero or vice versa. For what follows we do not need the algebraic numbers α_i but only the exponents u_i and v_j occuring in the factorizations. These exponents can be computed using symbolic algorithms, such as the Euclidean algorithm, and no field extensions or floating pointing computations are needed to find the integers u_i and v_j .

We abbreviate $u_0 = -u_1 - u_2 - \cdots - u_m$ and $v_0 = -v_1 - v_2 - \cdots - v_m$, and we consider the following collection of m+1 integer vectors in the plane:

$$\begin{pmatrix} u_0 \\ v_0 \end{pmatrix}, \begin{pmatrix} u_1 \\ v_1 \end{pmatrix}, \begin{pmatrix} u_2 \\ v_2 \end{pmatrix}, \dots, \begin{pmatrix} u_m \\ v_m \end{pmatrix}.$$
 (1.11)

We consider the rays spanned by these m+1 vectors. Each ray has a natural multiplicity namely the sum of the lattice lengths of all vectors $(u_i, v_i)^T$ lying on that ray. Since the vectors in (1.11) sum to zero, this configuration of rays satisfies the balancing condition: it is a tropical curve in the plane \mathbb{R}^2 .

The following result is an immediate consequence of the Fundamental Theorem of Tropical Geometry which will be stated and proved in Chapter 3:

Theorem 1.5.2. The tropical curve V(f) defined by the unknown polynomial coincides with the tropical curve determined by the vectors in (1.11).

The Newton polygon Newt(f) can be recovered from the tropical curve V(f) as follows. The first step is to rotate our vectors by 90 degrees:

$$\begin{pmatrix} v_0 \\ -u_0 \end{pmatrix}, \begin{pmatrix} v_1 \\ -u_1 \end{pmatrix}, \begin{pmatrix} v_2 \\ -u_2 \end{pmatrix}, \dots, \begin{pmatrix} v_m \\ -u_m \end{pmatrix}.$$
 (1.12)

Since these vectors sum to zero, there exists a convex polygon P whose edges are translates of these vectors. We construct P by sorting the vectors by increasing slope and then simply concatenating them. The convex polygon P is unique up to translation. Hence there exists a unique translate P^+ of the polygon P which lies in the non-negative orthant $\mathbb{R}_{>0}$ and which has non-empty intersection with both the x-axis and the y-axis. The latter requirements are necessary (and sufficient) for a convex polygon to arise as the Newton polygon of an irreducible polynomial in $\mathbb{C}[x, y]$. We conclude:

Corollary 1.5.3. The polygon P^+ coincides with the Newton polygon Newt(f) of the defining irreducible polynomial of the parametrized curve Image (Φ) .

This solves the tropical implicitization problem for plane curves over \mathbb{C} . We illustrate this solution for our running example.

Example 1.5.4. We write the map of Example 1.5.1 in factored form (1.10):

$$\phi_1(t) = (t-1)^{-1} t^1 (t+1)^{-1} (t+2)^2, \phi_2(t) = (t-1)^2 t^{-2} (t+1)^1 (t+2)^0.$$

The derived configuration of five vectors as in (1.11) equals

$$\begin{pmatrix} -1 \\ -1 \end{pmatrix}, \begin{pmatrix} -1 \\ 2 \end{pmatrix}, \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \end{pmatrix}.$$

We form their rotations as in (1.12), and we order them by increasing slope:

$$\begin{pmatrix} 2\\1 \end{pmatrix}, \begin{pmatrix} 1\\1 \end{pmatrix}, \begin{pmatrix} -1\\1 \end{pmatrix}, \begin{pmatrix} -2\\-1 \end{pmatrix}, \begin{pmatrix} 0\\-2 \end{pmatrix}$$

If we concatenate these vectors starting at the origin, then the resulting edges all remain in the non-negative orthant. The result is the convex pentagon P^+ in Corollary 1.5.3. As predicted, it coincides with the pentagon in (1.9).

The technique of tropical tmplicitization can be used, in principle, to compute the tropicalization of any parametrically presented algebraic variety. The details are more complicated than the simple curve case discussed here, and a proper treatment will require some toric geometry and concepts from resolution of singularities. We shall return to this subject in Chapter 7.

1.6 Group Theory

One of the origins in tropical geometry is the work of Bieri, Groves, Strebel and Neumann in group theory [BG84, BS80, BNS87]. Starting in the late 1970's, these authors associate polyhedral fans to certain classes of discrete
groups, and they establish remarkable results concerning generators, relations and higher cohomology of these groups in terms of these fans. This section aims to offer a first glimpse of this point of entry to the tropical universe.

We begin with an easy illustrative example. Fix a non-zero real number ξ and let G_{ξ} denote the group generated by the two invertible 2×2-matrices

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad X = \begin{pmatrix} 1 & 0 \\ 0 & \xi \end{pmatrix}. \tag{1.13}$$

What are the relations satisfied by these two generators? In particular, is the group G_{ξ} is finitely presented? Does this property depend on the number ξ ?

To answer this question, we explore some basic computations such as

$$X^{u}A^{c}X^{-u}X^{v}A^{d}X^{-v} = \begin{pmatrix} 1 & c\xi^{u} + d\xi^{v} \\ 0 & 1 \end{pmatrix}.$$

Here u, v, c and d can be arbitrary integers. This identity shows that the two matrices $X^u A^c X^{-u}$ and $X^v A^d X^{-v}$ commute, and this commutation relation is a valid relation among the two generators of G_{ξ} . If the number ξ is not algebraic over \mathbb{Q} then the set of all such commutation relations constitutes a complete presentation of G_{ξ} , and in this case the group G_{ξ} is never finitely presented. On the other hand, if ξ is an algebraic number then additional relations can be derived from the irreducible minimal polynomial $f \in \mathbb{Z}[x]$ of ξ . To show how this works, we consider the explicit example $\xi = \sqrt{2} + \sqrt{3}$. The minimal polynomial of this algebraic number is $f(x) = x^4 - 10x^2 + 1$. This polynomial translates into a word in the group generators as follows:

$$(X^{-4}A^{1}X^{4}) \cdot (X^{-2}A^{-10}X^{2}) \cdot (X^{0}A^{1}X^{0}) = X^{-4}AX^{2}A^{-10}X^{2}A.$$
(1.14)

We see that this word is a valid relation in G_{ξ} . Our question is whether the group of all relations is finitely generated. It turns out that the answer is affirmative for $\xi = \sqrt{2} + \sqrt{3}$, and we shall list the generators in Example 1.6.10 below. In general, finite presentation is characterized by the following result:

Proposition 1.6.1. The group $G_{\xi} = \langle A, X \rangle$ is finitely presented if and only if either the real number ξ or its reciprocal $1/\xi$ is an algebraic integer over \mathbb{Q} .

The condition that either ξ or $1/\xi$ is an algebraic integer says that either the highest term or the lowest term of f(x) has the coefficient +1 or -1. This is equivalent to saying that either the highest or the lowest term of the minimal polynomial f(x) is a unit in $\mathbb{Z}[x, x^{-1}]$. It is precisely this condition on leading terms that underlies the tropical thread in geometric group theory.

Bieri and Strebel introduced tropical varieties over the integers in their 1980 paper in metabelian groups [BS80]. Later work with Neumann [BNS87] extended their construction to a wider class of discrete groups. In what follows we restrict ourselves to metabelian groups whose corresponding module is cyclic. This special case suffices in order to explain the general idea, and to shed sufficient light on the mystery of why Proposition 1.6.1 might be true.

We begin with some commutative algebra definitions. Consider the Laurent polynomial ring $S = \mathbb{Z}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ over the integers \mathbb{Z} . The units in S are the monomials $\pm x^a = \pm x_1^{a_1} \cdots x_n^{a_n}$ where $a = (a_1, \ldots, a_n)$ runs over \mathbb{Z}^n . Let I be a proper ideal in S. If $w \in \mathbb{R}^n$ then the initial ideal $\operatorname{in}_w(I)$ is the ideal generated by all initial forms $\operatorname{in}_w(f)$ where f runs over I. The computation of the initial ideal $\operatorname{in}_w(I)$ from a generating set of I requires Gröbner bases over the integers. The relevant algorithm for computing $\operatorname{in}_w(I)$ from I is implemented in computer algebra systems such as Macaulay2 and Magma.

The tropical variety of the ideal I is defined as the following subset of \mathbb{R}^n :

$$\operatorname{trop}_{\mathbb{Z}}(I) = \{ w \in \mathbb{R}^n : in_w(I) \neq S \}.$$

Our construction of Gröbner fans over the field \mathbb{Q} in Chapter 2 will reveal that $\operatorname{trop}_{\mathbb{Z}}(I)$ is a polyhedral fan in \mathbb{R}^n . Moreover, it implies that the integral tropical variety contains the tropical variety over the field \mathbb{Q} as a subfan:

$$\operatorname{trop}_{\mathbb{Z}}(I) \supseteq \operatorname{trop}_{\mathbb{O}}(I),$$

but this containment is strict in general. For example, if n = 2 and I is the principal ideal $\langle x_1 + x_2 + 3 \rangle$, then $\operatorname{trop}_{\mathbb{Q}}(I)$ is the *tropical line*, which has three rays, but $\operatorname{trop}_{\mathbb{Z}}(I)$ additionally contains the positive quadrant.

We write R = S/I for the quotient Z-algebra, and, by mild abuse of notation, we write R^* for the multiplicative group generated by the images of the monomials. It follows from the results to be proved later in Chapter 3 that the complex variety of the ideal I is finite if and only if $\operatorname{trop}_{\mathbb{Q}}(I) = \{0\}$. Here we state the analogous result for tropical varieties over the integers.

Theorem 1.6.2 (Bieri-Strebel). The \mathbb{Z} -algebra R = S/I is finitely generated as a \mathbb{Z} -module if and only if

$$\operatorname{trop}_{\mathbb{Z}}(I) = \{0\}. \tag{1.15}$$

Proof. See [BS80, Theorem 2.4].

This raises the question of how to test this criterion in practise, and if (1.15) holds, how to determine a finite set of monomials $\mathcal{U} \subset R^*$ that generate R as an abelian group. It turns out that this can be done in Macaulay2.

Example 1.6.3. Fix integers m and n, where |m| > 1 and consider the ideal

$$J = \langle ms^{-1}t^{-1} + s^{-1} + t^{-1} + n + st, \, mst + s + t + n + s^{-1}t^{-1} \rangle \subset \mathbb{Z}[s^{\pm 1}, t^{\pm 1}].$$

This ideal is a variation on Example 43 in Strebel's exposition [Str84]. The condition (1.15) is satisfied. To find a generating set \mathcal{U} , we can run the following four lines of Macaulay2 code, for various fixed values of m and n:

```
R = ZZ[s,t,S,T]; m = 7; n = 13;
J = ideal(m*S*T+S+T+n+s*t,m*s*t+s+t+n+S*T,s*S-1,t*T-1);
toString leadTerm J
toString basis(R/J)
```

The output of this script is the same for all m and n, namely,

$$\mathcal{U} = \{1, s, st^{-1}, t, s^{-1}, s^{-1}t^{-1}, t^{-1}, t^{-2}\}.$$
(1.16)

For a proof that $\mathbb{Z}\mathcal{U} = \mathbb{R}/J$, it suffices to show that the initial monomial ideal of J with respect to the reverse lexicographic term order is generated by

{(m²-1)*S*T,t*T,m*s*T,S²,t*S,s*S,t²,s*t,s²,T³,S*T²,s*T²}

This proof amounts to computing a Gröbner basis over the integers ZZ. \Box

The integral tropical variety $\operatorname{trop}_{\mathbb{Z}}(I)$ is of interest even in the case n = 1.

Example 1.6.4. Suppose that ξ is an algebraic number of \mathbb{Q} and let I be the prime ideal of all Laurent polynomials f(x) in $\mathbb{Z}[x, x^{-1}]$ such that $f(\xi) = 0$. There are four possible cases of what the integral tropical variety can be:

- If ξ and $1/\xi$ are both algebraic integers then $\operatorname{trop}_{\mathbb{Z}}(I) = \{0\}$.
- If ξ is an algebraic integer but $1/\xi$ is not then $\operatorname{trop}_{\mathbb{Z}}(I) = \mathbb{R}_{\geq 0}$.
- If $1/\xi$ is an algebraic integer but ξ is not then $\operatorname{trop}_{\mathbb{Z}}(I) = \mathbb{R}_{\leq 0}$.
- If neither ξ nor $1/\xi$ are algebraic integers then $\operatorname{trop}_{\mathbb{Z}}(I) = \mathbb{R}$.

Examples of numbers for the first, third and last cases are $\xi = \sqrt{2} + \sqrt{3}$, $\xi = \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{3}}$, and $\xi = \sqrt{2} + \frac{1}{\sqrt{3}}$, respectively. In particular, we conclude from Proposition 1.6.1 that G_{ξ} is finitely presented if and only if $\operatorname{trop}_{\mathbb{Z}}(I) \neq \mathbb{R}$. \Box

We now come to the punchline of this section, namely, the extension of Example 1.6.4 to $n \ge 2$ variables. Let I be any ideal in $S = \mathbb{Z}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and set R = S/I. We associate with I a metabelian group G_I with a distinguished system of n + 1 of generators, namely, the group of 2×2 -matrices

$$G_I = \begin{pmatrix} 1 & R \\ 0 & R^* \end{pmatrix}$$

The elements of the group G_I are the matrices $\begin{pmatrix} 1 & f \\ 0 & m \end{pmatrix}$ where f is a Laurent polynomial and m is a Laurent monomial, but both considered modulo I.

The connection between tropical varieties and group theory is as follows.

Theorem 1.6.5 (Bieri-Strebel). The metabelian group G_I is finitely presented if and only if the integer tropical variety $\operatorname{trop}_{\mathbb{Z}}(I)$ contains no line.

This was the main result in the remarkable 1980 paper by Bieri and Strebel [BS80, Theorem A]. It predates the 1984 paper by Bieri and Groves [BG84] which is widely cited among tropical geometers for its resolution of problems left open in Bergman's 1971 paper on the logarithmic limit set.

In what follows we aim to shed some light on the presentation of the metabelian group G_I . We begin with the observation that G_I is always finitely generated, namely, by a natural set of n+1 matrices in R = S/I:

Lemma 1.6.6. The metabelian group G_I is generated by the n+1 matrices

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad and \quad X_i = \begin{pmatrix} 1 & 0 \\ 0 & x_i \end{pmatrix} \quad for \ i = 1, 2, \dots, n.$$

If n = 1 then we recover the group with two generators A and X seen at the beginning of this section. Indeed, if $I = \langle f(x) \rangle$ is the principal ideal generated by the minimal polynomial of an algebraic number ξ then $G_I = G_{\xi}$. In that special case, Theorem 1.6.5 is equivalent to Proposition 1.6.1.

Returning to the general case $n \ge 2$, we now examine the relations among the n + 1 generators in Lemma 1.6.6. Let us first assume that $I = \langle 0 \rangle$ is the zero ideal, so that R = S. Clearly, the matrices X_i and X_j commute, i.e., the commutator $[X_i, X_j] = X_i X_j X_i^{-1} X_j^{-1}$ is the 2×2-identity matrix. Next we consider the action of the group R^* on G_I by conjugation. For any monomial $m = x^u$ we have $X^u = \begin{pmatrix} 1 & 0 \\ 0 & x^u \end{pmatrix}$, and we find that the product $A^m = X^{-u}AX^u$ is equal to $\begin{pmatrix} 1 & m \\ 0 & 1 \end{pmatrix}$. Likewise, we have $A^{-m} = X^{-u}A^{-1}X^u = \begin{pmatrix} 1 & -m \\ 0 & 1 \end{pmatrix}$, so the same identify holds for monomials whose coefficient is -1. In particular, for any monomial m in S, the two matrices A and A^m commute. Hence, in the group $G_{(0)}$ we have

$$[X_i, X_j] = [A, A^m] = 1 \text{ for } 1 \le i < j \le n \text{ and monomials } m \in S^*.$$
(1.17)

Lemma 1.6.7. The relations (1.17) define a presentation of the group $G_{(0)}$.

For example, the following matrix lies in $G_{(0)}$ for any $f \in S$:

$$A^f = \begin{pmatrix} 1 & f \\ 0 & 1 \end{pmatrix}.$$

Indeed, if we write f as an alternating sum of monomials, say $\sum_{i=1}^{s} m_i x^{u_i}$, where $m_i \in \mathbb{Z}$ and $u_i \in \mathbb{Z}^n$, then this translates into the factorization

$$A^{f} = X^{-u_{1}}A^{m_{1}}X^{u_{1}}X^{-u_{2}}A^{m_{2}}X^{u_{2}}\cdots X^{-u_{s}}A^{m_{s}}X^{u_{s}}.$$

By applying the relations in (1.17), the word above can be transformed into the word for A^f that corresponds to any other way of writing f as an alternating sum of monomials. Hence the following statement makes sense:

Proposition 1.6.8. For any ideal I in S, the group G_I has the presentation

$$[X_i, X_j] = [A, A^m] = A^f = 1, (1.18)$$

where m runs over monomials, f runs over I, and $1 \le i < j \le n$.

As it stands, this presentation is infinite, and we are interested in the question whether the set of relations (1.18) can be replaced by a finite subset. We would like to know whether the group G_I is finitely presented. To answer this, we first note that the conjugation action satisfies the following relations:

$$A^f A^g = A^g A^f = A^{f+g}$$
 and $(A^f)^g = (A^g)^f = A^{fg}$ for $f, g \in S$.

This shows that it suffices to take f from any finite generating set of the ideal I. So, the question is whether there exists a finite subset $\mathcal{U} \subset \mathbb{Z}^n$ such that the monomials $m = \pm x^u$ with $u \in \mathcal{U}$ suffice in the presentation (1.18).

Theorem 1.6.5 offers a criterion for testing whether such a finite set \mathcal{U} exists. For instances where the answer is affirmative, we can use the techniques in [BS80, §3] to construct an explicit generating set \mathcal{U} . The underlying techniques are quite delicate and have not yet been developed into a practical algorithm. In what follows we sketch ideas on how one might approach this.

The first step is compute the integral tropical variety $\operatorname{trop}_{\mathbb{Z}}(I)$ from a given generating set of I. This can be done by first homogenizing the ideal I and then computing the Gröbner fan of the homogeneous ideal. The Gröbner fan of a homogeneous ideal J in $\mathbb{Z}[x_0, x_1, \ldots, x_n]$ is a polyhedral fan in \mathbb{R}^{n+1} such that the initial ideal $\operatorname{in}_w(I)$ is constant as w ranges over the relative interior of any cone in the fan. For coefficients in a field K, this will be explained in full detail in Chapter 2. The general theory is analogous over the integers \mathbb{Z} , except that Gröbner fans over \mathbb{Z} tend to be finer than over K. For example, if $I = \langle 2x_1, x_1x_2 - x_1x_3 \rangle$ then the Gröbner fan over \mathbb{Q} consists of single cone, while the Gröbner fan over \mathbb{Z} has a wall on the plane $\{w_2 = w_3\}$.

In the course of computing the Gröbner fan of I, one obtains a generating set for every initial ideal $in_w(I)$. This can be further extended to a finite generating set \mathcal{B} of I with the property that, for every $w \in \mathbb{R}^n$, either $in_w(I)$ is a proper ideal in S or the finite set $\{in_w(f) : f \in \mathcal{B}\}$ contains a unit. A subset \mathcal{B} of the ideal I that enjoys this property is called a *tropical basis*. Every Laurent polynomial in a tropical basis \mathcal{B} can be scaled by a unit, so we can always assume that the relevant leading monomial is the constant 1.

Suppose now that I is an ideal in S which satisfies the condition of Theorem 1.6.5, and that we have computed a tropical basis \mathcal{B} for I. Then

For all $w \in \mathbb{R}^n$ there exists $f \in \mathcal{B}$ with $\operatorname{in}_w(f) = 1$ or $\operatorname{in}_{-w}(f) = 1$. (1.19)

For each Laurent polynomial f in the tropical basis \mathcal{B} let support(f) denote the set of all vectors $a \in \mathbb{Z}^n$ such that the monomial x^a appears with non-zero coefficient in f. We define the Newton polytope of the tropical basis \mathcal{B} to be the convex hull of the union of these support sets for all f in \mathcal{B} :

Newton(
$$\mathcal{B}$$
) := conv $\left(\bigcup_{f \in \mathcal{B}} \operatorname{support}(f)\right)$

By examining the proof technique used in [BS80, §3.5], one can derive the following explicit version of the "if"-direction in the Bieri-Strebel Theorem:

Theorem 1.6.9. Fix a tropical basis \mathcal{B} satisfying (1.19) for the ideal I. Then the metabelian group G_I is presented by the relations (1.17) where fruns over the elements in the tropical basis \mathcal{B} and $m = x^u$ runs over the set Newton $(\mathcal{B}) \cap \mathbb{Z}^n$ of lattice points u in the Newton polytope the tropical basis.

We close this section with two simple examples:

Example 1.6.10. Let n = 1 and let I be the prime ideal of $\xi = \sqrt{2} + \sqrt{3}$. The singleton $\mathcal{B} = \{x^4 - 10x^2 + 1\}$ is a tropical basis of I satisfying (1.19). Then the group $G_{\xi} = G_I$ is presented by five relations. The first relation is the word in (1.14) and the other four required relations are the words

$$[A, A^{x}] = Ax^{-i}Ax^{i}A^{-1}x^{-i}A^{-1}x^{i}$$
 for $i = 1, 2, 3, 4$.

Example 1.6.11. We consider the group in [Str84, Example 43]. Let $S = \mathbb{Z}[s^{\pm 1}, t^{\pm 1}]$ with $I = \langle f \rangle$ generated by the polynomial in Example 1.6.3:

$$f(s,t) = ms^{-1}t^{-1} + s^{-1} + t^{-1} + n + st.$$

The tropical variety $\operatorname{trop}_{\mathbb{Z}}(I)$ contains no line. A minimal tropical basis satisfying the condition (1.19) consists of three Laurent polynomials:

$$\mathcal{B} = \{ s^{-1}t^{-1}f(s,t), sf(s,t), tf(s,t) \}$$

The corresponding polytope Newton(\mathcal{B}) is a planar convex 7-gon that has 14 lattice points, corresponding to the 14 Laurent monomials:

$$m = s^2 t, st^2, st, s, s/t, t, 1, 1/t, t/s, 1/s, 1/st, 1/st^2, 1/s^2 t, 1/s^2 t^2.$$

The metabelian group G_I has three generators A, X_1, X_2 . From the description in Theorem 1.6.9 we get a presentation with 17 = 3 + 14 relations.

In our view, it would be worthwhile to further develop the connection between Gröbner bases and tropical geometry over \mathbb{Z} , and to revisit the beautiful group theory results by Bieri, Groves, Neumann and Strebel from a computational point of view. Surely, there will be plenty of applications.

1.7 Curve Counting

One of the early successes that brought tropical methods to the attention of geometers was the work of Mikhalkin [Mik05] on Gromov-Witten invariants of

the plane. These invariants count the number of complex algebraic curves of a given degree and genus passing through a given number of points. Mikhalkin proved that complex curves can be replaced by tropical curves, and he then derived a combinatorial formula for the count in the tropical case. The objective of this section is to present the basic ideas and the main result.

As a warm-up, let us consider the question of how many singular quadratic curves pass through four general points in the plane. The answer to this question is three. A singular quadric is the union of two lines, and, since the four points are in general position, there are precisely three pairs of lines that pass through them. This analysis is valid both in classical geometry and in tropical geometry, and it yields the same result, namely 3, in both cases.

To state our general problem, we now review some classical facts about curves in the complex projective plane \mathbb{P}^2 . If C is a smooth curve of degree d in \mathbb{P}^2 then its genus is the number of handles of C when regarded as a two-dimensional Riemann surface over the real numbers. That genus equals

$$g(C) = \frac{1}{2}(d-1)(d-2).$$

Moreover, that same number counts the lattice points in the interior of the Newton polygon of the general curve of degree d. That Newton polygon is the triangle with vertices (0, 0, d), (0, d, 0) and (d, 0, 0). In symbols,

$$g(C) = #(int(Newt(C)) \cap \mathbb{Z}^3).$$

The set of all curves of degree d forms a projective space of dimension

$$\binom{d+2}{2} - 1 = \frac{1}{2}(d-1)(d-2) + 3d - 1.$$
(1.20)

As the $\binom{d+2}{2}$ coefficients of its defining polynomial vary, the curve C may acquire one or more singular points. The simplest type of singularity is a *node*. Each time the curve acquires a node, the genus drops by one. Thus for a singular curve C_{sing} with ν nodes and no other singularities, the genus is

$$g(C_{\rm sing}) = \frac{1}{2}(d-1)(d-2) - \nu.$$
 (1.21)

We are interested in the following problem of enumerative geometry: What is the number $N_{g,d}$ of irreducible curves of genus g and degree d that pass through g + 3d - 1 general points in the complex projective plane \mathbb{P}^2 ?

1.7. CURVE COUNTING

This question makes sense because the variety of curves of degree d and genus g is expected to have dimension g + 3d - 1, by (1.20) and (1.21), and each of the points poses one independent condition on the curve. Thus we expect the number $N_{g,d}$ of curves satisfying all constraints to be finite. Gromov-Witten theory offers the tools for proving that this is indeed the case.

There is also a closely related counting problem where all curves are allowed, not just irreducible ones. If we denote that number by $N_{g,d}^{\text{red}}$ then the reducible quadrics at the beginning of the section would give the count:

$$N_{-1,2}^{\text{red}} = 3.$$

In what follows we restrict our attention to the case of irreducible curves. The numbers $N_{g,d}$ are called *Gromov-Witten invariants* of the plane \mathbb{P}^2 . Their study has been a topic of considerable interest among geometers, and it has been a boon for the tropical approach. Here are some explicit such numbers:

Example 1.7.1. The simplest Gromov-Witten invariants of \mathbb{P}^2 are $N_{0,1} = 1$ and $N_{0,2} = 1$. This translates into saying that a unique line passes through two points, and that a unique conic passes through five points. We also have $N_{1,3} = 1$, which says that there is a unique cubic through nine points. \Box

Example 1.7.2. The first non-trivial number is $N_{0,3} = 12$, and we wish to explain this in some detail. It concerns curves defined by cubic polynomials

$$f = c_0 x^3 + c_1 x^2 y + c_2 x^2 z + c_3 x y^2 + c_4 x y z + c_5 x z^2 + c_6 y^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_7 y^2 z + c_8 y z^2 + c_9 z^3 + c_9$$

For general coefficients c_0, \ldots, c_9 , the curve $\{f = 0\}$ is smooth of genus g = 1. The curve becomes rational, i.e. the genus to drops to g = 0, precisely when there is a singular point, and this happens if and only if the *discriminant* of fvanishes. The discriminant $\Delta(f)$ is a homogeneous polynomial of degree 12 in the 10 unknown coefficients c_0, c_1, \ldots, c_9 . It is a sum of 2040 monomials:

$$\Delta(f) = 19683c_0^4c_6^4c_9^4 - 26244c_0^4c_6^3c_7c_8c_9^3 + 5832c_0^4c_6^3c_8^3c_9^2 + \dots - c_2^2c_3c_4^4c_5^3c_6^2 \quad (1.22)$$

The study of discriminants and resultants is the topic of the book by Gel'fand, Kapranov and Zelevinsky [GKZ08], which contains many formulas for computing them. Here is a simple determinant formula for (1.22). The Hessian H of the quadrics $\frac{\partial f}{\partial x}$, $\frac{\partial f}{\partial y}$ and $\frac{\partial f}{\partial z}$ is a polynomial of degree 3. Form the 6×6 matrix M(f) whose entries are the coefficients of the six quadrics $\frac{\partial f}{\partial x}$, $\frac{\partial f}{\partial y}$, $\frac{\partial f}{\partial z}$, $\frac{\partial H}{\partial x}$, $\frac{\partial H}{\partial y}$, $\frac{\partial H}{\partial z}$. Then $\Delta(f) = \det(M(f))$. Now, suppose we are given eight points in \mathbb{P}^2 and we require them to lie on the cubic $\{f = 0\}$. This translates into eight linear equations in c_0, c_1, \ldots, c_9 . Combining the eight linear equations with the degree 12 equation $\Delta(f) = 0$, we obtain a system of equations that has 12 solutions in \mathbb{P}^9 . These solutions are the coefficient vectors of the $N_{0,3} = 12$ rational cubics we seek to find. \Box

Example 1.7.3. Quartic curves in the plane \mathbb{P}^2 can have genus 3, 2, 1 or 0. The Gromov-Witten numbers corresponding to these four cases are

$$N_{3,4} = 1$$
, $N_{2,4} = 27$, $N_{1,4} = 225$, and $N_{0,4} = 620$.

Here 27 is the degree of the discriminant of a ternary quartic. The last entry means that there are 620 rational quartics through 11 general points. \Box

The result of Mikhalkin [Mik05] can be stated informally as follows:

Theorem 1.7.4. The Gromov-Witten numbers $N_{q,d}$ can be found tropically.

The following discussion is aimed at stating precisely what this means. We consider tropical curves of degree d in \mathbb{R}^2 . Each such curve C is the planar dual graph to a regular subdivision of the triangle with vertices (0,0), (0,d)and (d,0). We say that the curve C is *smooth* if this subdivision consists of d^2 triangles each having unit area 1/2. Equivalently, the tropical curve C is *smooth* if it has d^2 vertices. These vertices are necessarily trivalent.

A tropical curve C is called *simple* if each vertex is either trivalent or is locally the intersection of two line segments. Equivalently, C is simple if the corresponding subdivision consists only of triangles and parallelograms. Here the triangles are allowed to have large area. Let t(C) be the number of trivalent vertices and let r(C) be the number of bounded edges of C.

We define the *genus* of a simple tropical curve C by the formula

$$g(C) = \frac{1}{2}t(C) - \frac{1}{2}r(C) + 1.$$
(1.23)

It is instructive to check that this definition makes sense for smooth tropical curves. Indeed, if C is smooth then $t(C) = d^2$ and r(C) = 3d, and we recover the formula for the genus of a classical complex curve that is smooth:

$$g(C) = \frac{1}{2}d^2 - \frac{1}{2}3d + 1 = \frac{1}{2}(d-1)(d-2)$$

We finally define the *multiplicity* of a simple curve C as the product of the normalized areas of all triangles in the corresponding subdivision. Thus, in

computing the multiplicity of C, we disregard the "nodal singularities", which correspond to 4-valent crossings. We just multiply positive integers attached to the trivalent vertices. The contribution of a trivalent vertex can also be computed by the formula $w_1w_2|\det(u_1, u_2)|$ where w_1, w_2, w_3 are the weights of the adjacent edges and u_1, u_2, u_3 are their primitive edge directions. That formula is independent of the choice made because of the balancing condition $w_1u_1 + w_2u_2 + w_3u_3 = 0$. If the curve is smooth then its multiplicity equals 1.

Here now is the precise statement of what was meant in Theorem 1.7.4:

Theorem 1.7.5 (Mikhalkin's Correspondence Principle). The number of simple tropical curves of degree d and genus g that pass through g + 3d - 1 general points in \mathbb{R}^2 , where each curve is counted with its multiplicity, is equal to the Gromov-Witten number $N_{g,d}$ of the complex projective plane \mathbb{P}^2 .

Example 1.7.6. We show the tropical count for the number $N_{0,3} = 12$ in Example 1.7.2.

The proof of Theorem 1.7.5 given by Mikhalkin in [Mik05] uses methods from complex geometry, specifically, deformations of *J*-holomorphic curves. Subsequently, Gathmann and Markwig [GM07a, GM07b] developed a more algebraic approach, and this has led to an ongoing systematic development of tropical moduli spaces and tropical intersection theory on such spaces.

We close with one more example of what can be done with tropical curves in enumerative geometry. The Gromov-Witten invariants $N_{0,d}$ for rational curves (genus g = 0) satisfy the following remarkable recursive relations:

$$N_{0,d} = \sum_{\substack{d_1+d_2=d\\d_1,d_2>0}} \left(d_1^2 d_2^2 \binom{3d-4}{3d_1-2} - d_1^3 d_2 \binom{3d-4}{3d_1-1} \right) N_{0,d_1} N_{0,d_2}.$$
(1.24)

This equation is due to Kontsevich, who derived them from the WDVV equations, named after the theoretical physicists Witten, Dijkgraaf, Verlinde and Verlinde, which express the associativity of quantum cohomology of \mathbb{P}^2 .

Using Mikhalkin's Correspondence Principle, Gathmann and Markwig [GM08] gave a proof of this formula using tropical methods. Namely, they establish the combinatorial result that simple tropical curves of degree d and genus 0 passing through 3d-1 points satisfy the Kontsevich relations (1.24).

1.8 Compactifications

Many of the advanced tools of algebraic geometry, such as intersection theory, are custom-taylored for varieties that are compact, such as complex projective varieties. On the other hand, in concrete problems, the given spaces are quite often not compact. In such a case one first needs to replace the given variety X by a nice compact variety \bar{X} that contains X as dense subset. Here the emphasis lies on the adjective "nice" because the advanced tools will not work or will give incorrect answers if the boundary $\bar{X} \setminus X$ is not good enough.

We begin by considering a non-singular curve X in the *n*-dimensional complex torus $(\mathbb{C}^*)^n$. The curve X is not compact, and we wish to add a finite set of points to X so as to get a smooth compactification \overline{X} of X.

From a geometric point of view, it is clear what must be done. When viewed over the field of real numbers \mathbb{R} , the curve X is a surface. More precisely, X is a non-compact Riemann surface. It is an orientable smooth compact surface of some genus g with a certain number m of points removed. The problem is to identify the m missing points and to fill them back in. What is the algebraic procedure that accomplishes this geometric process?

To illustrate the algebraic complications, we begin with a plane curve

$$X = \{ (x, y) \in (\mathbb{C}^*)^2 : f(x, y) = 0 \}$$

Our smoothness hypothesis says that the three Laurent polynomial equations

$$f(x,y) = \frac{\partial f}{\partial x}(x,y) = \frac{\partial f}{\partial y}(x,y) = 0$$
(1.25)

have no common solutions (x, y) in the algebraic torus $(\mathbb{C}^*)^2$. The first thing one might try to compactify X is to homogenize the given Laurent polynomial

$$f^{hom}(x,y,z) = z^N \cdot f\left(\frac{x}{z},\frac{y}{z}\right).$$

Here N is the smallest integer such that this expression is a polynomial. This homogeneous polynomial defines a curve in the complex projective plane \mathbb{P}^2 :

$$X^{hom} = \{ (x:y:z) \in \mathbb{P}^2 : f^{hom}(x,y,z) = 0 \}.$$

This curve is a compactification of X but it usually not what we had in mind.

Example 1.8.1. Let X be the curve in $(\mathbb{C}^*)^2$ defined by the polynomial

$$f(x,y) = c_1 + c_2 xy + c_3 x^2 y + c_4 x^3 y + c_5 x^3 y^2.$$
(1.26)

Here c_1, c_2, c_3, c_4, c_5 are any complex numbers that satisfy

$$c_2c_3^4 - 8c_2^2c_3^2c_4 + 16c_2^3c_4^2 - c_1c_3^3c_5 + 36c_1c_2c_3c_4c_5 - 27c_1^2c_4c_5^2 \neq 0.$$
(1.27)

This inequation ensures that the given non-compact curve X is smooth. The discriminant polynomial in (1.27) is computed by eliminating x and y from (1.25). The homogenization of the polynomial f(x, y) equals

$$f^{hom}(x,y,z) = c_1 z^5 + c_2 x y z^3 + c_3 x^2 y z^2 + c_4 x^3 y z + c_5 x^3 y^2.$$

The corresponding projective curve X^{hom} in \mathbb{P}^2 is compact but it is not smooth. The boundary we have added to compactify consists of two points

$$X^{hom} \setminus X = \{ (1:0:0), (0:1:0) \}.$$

Both of these points are highly singular on the compact curve X^{hom} . Their respective multiplicities are 7 and 14.

Another thing one might try is the closure of our curve $X \subset (\mathbb{C}^*)^2$ in the product of two projective lines $\mathbb{P}^1 \times \mathbb{P}^1$. Then the ambient coordinates are $((x_0 : x_1), (y_0 : y_1))$, and our polynomial is replaced by its *bihomogenization*

$$x_0^3 y_0^2 f(\frac{x_1}{x_0}, \frac{y_1}{y_0}) = c_1 x_0^3 y_0^2 + c_2 x_1 y_1 x_0^2 y_0 + c_3 x_1^2 y_1 x_0 y_0 + c_4 x_1^3 y_1 y_0 + c_5 x_1^3 y_1^2.$$

The compactification X^{bihom} of X is the zero set of this polynomial in $\mathbb{P}^1 \times \mathbb{P}^1$. Now, the boundary we have added to compactify consists of three points

$$X^{bihom} \setminus X = \left\{ \left((1:0), (0:1) \right), \left((0:1), (1:0) \right), \left((0:1), (c_5, -c_4) \right) \right\}.$$
(1.28)

The compactification X^{bihom} is still singular but it is better than X^{hom} . The first two points in (1.28) are singular, of multiplicity 5 and 3 respectively, but the third point is smooth. It correctly fills in one of the holes in X. \Box

The solution to our problem offered by tropical geometry is to replace a given non-compact variety $X \subset (\mathbb{C}^*)^n$ by a *tropical compactification* X^{trop} . Each such tropical compactification of X is characterized by a polyhedral fan in \mathbb{R}^n whose support is the tropical variety corresponding to X. In small and low-dimensional examples, including all curves and all hypersurfaces, there is a unique coarsest fan structure, and in these cases we obtain a canonical tropical compactification. However, in general, picking a tropical compactification requires making choices, and X^{trop} will depend on these choices.

Tropical compactifications were introduced by Jenia Tevelev in [Tev07]. The geometric foundation for his construction is the theory of *toric varieties*. In this introductory section we do not assume any familiarity with toric geometry, but we do encourage the reader to start perusing one of the text books on this topic. In Chapter 6, we shall give a brief introduction to toric geometry, and we shall explain its relationship to tropical geometry. In that later chapter, we shall see the precise definition of the tropical compactification X^{trop} of a variety $X \subset (\mathbb{C}^*)^n$, and we shall prove its key geometric properties. In what follows, we keep the discussion informal and entirely elementary, and we simply go over a few examples of tropical compactifications.

Example 1.8.2. Let X be the plane complex curve in (1.26). Its tropical compactification X^{trop} is a smooth elliptic curve, that is, it is a Riemann surface of genus g = 1. The boundary $X^{trop} \setminus X$ consists of m = 4 points. Unlike the extra points in the bad compactifications X^{hom} and X^{bihom} in Example 1.8.1, these four new points are smooth on X^{trop} . Thus, with hindsight, we see that the complex curve X is a real torus with m = 4 points removed.

The tropical compactification of a plane curve is nothing but the classical compactification derived from its Newton polygon. Here, the Newton polygon is the quadrangle Newt $(f) = \operatorname{conv}\{(0,0), (1,1), (3,2), (3,1)\}$. The genus g of the curve X is the number of interior lattice points of Newt(f).

The tropical curve is the union of the inner normal rays to the four edges of this quadrangle. In other words, $\operatorname{Trop}(X)$ consists of the four rays spanned by (1, -1), (1, -2), (-1, 0) and (-1, 3). Each ray has multiplicity one because the edges of Newt(f) have no interior lattice points. This shows that m =4 points need to be added to X to get X^{trop} . The directions of the rays specifies how these new points should be glued into X in order to make them smooth in X^{trop} . Algebraically, this process can be described by replacing the given polynomial f by a certain homogeneous polynomial f^{trop} , but the homogenization process is now a little bit more tricky. One uses homogeneous coordinates, in the sense of David Cox, on the toric surface given by Newt(f). These generalize the homogeneous coordinates we used for \mathbb{P}^2 and $\mathbb{P}^1 \times \mathbb{P}^1$. \Box

The examples of plane curves has two natural generalizations, namely, curves in $(\mathbb{C}^*)^n$ and hypersurfaces in $(\mathbb{C}^*)^n$. We briefly discuss both of these.

1.8. COMPACTIFICATIONS

If X is a curve in $(\mathbb{C}^*)^n$ then the geometry is still easy. All we are doing is to fill in m missing points in a punctured Riemann surface of genus g. However, the algebra is more complicated now than in Example 1.8.2. The curve X is given by an ideal $I \subset \mathbb{C}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and our primary challenge is to determine the number m from I. The number m is the sum of the multiplicities of the rays in the tropicalization of X. The tropical curve $\operatorname{Trop}(X)$ is a finite union of rays in \mathbb{R}^n but it is generally impossible to determine these rays from (the Newton polytopes of) the given generators of I. To understand how $\operatorname{Trop}(X)$ arises from I, one needs the concepts pertaining to Gröbner bases and initial ideals to be introduced in Chapters 2 and 3. In practise, the software GFan, due to Anders Jensen, can be used to compute the tropical curve $\operatorname{Trop}(X)$ and the multiplicity of each of its rays.

If X is a hypersurface in $(\mathbb{C}^*)^n$ then the roles are reversed. The algebra is still easy but the geometry is more complicated now than in Example 1.8.2. Let $f = f(x_1, \ldots, x_n)$ be the polynomial that defines X. We compute its Newton polytope Newt $(f) \subset \mathbb{R}^n$. The tropical compactification X^{trop} has one boundary divisor for each facet of Newt(f). These boundary divisors are varieties of dimension n-2 that get glued to the (n-1)-dimensional variety X in order to create the compact (n-1)-dimensional variety X^{trop} . The precise nature of this gluing is determined by the ray normal to the facet. What is different from the curve case is that the boundary divisors are themselves non-trivial hypersurfaces, and they are no longer pairwise disjoint. In fact, describing their intersection pattern in $X^{trop} \setminus X$ is an essential part of the construction. The relevant combinatorics is encoded in the facial structure of the polytope Newt(f), and we record this data in the tropical hypersurface.

Tropical geometry provides the tools to generalize these constructions to an arbitrary *d*-dimensional subvariety *X* of the algebraic torus $(\mathbb{C}^*)^n$. The variety *X* is presented by an ideal *I* in $\mathbb{C}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Given any generating set of *I*, we can compute the tropical variety $\operatorname{Trop}(X)$. For small examples this can be done by hand, but for larger examples we use software such as **GFan** for that computation. The output is a polyhedral fan Δ in \mathbb{R}^n whose support $|\Delta|$ equals $\operatorname{Trop}(X)$. That fan determines a tropical compactification $X^{trop}(\Delta)$ of the variety *X*. Now, this compactification may not be quite nice enough, so one sometimes has to replace the fan Δ by a suitable refinement Δ' . This induces a map $X^{trop}(\Delta') \to X^{trop}(\Delta)$, and now $X^{trop}(\Delta')$ may satisfy the technical conditions for tropical compactifications required by Tevelev in [Tev07]. For example, Δ may not be a simplicial fan, and, as is customary in toric geometry, we could replace Δ' by a triangulation of Δ .

Let us consider the case when X is an irreducible surface in $(\mathbb{C}^*)^n$. In any compactification \bar{X} of X, the boundary $\bar{X} \setminus X$ is a finite union of irreducible curves. What is desired is that these curves are smooth and that they intersect each other transversally. If this holds then the boundary $\bar{X} \setminus X$ is said to have *normal crossings*. The tropical compactifications of a surface X usually have the normal crossing property. Here the tropical variety $\operatorname{Trop}(X)$ supports a two-dimensional fan in \mathbb{R}^n . Such a fan has a unique coarsest fan structure. We identify the tropical surface $\operatorname{Trop}(X)$ with that coarsest fan Δ , and we abbreviate $X^{trop} = X^{trop}(\Delta)$. The rays in the fan $\operatorname{Trop}(X)$ correspond to the irreducible curves in $\bar{X} \setminus X$, and two such curves intersect if and only if the corresponding rays span a two-dimensional cone. Since the fan $\operatorname{Trop}(X)$ is two-dimensional, it has no cones of dimension ≥ 3 . This means that the intersection of any three of the irreducible curves in $\bar{X} \setminus X$ is empty.

Example 1.8.3. Let I be the ideal minimally generated by three linear polynomials $a_1x_1+a_2x_2+a_3x_3+a_4x_4+a_5x_5+a_6$ in $\mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}, x_3^{\pm 1}, x_4^{\pm 1}, x_5^{\pm 1}]$. Its variety X is an non-compact surface in $(\mathbb{C}^*)^5$. If we took the variety of I in affine space \mathbb{C}^5 then this would simply be an affine plane \mathbb{C}^2 . But the torus $(\mathbb{C}^*)^5$ is obtained from \mathbb{C}^5 by removing the hyperplanes $\{x_i = 0\}$. Hence our non-compact surface X equals the affine plane \mathbb{C}^2 with five lines removed. Equivalently, X is the complex projective plane \mathbb{P}^2 with six lines removed.

If the three generators of I are linear polynomials with random coefficients, then the six lines form a normal crossing configuration in \mathbb{P}^2 , i.e., no three of the lines intersect. In that generic case, the tropical compactification is constructed by simply filling the six lines back in, that is, we have $X^{trop} = \mathbb{P}^2$. Here the tropical variety Trop(X) consists of six rays and the 15 two-dimensional cones spanned by any two of the rays. Five of the rays are spanned by the standard basis vectors e_1, e_2, e_3, e_4, e_5 of \mathbb{R}^5 , and the sixth rays is spanned by their negated sum $-e_1 - e_2 - e_3 - e_4 - e_5$.

The situation is more interesting if the generators of I are special, e.g.,

$$I = \langle x_1 + x_2 - 1, x_3 + x_4 - 1, x_1 + x_3 + x_5 - 1 \rangle.$$

For this particular ideal, the configuration of six lines in \mathbb{P}^2 has four triples of lines that meet in one point. Two of these special intersection points are

$${x_1 = x_4 = x_5 = 0, x_2 = x_3 = 1}$$
 and ${x_2 = x_3 = x_5 = 0, x_1 = x_4 = 1}.$

The other two points lie on the line at infinity, where they are determined by $\{x_1 = x_2 = 0\}$ and $\{x_3 = x_4 = 0\}$ respectively. The tropical compactification is constructed geometrically by blowing up these four special points. This process replaces each triple intersection point by a new line that meets the three old lines transversally at three distinct points. Thus X^{trop} is a compact surface whose boundary $X^{trop} \setminus X$ consists of ten lines, namely, the six old lines that had been removed from \mathbb{P}^2 plus the four new lines from blowing up. Now, no three lines intersect, so the boundary $X^{trop} \setminus X$ is normal crossing. There are 15 pairwise intersection points, three on each of the four new lines, and three old intersection points. The latter are determined by $\{x_1 = x_3 = 0\}$, $\{x_2 = x_4 = 0\}$ and by intersecting $\{x_5 = 0\}$ with the line at infinity.

The combinatorics of this situation is encoded in the tropical surface $\operatorname{trop}(X)$. It consists span 15 two-dimensional cones which are spanned by 10 rays. The rays correspond to the ten lines, and their primitive generators are

$$e_1, e_2, e_3, e_4, e_5, -e_1 - e_2 - e_3 - e_4 - e_5$$

 $e_1 + e_4 + e_5, e_2 + e_3 + e_5, -e_3 - e_4 - e_5, -e_1 - e_2 - e_5.$

The tropical surface $\operatorname{trop}(X)$ is the cone over a graph which is known as the *Petersen graph*. The ten vertices of the Petersen graph correspond to the ten lines in $X^{trop} \setminus X$, and the 15 edges of the Petersen graph correspond to the pairs of lines that intersect on the tropical compactification X^{trop} .

The previous example shows that tropical compactifications are nontrivial and interesting even for linear ideals I. Since linear ideals cut out linear spaces, we refer to the tropical variety $\operatorname{trop}(X)$ as a *tropical linear* space. The combinatorics of tropical linear spaces is governed by the theory of matroids. This will be explained in Chapter 4. In the linear case, the open variety $X \subset (\mathbb{C}^*)^n$ is the complement of an arrangement of n+1 hyperplanes in a projective space, and the tropical compactification X^{trop} was already known before the advent of tropical geometry. It is essentially equivalent to the wonderful compactifications of a hyperplane arrangement complement due to De Concini and Procesi. This was shown in [FS05, Theorem 6.1].

1.9 Exercises

1. Let $p : \mathbb{R}^n \to \mathbb{R}$ be a function that is continuous, concave and piecewiselinear, with finitely many linear functions having integer coefficients. Show that p can be represented by a tropical polynomial in x_1, \ldots, x_n .

- 2. Formulate the Fundamental Theorem of Algebra in the tropical setting, and give a proof. Why is the tropical semiring "algebraically closed"?
- 3. Prove Proposition 1.2.3. This concerns the tropical interpretation of the dynamic programming method for integer programming.
- 4. Show that for any collection of five points in the plane there is a unique tropical conic passing through them. If the five points are in special position then you will need to use stable intersections to get uniqueness.
- 5. Let $D = (d_{ij})$ be a symmetric $n \times n$ -matrix with zeros on the diagonal and positive off-diagonal entries. We say that D represents a metric space if the triangle inequalities $d_{ik} \leq d_{ij} + d_{jk}$ hold for all indices i, j, k. Show that D represents a metric space if and only if $D \odot D = D$.
- 6. The tropical 3×3-determinant is a piecewise-linear real-valued function ℝ^{3×3} → ℝ on the 9-dimensional vectors space of 3 × 3-matrices. Describe all the regions of linearity of this function and their boundaries. In tropical geometry, what does it mean for a matrix to be singular?
- 7. How many distinct combinatorial types of quadratic curves are there?
- 8. Prove that the stable self-intersection of a plane curve is precisely its set of vertices. What does this mean for classical algebraic geometry?
- 9. Given five general points in \mathbb{R}^2 , there exists a unique tropical quadric passing through these points. Compute and draw the quadratic curve passing through the points (0,5), (1,0), (4,2), (7,3) and (9,4).
- 10. A tropical cubic curve in \mathbb{R}^2 is *smooth* if it has precisely nine nodes. Prove that every smooth cubic curve has a unique bounded region, and that this region can have either three, four, five, six, seven, eight, or nine edges. Draw examples for all seven cases.
- 11. Install Anders Jensen's software **GFan** on your computer. Download the manual and try running one example.
- 12. The amoeba of a curve of degree four in the plane C² can have either 0,
 1, 2 or 3 bounded convex regions in its complement. Construct explicit examples for all four cases.

- 13. Prove Theorem 1.4.2 on the logarithmic limit set, at least for curves.
- 14. Consider the plane curve given by the parametrization

 $x = (t-1)^{13}t^{19}(t+1)^{29}$ and $y = (t-1)^{31}t^{23}(t+1)^{17}$.

Find the Newton polygon of the implicit equation f(x, y) = 0 of this curve. How many terms do you expect the polynomial f(x, y) to have?

- 15. Let v_1, v_2, \ldots, v_m be vectors in \mathbb{Z}^n that sum to zero: $v_1 + v_2 + \cdots + v_m = 0$. Show that there exists a algebraic curve in $(\mathbb{C}^*)^n$ whose corresponding tropical curve in \mathbb{R}^n consists of the rays spanned by v_1, v_2, \ldots, v_m .
- 16. Let $\xi = \frac{1}{4}(1 + \sqrt{33})$ and consider the group generated by the matrices A and X in (1.13). Can you construct a finite presention of this group?
- 17. Let I be any ideal generated by two linear forms in $\mathbb{Z}[x, y, z]$. What can the integral tropical variety $\operatorname{trop}_{\mathbb{Z}}(I)$ look like? List all possibilities.
- 18. Given 14 general points in the plane \mathbb{C}^2 , what is the number of rational curves of degree five that pass through these 14 points?
- 19. Consider the curve X in $(\mathbb{C}^*)^3$ cut out by two general polynomials of degree two. What is the genus g and the number m of punctures of this Riemann surface? Describe its tropical compactification X^{trop} .
- 20. The set of all singular 3×3 -matrices with non-zero complex entries is a hypersurface X in the 9-dimensional algebraic torus $(\mathbb{C}^*)^{3\times 3}$. Describe its tropical compactification X^{trop} . How many irreducible components does the boundary $X^{\text{trop}} \setminus X$ have? How do these components intersect?

Chapter 2

Building Blocks

Tropical geometry is a marriage between algebraic and polyhedral geometry. In order to develop this properly, we need some tools and building blocks from various mathematical disciplines, such as abstract algebra, discrete mathematics, elementary algebraic geometry, and symbolic computation. The first four sections of this chapter will introduce these building blocks. They are fields and valuations, algebraic varieties, polyhedral geometry, and Gröbner bases. In the last section we take a first step into the tropical world by defining a family of bases for ideals in Laurent polynomial ring.

2.1 Fields

Let K be a field. We denote by K^* the nonzero elements of K. A valuation on K is a function val: $K \to \mathbb{R} \cup \{\infty\}$ satisfying the following three axioms:

- 1. $\operatorname{val}(a) = \infty$ if and only if a = 0,
- 2. $\operatorname{val}(ab) = \operatorname{val}(a) + \operatorname{val}(b)$ and
- 3. $\operatorname{val}(a+b) \ge \min\{\operatorname{val}(a), \operatorname{val}(b)\}\$ for all $a, b \in K^*$.

The image of the valuation map is denoted Γ_{val} . This is an additive subgroup of the real numbers \mathbb{R} which is called the *value group*. We will usually assume that the value group Γ_{val} contains 1. Since $(\lambda \text{ val}): K \to \mathbb{R}$ is a valuation for any valuation val and $\lambda \in \mathbb{R}_{>0}$, this is not a serious restriction.

Lemma 2.1.1. If $val(a) \neq val(b)$ then val(a + b) = min(val(a), val(b)).

Proof. Without loss of generality we may assume that val(b) > val(a). Since $1^2 = 1$, we have val(1) = 0, and so $(-1)^2 = 1$ implies val(-1) = 0 as well. This implies val(-b) = val(b) for all $b \in K$. The third axiom implies

$$\operatorname{val}(a) \ge \min(\operatorname{val}(a+b), \operatorname{val}(-b)) = \min(\operatorname{val}(a+b), \operatorname{val}(b)),$$

and therefore $val(a) \ge val(a+b)$. But we also have

$$\operatorname{val}(a+b) \ge \min(\operatorname{val}(a), \operatorname{val}(b)) = \operatorname{val}(a),$$

and hence val(a + b) = val(a) as desired.

Consider the set of all field elements with non-negative valuation:

$$R_K = \{ c \in K : \operatorname{val}(c) \ge 0 \}.$$

The set R_K is a local ring. Its unique maximal ideal equals

$$\mathfrak{m}_K = \{ c \in K : \operatorname{val}(c) > 0 \}.$$

The quotient ring $\mathbb{k} = R_K / \mathfrak{m}_K$ is a field, called the *residue field* of (K, val). Our main example of a field with a valuation is the field of Puiseux series.

Example 2.1.2. Let K be the field of *Puiseux series* with coefficients in the complex numbers \mathbb{C} . The scalars in this field are the formal power series

$$c(t) = c_1 t^{a_1} + c_2 t^{a_2} + c_3 t^{a_3} + \cdots, \qquad (2.1)$$

where the c_i are nonzero complex numbers for all i, and $a_1 < a_2 < a_3 < \cdots$ are rational numbers that have a common denominator. We use the notation $\mathbb{C}\{\{t\}\}\$ for the field of Puiseux series over \mathbb{C} . Note that we can write

$$\mathbb{C}\{\{t\}\} = \bigcup_{n \ge 1} \mathbb{C}((t^{1/n})),$$

where $\mathbb{C}((t^{1/n}))$ is the field of Laurent series in the formal variable $t^{1/n}$.

This field has a natural valuation val: $\mathbb{C}\{\{t\}\} \to \mathbb{R}$ given by taking a nonzero scalar $c(t) \in \mathbb{C}\{\{t\}\}^*$ to the lowest exponent a_1 that appears in the series expansion of c(t). The field of rational functions $\mathbb{C}(t)$ is a subfield of $\mathbb{C}\{\{t\}\}$ because every rational function c(t) in one variable t has a unique expansion as a Laurent series in t. The valuation of a rational function c(t) is a positive

2.1. FIELDS

integer if c(t) has a zero at t = 0, and it is a negative integer if c(t) has a pole at t = 0. The value of val(c(t)) indicates the order of the zero or pole. Here are three examples that illustrate the inclusion of $\mathbb{C}(t)$ into $\mathbb{C}\{\{t\}\}$:

$$c(t) = \frac{4t^2 - 7t^3 + 9t^5}{6 + 11t^4} = \frac{2}{3}t^2 - \frac{7}{6}t^3 + \frac{3}{2}t^5 + \cdots \text{ has } \operatorname{val}(c(t)) = 2,$$

$$\tilde{c}(t) = \frac{14t + 3t^2}{7t^4 + 3t^7 + 8t^8} = 2t^{-3} + \frac{3}{7}t^{-2} + \cdots \text{ has } \operatorname{val}(\tilde{c}(t)) = -3,$$

$$\pi = 3.1415926535897932385.... \text{ has } \operatorname{val}(\pi) = 0.$$

We shall see in Theorem 2.1.4 that the field of Puiseux series is algebraically closed, so we also get an inclusion of $\overline{\mathbb{C}(t)}$ into $\mathbb{C}\{\{t\}\}$. Here is an illustration: Consider the two roots of the algebraic equation $x^2 - x + t = 0$. They are:

$$x_1(t) = \frac{1 + \sqrt{1 - 4t}}{2} = 1 - \sum_{k=1}^{\infty} \frac{1}{k + 1} \binom{2k}{k} t^k \text{ with } \operatorname{val}(x_1(t)) = 0,$$

$$x_2(t) = \frac{1 - \sqrt{1 - 4t}}{2} = \sum_{k=1}^{\infty} \frac{1}{k + 1} \binom{2k}{k} t^k \text{ with } \operatorname{val}(x_2(t)) = 1.$$

Combinatorialists will recognize the coefficients as *Catalan numbers*. Similarly, every univariate polynomial equation with coefficients in $\mathbb{C}(t)$ can be solved in $\mathbb{C}\{\{t\}\}$. This algorithm for computing such series solutions is implemented in computer algebra systems such as maple and Mathematica.

Remark 2.1.3. We can replace \mathbb{C} by another field \Bbbk in Example 2.1.2 and construct the field $\Bbbk\{\{t\}\}$ of Puiseux series over \Bbbk . If \Bbbk is algebraically closed of characteristic zero then so is $\Bbbk\{\{t\}\}$. However, if \Bbbk is algebraically closed of positive characteristic p, then the Puiseux series field $\Bbbk\{\{t\}\}$ would not be algebraically closed. Explicitly, if char(\Bbbk) = p > 0, then the Artin-Schreier polynomial $x^p - x - t^{-1}$ has no roots (see Remark 2.1.9 below for details).

Here is now the promised key property of the Puiseux series field:

Theorem 2.1.4. The field $K = \mathbb{k}\{\{t\}\}$ of Puiseux series is algebraically closed when \mathbb{k} is an algebraically closed field of characteristic zero.

Proof. We need to show that given a polynomial $F = \sum_{i=0}^{n} c_i x^i \in K[x]$ there is $y \in K$ with $F(y) = \sum_{i=0}^{n} c_i y^i = 0$. We shall describe an algorithm for constructing y as a Puiseux series, by successively adding higher powers of t.

We first note that we may assume the following properties of F:

- 1. $\operatorname{val}(c_i) \ge 0$ for all i,
- 2. There is some j with $val(c_j) = 0$,
- 3. $c_0 \neq 0$, and
- 4. $val(c_0) > 0$.

To see this, note that if $\alpha = \min\{\operatorname{val}(c_i) : 0 \le i \le n\}$ then multiplying F by $t^{-\alpha}$ does not change the existence of a root of F. This justifies the first two properties. If $c_0 = 0$ then y = 0 is a root so there is nothing to prove.

To make the last assumption, suppose that F satisfies the first three assumptions but $\operatorname{val}(c_0) = 0$. If $\operatorname{val}(c_n) > 0$ then we can form $G(x) = x^n F(1/x) = \sum_{i=0}^n c_{n-i}x^i$, which has the desired form, and if G(y') = 0 then F(1/y') = 0. If $\operatorname{val}(c_0) = \operatorname{val}(c_n) = 0$ then consider the polynomial $f := \overline{F} \in \mathbb{k}[x]$ that is the image of F modulo \mathfrak{m} . This is not constant since $\operatorname{val}(c_n) = 0$. Since \mathbb{k} is algebraically closed, the polynomial f has a root $\lambda \in \mathbb{k}$. Then

$$\tilde{F}(x) := F(x+\lambda) = \sum_{i=0}^{n} (\sum_{j=i}^{n} c_j {j \choose i} \lambda^{j-i}) x^i$$

has constant term $\tilde{F}(0) = F(\lambda)$ with positive valuation, and \tilde{F} still satisfies the first three properties. If y' is a root of \tilde{F} , then $y' + \lambda$ is a root of F.

Set $F_0 = F$. We will construct a sequence of polynomials $F_i = \sum_{j=0}^n c_j^i x^j$. Each of the F_i is assumed to satisfy conditions 1 to 4 above, by the same reasoning we employed for i = 0 above. The Newton polygon of F_i is the convex hull of the points $\{(i, j) :$ there is k with $k \leq i$, $val(c_k^i) \leq j\} \subset \mathbb{R}^2$. The Newton polygon has an edge with negative slope connecting the vertex $(0, val(c_0^i))$ to a vertex $(k_i, val(c_{k_i}^i))$. The absolute value of that slope equals

$$w_i = \frac{\operatorname{val}(c_0^i) - \operatorname{val}(c_{k_i}^i)}{k_i}.$$

Let f_i be the image in $\mathbb{k}[x]$ of the polynomial $t^{-\operatorname{val}(c_0^i)}F_i(t^{w_i}x) \in K[x]$. Note that f_i has degree k_i , and has nonzero constant term. Since \mathbb{k} is algebraically closed we can find a root λ_i of f_i . Let r_{i+1} be the multiplicity of λ_i as a root of f_i . Hence $f_i = (x - \lambda_i)^{r_{i+1}}g_i(x)$, where $g_i(\lambda_i) \neq 0$. We define

$$F_{i+1}(x) = t^{-\operatorname{val}(c_0^i)} F_i(t^{w_i}(x+\lambda_i)) = \sum_{j=0}^n c_j^{i+1} x^j.$$

60

2.1. FIELDS

The coefficients c_i^{i+1} of the new polynomial $F_{i+1}(x)$ are given by the formula

$$c_{j}^{i+1} = \sum_{l=j}^{n} c_{l}^{i} t^{lw_{i}-\operatorname{val}(c_{0}^{i})} {l \choose j} \lambda_{i}^{l-j}.$$
 (2.2)

The image of this Puiseux series in the residue field k equals

$$\overline{c_j^{i+1}} = \frac{1}{j!} \frac{\partial^j f_i}{\partial x^j} (\lambda_i).$$

For $0 \leq j < r_{i+1}$ this is zero, since λ_i is a root of f_i of multiplicity r_{i+1} . For $j = r_{i+1}$ this is nonzero. Thus $\operatorname{val}(c_j^{i+1}) > 0$ for $0 \leq i \leq r_{i+1}$, and $\operatorname{val}(c_j^{i+1}) = 0$ for $j = r_{i+1}$. Note that here we have used the fact that $\operatorname{char}(\Bbbk) = 0$.

If $c_0^{i+1} = 0$ then x = 0 is a root of F_{i+1} , so $\lambda_i t^{w_i}$ is root of F_i , and further back substitutions reveal that $\sum_{j=0}^i \lambda_j t^{w_0+\cdots+w_j}$ is a root of $F_0 = F$, and we are done. Thus we may assume for each *i* that $c_0^{i+1} \neq 0$, so F_{i+1} satisfies conditions 1 to 4 above. This ensures that the construction can be continued.

The observation above on $\operatorname{val}(c_j^{i+1})$ implies that $k_{i+1} \leq r_{i+1} \leq k_i$. Since n is finite, the value of k_i can only drop a finite number of times. Hence there exist $k \in \{1, \ldots, n\}$ and $m \in \mathbb{N}$ such that $k_i = k$ for all $i \geq m$. This means that $r_i = k$ for all i > m, so $f_i = \mu_i (x - \lambda_i)^k$ for all i > m, and some $\mu_i \in \mathbb{k}$.

Let N_i be such that $c_j^i \in \mathbb{k}((t^{1/N_i}))$ for $0 \leq j \leq n$. By Equation (2.2), we can take N_{i+1} to be the least common denominator of N_i and w_i . Let $y_i = \sum_{j=0}^i \lambda_j t^{w_0 + \dots + w_j} \in \mathbb{k}((t^{1/N_i}))$. We claim that $N_{i+1} = N_i$ works for i > m. Indeed, we have $w_{i+1} = \operatorname{val}(c_0^i)/k$, so it suffices to show $\operatorname{val}(c_0^i) \in \frac{k}{N_i}\mathbb{Z}$ for i > m. This follows from the fact that f_i is a pure power, so $\operatorname{val}(c_j^i) = (k - j)/k \operatorname{val}(c_0^j)$ for $1 \leq j \leq k$, and in particular $\operatorname{val}(c_{k-1}^i) = 1/k \operatorname{val}(c_0^j) \in \frac{1}{N_i}\mathbb{Z}$. Thus there is an N for which $y_i \in \mathbb{k}((t^{1/N}))$ for all i, and so the limit

$$y = \sum_{j \ge 0} \lambda_j t^{w_0 + \dots + w_j}$$
 lies in $\mathbb{k}((t^{1/N})).$

It remains to show that y is a root of F. To see this, consider $z_i = \sum_{j\geq i} \lambda_j t^{w_i+\cdots+w_j}$, and note that $y = y_{i-1} + t^{w_0+\cdots+w_{i-1}} z_i$ for i > 0. We have

$$F_i(z_i) = t^{\operatorname{val}(c_0^i)} F_{i+1}(z_{i+1}).$$

Since $z_0 = y$, it follows that

$$\operatorname{val}(F(y)) = \sum_{j=0}^{i} \operatorname{val}(c_0^j) + \operatorname{val}(F_{i+1}(z_{i+1})) > \sum_{j=0}^{i} \operatorname{val}(c_0^j) \quad \text{for all } i > 0.$$

Since $\operatorname{val}(c_0^j) \in \frac{1}{N}\mathbb{Z}$, we conclude $\operatorname{val}(F(y)) = \infty$, so F(y) = 0 as required. \Box

Remark 2.1.5. When char(\Bbbk) = 0, the Puiseux series field \Bbbk {{*t*}} is the algebraic closure of the Laurent series field \Bbbk ((*t*)). See [Rib99, 7.1.A(β), p186].

The fact that the field of Puiseux series is not algebraically closed when $\operatorname{char}(\Bbbk) > 0$ motivates the following definition. Recall that a group G is *divisible* if for all $g \in G$ and positive integers n there is a g' with ng' = g.

Example 2.1.6. Fix an algebraically closed field \mathbb{k} , and a divisible group $G \subset \mathbb{R}$. The *Mal'cev-Neumann ring* $K = \mathbb{k}((G))$ of generalized power series is the set of formal sums $\alpha = \sum_{g \in G} \alpha_g t^g$, where $\alpha_g \in \mathbb{k}$ and t is a variable, with the property that $\operatorname{supp}(\alpha) := \{g \in G : \alpha_g \neq 0\}$ is a well-ordered set.

If $\beta = \sum_{g \in G} \beta_g t^g$ then we set $\alpha + \beta = \sum_{g \in G} (\alpha_g + \beta_g) t^g$, and $\alpha\beta = \sum_{h \in G} (\sum_{g+g'=h} \alpha_g \beta_{g'}) t^h$. Then $\operatorname{supp}(\alpha + \beta) \subseteq \operatorname{supp}(\alpha) \cup \operatorname{supp}(\beta)$ is well-ordered, and thus $\alpha + \beta$ is well-defined. For $\alpha\beta$, define $\operatorname{supp}(\alpha) + \operatorname{supp}(\beta)$ to be the set $\{g + g' : g \in \operatorname{supp}(\alpha), g' \in \operatorname{supp}(\beta)\}$. This set is well-ordered, and hence $\{(g,g') : g+g'=h\}$ is finite for all $h \in G$. Thus, multiplication is well-defined. The same holds for division, so K is a field. For details, see [Pas85, Theorem 13.2.11, p. 601]. Moreover, it is known the field K is algebraically closed. For a non-constructive proof see [Poo93, Corollary 4].

Remark 2.1.7. One might be tempted to define the elements of a ring of generalized power series to be formal sums $\alpha = \sum_{g \in G} \alpha_g t^g$ with no restriction on $\operatorname{supp}(\alpha)$. However, with that definition, multiplication is not well-defined. Without the well-ordering hypothesis, the set $\{(g, g') : g + g' = h\}$ summed over in the definition of the product of two series may be infinite.

The field of generalized power series is the most general field with valuation we need to consider in the following sense.

Theorem 2.1.8. [Poo93, Corollary 5] Fix a divisible group G and an algebraically closed residue field k. Let K be a field with a valuation val with value group $G = \Gamma_{\text{val}}$ and residue field k. If val is trivial on the prime field $(\mathbb{F}_p \text{ or } \mathbb{Q})$ of K, then K is isomorphic to a subfield of $\Bbbk((G))$.

Remark 2.1.9. Consider the case when k has characteristic p > 0. Then the Artin-Schreier polynomial $x^p - x - t^{-1}$ has the roots

$$(t^{-1/p} + t^{-1/p^2} + t^{-1/p^3} + \dots) + c$$

2.1. FIELDS

where c runs over the prime field \mathbb{F}_p of \mathbb{k} . These are well-defined elements of the ring of generalized power-series, since $\{1/p^i : i \leq 0\} \cup \{0\}$ is wellordered, but they are not Puiseux series. Since there are p such roots to the Artin-Schreier polynomial, which is of degree p, and the Puiseux series are a subfield of the generalized power series, we see that there are no Puiseux series roots. Hence the Puiseux series field over \mathbb{k} is not algebraically closed. See [Ked01] for a subfield of the field of generalized power series that contains the algebraic closure of the field of Laurent series in positive characteristic.

Example 2.1.10. Let $K = \mathbb{Q}(t)$ be the algebraic closure of the field of rational functions in one variable with coefficients in \mathbb{Q} . Since $\mathbb{Q}(t) \subset \mathbb{C}((t))$ the field K is a subfield of $\mathbb{C}\{\{t\}\}$. An advantage of K over $\mathbb{C}\{\{t\}\}$ is that elements of K can be described in finite space as the roots of polynomials $g = \sum_{i=0}^{r} a_i x^i$ with coefficients $a_i \in \mathbb{Q}(t)$. This allows them to be represented in a computer. The valuation val: $K \to \mathbb{R}$ is inherited from $\mathbb{C}\{\{t\}\}$. The valuations of the roots of g can also be read from g as follows. Write $a_i = p_i/q_i$ for $1 \leq i \leq n$ where $p_i, q_i \in \mathbb{Q}[t]$. The valuation of $p = \sum_{j=0}^{s} b_j t^j \in \mathbb{Q}[t]$ is $\min\{j: b_j \neq 0\}$, and $\operatorname{val}(a_i) = \operatorname{val}(p_i) - \operatorname{val}(q_i)$. Then the valuations of the roots α of g are the $w \in \mathbb{R}$ for which the graph of the function $\min\{\operatorname{val}(a_i) + ix: 0 \leq i \leq r\}$ is not differentiable. Note that there are most r such values w. We picture this as shown in Figure 1.1. The polynomial g is replaced by an associated tropical polynomial, as in Section 1.1.

Lemma 2.1.11. Let K be algebraically closed with non-trivial valuation. Then the value group Γ_{val} is a divisible subgroup of \mathbb{R} that is dense in \mathbb{R} .

Proof. The fact that $\Gamma_{\text{val}} = \text{val}(K^*)$ is divisible follows from $\text{val}(a^{1/n}) = 1/n \text{val}(a)$. We assume for all valuations that $1 \in \Gamma_{\text{val}}$, so this means in addition that $\mathbb{Q} \subseteq \Gamma_{\text{val}}$, which implies that Γ_{val} is dense in \mathbb{R} .

Example 2.1.12. In [Mar07] Thomas Markwig proposes using a subfield of $\mathbb{k}(\mathbb{R})$) that contains the Puiseux series when $\operatorname{char}(\mathbb{k}) = 0$. His field has the advantage that the valuation map $K^* \to \mathbb{R}$ is surjective, which is not the case for the Puiseux series, since the valuation of any series is rational. \Box

We will make frequent use of the fact that the surjection $K^* \to \Gamma_{\text{val}}$ splits.

Lemma 2.1.13. If K is algebraically closed then the surjection $K^* \to \Gamma_{\text{val}}$ splits: there is a group homomorphism $\psi \colon \Gamma_{\text{val}} \to K^*$ with $\operatorname{val}(\psi(w)) = w$. Proof. Since K is algebraically closed, it contains the nth roots of all of its elements. Thus K^* , and so Γ_{val} are divisible abelian groups. Since Γ_{val} is an additive subgroup of \mathbb{R} , it is torsion-free, so Γ_{val} is a torsion-free divisible group, and thus isomorphic to a (possibly uncountable) direct sum of copies of \mathbb{Q} (see, for example, [Hun80, Exercise 8, p198]). Given any summand isomorphic to \mathbb{Q} , with $w \in \Gamma_{\text{val}}$ taken to 1 by the isomorphism, and any $a \in K^*$ with val(a) = w, there is a group homomorphism $\psi(\mathbb{Q}) \to K^*$ taking w to K^* . By construction this homomorphism satisfies $\text{val}(\psi(m/nw)) = m/nw$. The universal property of the direct product then implies the existence of a homomorphism $\Gamma_{\text{val}} \to K^*$ with the desired property. \Box

Throughout this book, we use the notation t^w to denote the element $\psi(w) \in K^*$. This is consistent with the canonical splitting for the Puiseux series field $\mathbb{C}\{\{t\}\}$. Here $\Gamma_{\text{val}} = \mathbb{Q}$, and the elements t^w are the powers of t.

Consider any field K with a valuation val: $K \to \mathbb{R} \cup \{\infty\}$. The valuation induces a *norm* $|\cdot|: K \to \mathbb{R}$ by setting $|a| = \exp(-\operatorname{val}(a))$ for $a \neq 0$, and |0| = 0. Here "exp" can be the exponential function for any base. This norm on the field K satisfies the standard norm axioms: |a| = 0 if and only if a = 0, |ab| = |a||b|, and $|a + b| \leq |a| + |b|$. The last condition can be strengthened to $|a + b| \leq \max\{|a|, |b|\}$. Norms satisfying this are called *non-archimedean*.

The norm on K allows the use of analytical and topological arguments. The field K is now a *metric space* with distance |a - b| between two elements $a, b \in K$. A *ball* is the set of all elements whose distance to a fixed element is bounded by some real constant. Our metric space K has the following remarkable property: if two balls intersect then one must be contained in the other. This structure suggest that K can be drawn as the leaves of a rooted tree, and that is why pictures of trees are ubiquitous in arithmetic geometry.

Example 2.1.14. One of the original motivations for the study of valuations is the *p*-adic valuation on the field \mathbb{Q} of rational numbers for a prime number p. The valuation val: $\mathbb{Q} \to \mathbb{R}$ given by setting $\operatorname{val}_p(q) = k$, for $q = p^k a/b$, where p does not divide a or b. For example,

$$\operatorname{val}_2(4/7) = 2$$
, $\operatorname{val}_2(3/16) = -4$.

We use this valuation to construct the completion \mathbb{Q}_p of \mathbb{Q} . Algebraically, this is the field of fractions of the completion \mathbb{Z}_p of \mathbb{Z} at the prime p. See [Eis95, Chapter 7] for details on completions. More analytically, the field \mathbb{Q}_p

2.1. FIELDS

is the completion of \mathbb{Q} with respect to the norm $|\cdot|_p$ induced by the *p*-adic valuation val_p. An element $a \in \mathbb{Q}_p$ can be written in the form

$$a = \sum_{i=m}^{\infty} a_i p^i,$$

where $a_i \in \{0, \ldots, p-1\}$ and $m \in \mathbb{Z}$. The *p*-adic integers \mathbb{Z}_p have the same representation but with $m \in \mathbb{N}$. The valuation val extends to \mathbb{Q}_p by setting val $(a) = \min\{i : a_i \neq 0\}$. This is consistent with valuation on \mathbb{Q} and the inclusion of \mathbb{Q} into \mathbb{Q}_p ; for example, val₂(6) = 1, and 6 = 1 \cdot 2^1 + 1 \cdot 2^2.

It is instructive to explore the topological properties of \mathbb{Q}_p . The ball with center 0 and radius 1 in this metric space equals \mathbb{Z}_p . The topology on \mathbb{Z}_p is fractal in nature, and, in fact, \mathbb{Z}_2 is homeomorphic to the Cantor set.

The field \mathbb{Q}_p is not algebraically closed. For instance, $x^p - x - p^{-1}$ has no roots. Its algebraic closure $\overline{\mathbb{Q}_p}$ inherits the norm but it is no longer complete. The completion of $\overline{\mathbb{Q}_p}$ is the field \mathbb{C}_p , which is both complete and algebraically closed. Performing arithmetic with scalars in these fields is a challenge. \Box

In Theorem 2.1.8 we had assumed that val is trivial on the prime field. That result does not apply to the fields K in Example 2.1.14, where the prime field is \mathbb{Q} but with the *p*-adic valuation. There exists a generalization of the field $\mathbb{k}((G))$ of generalized power series which allows an extension of Theorem 2.1.8 to the case where val is the *p*-adic valuation on \mathbb{Q} . However, the arithmetic in such fields is really tricky. See [Poo93] for details.

We close this section with a remark about computational issues. It is impossible to enter a generalized power series or arbitrary Puiseux series into a computer, as it cannot be described by a finite amount of information. This suggests that the best pure characteristic zero field we can hope to compute with the algebraic closure $\overline{\mathbb{Q}(t)}$ of the ring of rational functions in t with coefficients in \mathbb{Q} . Almost all of the examples in the examples will be defined and computed over the field $\mathbb{Q}(t)$ of rational functions.

A typical computation one may wish to do is compute a Gröbner basis of a homogeneous ideal in a polynomial ring, as in Section 2.4 below, or perhaps even a tropical basis of an ideal in a Laurent polynomial ring, as in Section 2.5 below. If $K = \mathbb{Q}(t)$ then this computation can be reduced to working over the field of constants $\mathbb{k} = \mathbb{Q}$. Namely, given an ideal $I \subset \mathbb{Q}(t)[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, we may consider instead do our computation for $I' = I \cap \mathbb{Q}[t^{\pm 1}, x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$.

2.2 Algebraic Varieties

For the first four chapters of this book we assume no background in algebraic geometry beyond that covered in the highly recommended undergraduate text book by Cox, Little and O'Shea [CLO07]. We now recall their notation.

Definition 2.2.1. Let K be a field. Affine space over K of dimension n is

$$\mathbb{A}_{K}^{n} = \mathbb{A}^{n} = \{(a_{1}, a_{2}, \dots, a_{n}) : a_{i} \in K\} = K^{n}.$$

The *n*-dimensional *projective space* over the field K is

$$\mathbb{P}^n_K = \mathbb{P}^n = (K^{n+1} \setminus \mathbf{0}) / \sim$$

where $\mathbf{v} \sim \lambda \mathbf{v}$ for all $\lambda \neq 0$. The points of \mathbb{P}^n are the equivalence classes of lines through the origin **0**. We write $[v_0 : v_1 : \cdots : v_n]$ for the equivalence class of $\mathbf{v} = (v_0, v_1, \ldots, v_n) \in K^{n+1}$. The *n*-dimensional algebraic torus is

$$\mathbb{T}_{K}^{n} = \mathbb{T}^{n} = \{ (a_{1}, a_{2}, \dots, a_{n}) : a_{i} \in K^{*} \}.$$

Definition 2.2.2. The coordinate ring of the affine space \mathbb{A}^n is the polynomial ring $K[x_1, \ldots, x_n]$. The homogeneous coordinate ring of the projective space \mathbb{P}^n is $K[x_0, x_1, \ldots, x_n]$, and the coordinate ring of the algebraic torus \mathbb{T}^n is the Laurent polynomial ring $K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$.

The affine variety defined by an ideal $I \subset K[x_1, \ldots, x_n]$ is

$$V(I) = \{a \in \mathbb{A}_K^n : f(a) = 0 \text{ for all } f \in I\}.$$

An ideal $I \subset K[x_0, \ldots, x_n]$ is homogeneous if it has a generating set consisting of homogeneous polynomials. The projective variety defined by a homogeneous ideal $I \subset K[x_0, \ldots, x_n]$ is

$$V(I) = \{ x \in \mathbb{P}_K^n : f(x) = 0 \text{ for all } f \in I \}.$$

Any ideal I in $K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ defines a very affine variety in the torus:

$$V(I) = \{ x \in \mathbb{T}_K^n : f(x) = 0 \text{ for all } f \in I \}.$$

For any variety X we consider the ideal I_X of all polynomials (or homogeneous polynomials, or Laurent polynomials) that vanish on X. The *coordinate ring* K[X] of a variety X is the quotient of the coordinate ring of the ambient space, namely \mathbb{A}^n , \mathbb{P}^n or \mathbb{T}^n , by the defining ideal I_X . In tropical geometry, we are mostly concerned with Laurent polynomials and the very affine varieties they define. Frequently, our ground field will be $K = \mathbb{C}$, the complex numbers. Very affine affine varieties are non-compact, as was discussed in Section 1.8. Here is one more example along these lines.

Example 2.2.3. Let $K = \mathbb{C}$, n = 3, and $f = f(x_1, x_2, x_3)$ and $g = g(x_1, x_2, x_3)$ be random polynomials of degree two, and consider the ideal $I = \langle f, g \rangle$ they generate in the Laurent polynomial ring $\mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}, x_3^{\pm 1}]$. The very affine variety V(I) is a curve in the torus $\mathbb{T}^3_{\mathbb{C}} = (\mathbb{C}^*)^3$. This is an elliptic curve with 16 punctures. In other words, V(I) is a non-compact Riemann surface of genus g = 1 but with m = 16 points removed. Students of number theory may wish to contemplate the same example for $K = \mathbb{Q}_p$. \Box

The map that takes an ideal to its variety is not a bijection; for example, $V(\langle x \rangle) = V(\langle x^2 \rangle) \subset A^1$. Two ideals I and J satisfy V(J) = V(I) if have the same radical $\sqrt{J} = \sqrt{I}$. When K is algebraically closed, *Hilbert's Nullstel lensatz* states that $\sqrt{I} = I_X$ where X = V(I) is the variety of I. For details, see any book on commutative algebra (for example, [Eis95], or [CLO07]).

Experts should note that we do not assume that our varieties are irreducible. A variety X is *irreducible* if it cannot be written as the union of two proper subvarieties. Every variety can be decompose into a finite union of irreducible varieties. This can be computed algebraically (e.g. in Macaulay2) by means of the *primary decomposition* of the corresponding ideals. If X is an irreducible variety then its vanishing ideal I_X is a prime ideal.

The simplest prime ideals are those generated by linear polynomials. The corresponding varieties are called *linear spaces*. An ideal is *principal* if it is generated by one polynomial, and in this case the variety is a *hypersurface*. Hypersurfaces are varieties of codimension one. The *dimension* of a variety is its most basic invariant. The *codimension* is n minus the dimension. See Chapter 9 in [CLO07] for the definition of dimension and how to compute it.

Linear algebra furnishes many interesting examples of varieties. For example, the set X of all $m \times n$ -matrices of rank $\leq r$ is an irreducible variety. Its prime ideal I_X is generated by all $(r+1) \times (r+1)$ -minors of an $m \times n$ -matrix of variables. Such varieties are called *determinantal varieties*, and they frequently occur as irreducible components of other interesting varieties.

Example 2.2.4. Let n = 8, fix any field K, and consider the affine space

 $\mathbb{A}^8 = \mathbb{A}^8_K$ whose points are pairs (A, B) of 2×2-matrices with entries in K:

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \text{ and } B = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}.$$

The commuting variety is defined by the matrix equation $A \cdot B = B \cdot A$. This variety is irreducible and has dimension 4. (It's a complete intersection). The four matrix entries of the commutator $A \cdot B - B \cdot A$ generate a prime ideal.

For an example with different properties consider the 5-dimensional variety defined by the matrix equation $A \cdot B = 0$. This is the smallest instance of what is known as a *variety of complexes*. Its radical ideal equals

$$I = \langle a_{11}b_{11} + a_{12}b_{21}, a_{11}b_{12} + a_{12}b_{22}, a_{21}b_{11} + a_{22}b_{21}, a_{21}b_{12} + a_{22}b_{22} \rangle.$$

This ideal I has the prime decomposition

$$(I + \langle a_{11}a_{22} - a_{12}a_{21}, b_{11}b_{22} - b_{12}b_{21} \rangle) \cap \langle b_{11}, b_{21}, b_{12}, b_{22} \rangle \cap \langle a_{11}, a_{12}, a_{21}, a_{22} \rangle$$

Hence the variety has three irreducible components, corresponding to what the ranks of A and B are. The components have dimensions 5, 4 and 4.

Tropical geometers would study the variety $\{A \cdot B = 0\}$ not in affine space \mathbb{A}^8 but in the torus \mathbb{T}^8 . The variety in \mathbb{T}^8 is irreducible because the components $\{A = 0\}$ and $\{B = 0\}$ disappear. In terms of algebra, the ideal I is a prime ideal in the Laurent polynomial ring $\mathbb{C}[a_{11}^{\pm 1}, a_{12}^{\pm 1}, \ldots, b_{22}^{\pm 1}]$. \Box

We place a topology on affine space \mathbb{A}^n by taking the closed sets to be $\{V(I) : I \text{ is an ideal of } K[x_1, \ldots, x_n]\}$. This is the Zariski topology. To check that \emptyset and \mathbb{A}^n are closed, note that $\emptyset = V(1)$, and $\mathbb{A}^n = V(0)$. It is an exercise to check that the finite union of closed sets and the arbitrary intersection of closed sets are closed. We denote by \overline{U} the closure in the Zariski topology of a set U. This is the smallest set of the form V(I) for some I that contains U. Similarly we can define the Zariski topology on \mathbb{P}^n and \mathbb{T}^n .

There are inclusions $\mathbb{T}^n \xrightarrow{i} \mathbb{A}^n \xrightarrow{j} \mathbb{P}^n$, where the second map sends $x \in \mathbb{A}^n$ to $(1:x) \in \mathbb{P}^n$. The *affine closure* of a variety $X \subset \mathbb{T}^n$ is the Zariski closure $\overline{i(X)}$ of $\underline{i(X)} \subset \mathbb{A}^n$. The *projective closure* of a variety $X \subset \mathbb{A}^n$ is the Zariski closure $\overline{j(X)}$ of $j(X) \subset \mathbb{P}^n$. We now recall their algebraic descriptions.

Definition 2.2.5. The degree of a polynomial $f = \sum c_u x^u$ in $K[x_1, \ldots, x_n]$ is $W = \max\{|u| : c_u \neq 0\}$, where $|u| = \sum_{i=1}^n u_i$. The homogenization \tilde{f} of fis the homogeneous polynomial $\tilde{f} = \sum c_u x_0^{W-|u|} x^u \in K[x_0, x_1, \ldots, x_n]$. Given an ideal I in $K[x_1, \ldots, x_n]$, its homogenization is the ideal $I_{\text{proj}} = \langle \tilde{f} : f \in I \rangle$.

2.2. ALGEBRAIC VARIETIES

Proposition 2.2.6. Let X = V(I) be a subvariety of the torus \mathbb{T}^n for an ideal $I \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Then $\overline{i(X)} = V(I_{\text{aff}})$, where $I_{\text{aff}} = I \cap K[x_1, \ldots, x_n]$. For an ideal $I \subset K[x_1, \ldots, x_n]$, the projective closure $\overline{j(X)}$ of V(I) is the subvariety of projective space \mathbb{P}^n defined by the homogeneous ideal I_{proj} .

Proof. The sets $V(I_{\text{aff}})$ and $V(I_{\text{proj}})$ are Zariski closed subsets of \mathbb{A}^n and \mathbb{P}^n respectively that contain i(X) and j(X), so they contain $\overline{i(X)}$ and $\overline{j(X)}$. Conversely, suppose that $f \in K[x_1, \ldots, x_n]$ vanishes on $\overline{i(X)}$. Then f(y) = 0 for all $y \in X$, so $f \in I(X)$ when regarded as a polynomial in $K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and thus $f \in I_{\text{aff}}$. Similarly, if a homogeneous polynomial $g \in K[x_0, \ldots, x_n]$ vanishes on the projective variety $\overline{j(X)}$ then $g(1, y_1, \ldots, y_n) = 0$ for all $y = (y_1, \ldots, y_n) \in X$, so $g(1, x) \in I$, and thus $g \in I_{\text{proj}}$.

Example 2.2.7. Consider the very affine variety X = V(I) in \mathbb{T}^3 defined by

$$I = \left\langle \frac{1}{x_1} + \frac{1}{x_2} + \frac{1}{x_3} - 1, \frac{1}{x_1} + \frac{2}{x_2} + \frac{3}{x_3} \right\rangle.$$

Then its affine closure $\overline{i(X)} = V(I_{\text{aff}})$ in \mathbb{A}^3 is defined by the ideal

$$I_{\text{aff}} = I \cap K[x_1, x_2, x_3] = \langle x_2 x_3 + 2x_2 + x_3, 2x_1 x_3 + x_1 - x_3 \rangle,$$

and its projective closure $\overline{j(X)} = V(I_{\text{proj}})$ in \mathbb{P}^3 is defined by the ideal

$$I_{\text{proj}} = \langle x_2 x_3 + 2x_0 x_2 + x_0 x_3, \ 2x_1 x_3 + x_0 x_1 - x_0 x_3, \ 3x_1 x_2 - x_0 x_1 - 2x_0 x_2 \rangle.$$

Such computations are based on *ideal quotients* as in [CLO07, $\S4.4$].

A morphism $\phi : X \to Y$ of affine or very affine varieties is induced by a ring homomorphism $\phi^* : K[Y] \to K[X]$ between the respective coordinate rings. Note that the homomorphism ϕ^* takes the coordinate ring of Y to that of X. The transformation $X \mapsto K(X)$ is a contravariant functor. Computing the image of a morphism is known as *implicatization* (cf. Section 1.5).

For a morphism of tori $\phi : \mathbb{T}^n \to \mathbb{T}^m$ we place the additional constraint that the map ϕ be a homomorphism of algebraic groups. This means that the map $\phi^* : K[x_1^{\pm 1}, \ldots, x_m^{\pm 1}] \to K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ is (after appropriate choice of coordinates) a monomial map, so $\phi^*(x_i)$ is a monomial for $1 \leq i \leq m$. Equivalently, the ring homomorphism ϕ^* is induced by a group homomorphism, which we also denote by ϕ^* , from \mathbb{Z}^m to \mathbb{Z}^n . If X = V(I) is a subvariety of \mathbb{T}^n , then the Zariski closure $\phi(X)$ of $\phi(X)$ in \mathbb{T}^m is the variety $V(\phi^{*-1}(I))$. **Example 2.2.8.** Let n = 1 and m = 2. The curve in Example 1.5.1 is given as the image of a morphism Φ of very affine varieties $\mathbb{T}^1 \to \mathbb{T}^2$. However, that morphism is not a morphism of tori because the coordinates of Φ are not monomials in t. A morphism of tori has the form $\phi : \mathbb{T}^1 \to \mathbb{T}^2$, $t \mapsto (t^a, t^b)$, where a and b are integers. Assuming that a and b are relatively prime, the image of ϕ is the binomial curve $\{(x, y) \in \mathbb{T}^2 : x^b - y^a = 0\}$. \Box

Recall that the group of group automorphisms of the lattice \mathbb{Z}^n is isomorphic to $\operatorname{GL}(n,\mathbb{Z})$, the group of invertible matrices with integer entries and determinant ± 1 . We denote by $\mathbf{e}_1, \ldots, \mathbf{e}_n$ the standard basis for \mathbb{Z}^n .

Lemma 2.2.9. Given any vector $\mathbf{v} \in \mathbb{Z}^n$ with the greatest common divisor of the $|v_i|$ equal to one, there is a matrix $U \in \operatorname{GL}(n,\mathbb{Z})$ with $U\mathbf{v} = \mathbf{e}_1$. Further, if L is a rank k subgroup of \mathbb{Z}^n with \mathbb{Z}^n/L torsion-free then there is a matrix $U \in \operatorname{GL}(n,\mathbb{Z})$ with UL equal to the subgroup generated by $\mathbf{e}_1, \ldots, \mathbf{e}_k$.

Proof. The first statement follows from the second, as if the greatest common divisor of the $|v_i|$ is one, the group \mathbb{Z}^n/\mathbf{v} is torsion-free. Let A be a $k \times n$ matrix with rows an integer basis for the subgroup L. The condition that \mathbb{Z}^n/L is torsion-free implies that the Smith normal form of A is the $k \times n$ matrix A' with first $k \times k$ block the identity matrix, and all other entries zero. There are matrices $V \in \operatorname{GL}(k,\mathbb{Z}), U' \in \operatorname{GL}(n,\mathbb{Z})$ with A' = VAU'. Multiplying on the left by an element of $\operatorname{GL}(k,\mathbb{Z})$ does not change the integer row span, so the integer row span of VA equals L. We now take $U = U'^T$. \Box

An automorphism of the torus \mathbb{T}^n is an invertible map specified by nLaurent monomials in x_1, \ldots, x_n . Thus the automorphism group of \mathbb{T}^n is isomorphic to $\operatorname{GL}(n, \mathbb{Z})$. Here the matrix entries are the exponents of the monomials. We we speak of a coordinate change in \mathbb{T}^n we mean the transformation given by such an invertible monomial map. These multiplicative changes of variables behave very differently from the more familiar linear changes of variables in affine space \mathbb{A}^n or projective space \mathbb{P}^n . Automorphisms of \mathbb{T}^n are essential for tropical geometry, and we already encountered them in Bergman's solution to Zalessky's problem in Corollary 1.4.3.

Example 2.2.10. The invertible integer map $U = \begin{pmatrix} 1 & -1 \\ 1 & -2 \end{pmatrix}$ represents the automorphism $(x, y) \mapsto (xy, x^{-1}y^{-2})$ of the torus \mathbb{T}^2 and of its coordinate ring $\mathbb{C}[x^{\pm 1}, y^{\pm 1}]$. The image of the curve $X = \{(x, y) \in \mathbb{T}^2 : f(x, y) = 0\}$ in

Example 1.8.1 under the automorphism U is the curve defined by

$$(U \circ f)(x, y) = c_2 + c_5 x + c_1 y + c_3 x y + c_4 x^2 y^2$$

Note how the linear map U moves the tropical curve $\operatorname{trop}(X)$. The compactifications $X^{\text{hom}} = \overline{j(X)} \subset \mathbb{P}^2$ and $X^{\text{bihom}} \subset \mathbb{P}^1 \times \mathbb{P}^1$ are changed under this automorphism, but the tropical compactification X^{trop} remains the same. \Box

2.3 Polyhedral Geometry

We review here the notions from polyhedral geometry that are needed in this book. Polyhedral geometry is a rich and beautiful area of discrete mathematics. The reader unfamiliar with this area is urged to spend some time with a reference such as the first few chapters of Ziegler's text book [Zie95].

Definition 2.3.1. A set $X \subseteq \mathbb{R}^n$ is *convex* if for all $\mathbf{u}, \mathbf{v} \in X$ and all $0 \leq \lambda \leq 1$ we have $\lambda \mathbf{u} + (1 - \lambda)\mathbf{v} \in X$. The *convex hull* conv(U) of a set $U \subseteq \mathbb{R}^n$ is the smallest convex set containing U. If $U = {\mathbf{u}_1, \ldots, \mathbf{u}_r}$ is finite then $conv(U) = {\sum_{i=1}^r \lambda_i \mathbf{u}_i : 0 \leq \lambda_j \leq 1, \sum_{i=1}^r \lambda_i = 1}$ is called a *polytope*.

A polyhedral cone in \mathbb{R}^n is the positive hull of a finite set of vectors in \mathbb{R}^n :

$$C = \operatorname{pos}(\mathbf{v}_1, \dots, \mathbf{v}_r) := \{\sum_{i=1}^r \lambda_i \mathbf{v}_i : \lambda_i \ge 0\}.$$

Every polyhedral cone has the alternate description as a set of the form

$$C = \{ \mathbf{x} : A\mathbf{x} \le 0 \}$$

where A is a $d \times n$ matrix. For a proof see [Zie95, Theorem 1.3].

A face of a cone is determined by a linear functional $w \in \mathbb{R}^{n \vee}$, via

$$face_w(C) = \{ \mathbf{x} \in C : w \cdot \mathbf{x} \le w \cdot \mathbf{y} \text{ for all } \mathbf{y} \in C \}$$

This has the alternate description as $face_w(C) = {\mathbf{x} \in C : A'\mathbf{x} = 0}$, where A' is a $d' \times n$ submatrix of A. A polyhedral fan is a collection of polyhedral cones, the intersection of any two of which is a face of each.

A convex set is, by definition, the intersection of half spaces in some \mathbb{R}^n . A polyhedron $P \subset \mathbb{R}^n$ is the intersection of finitely many closed half spaces:

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



Figure 2.1: Polyhedral fans



Figure 2.2: Not a polyhedral fan

where A is a $d \times n$ matrix, and $\mathbf{b} \in \mathbb{R}^d$. Polytopes are those polyhedra that are bounded [Zie95, §1.1]. A face of a polyhedron is determined by a linear functional $w \in \mathbb{R}^{n^{\vee}}$, via $\operatorname{face}_w(P) = \{\mathbf{x} \in P : w \cdot \mathbf{x} \leq w \cdot \mathbf{y} \text{ for all } \mathbf{y} \in P\}$. A face of P that is not contained in any larger proper face is called a *facet*. A *polyhedral complex* is a collection of polyhedra satisfying two conditions: if a polyhedron P is in the collection, then so is any face of P, and if P and Q lie in the collection then $P \cap Q$ is a face of both P and Q. The polyhedra in a polyhedral complex Σ that are not faces of any larger polyhedra are called *facets* of the complex. Their facets are called ridges of the complex. The support $\sup(\Sigma)$ of a polyhedral complex Σ is the set $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \in$ P for some $P \in \Sigma\}$. Polyhedral complexes.

The lineality space of a polyhedron is the largest affine subspace contained in P. Equivalently, it is the largest subspace $V \subset \mathbb{R}^n$ for which $\mathbf{x} + \mathbf{v} \in P$ for all $\mathbf{x} \in P, \mathbf{v} \in V$. The lineality space of a polyhedral complex Σ is the intersection of all the lineality spaces of the polyhedra in the complex. The *affine span* of a polyhedron P is the smallest affine subspace containing P. The *dimension* of P is the dimension of its affine span. A polyhedral complex Σ is *pure* of dimension d if every polyhedron in Σ that is not the face of any other polyhedron in Σ has dimension d. The *relative interior* of P is the


Figure 2.3: A polyhedral complex

interior of P inside its affine span. If $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, A'\mathbf{x} \leq \mathbf{b}'\}$, where each of the inequalities represented by $A'\mathbf{x} \leq \mathbf{b}'$ can be strict for some $\mathbf{x} \in P$, then relint $(P) = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, A'\mathbf{x} < \mathbf{b}'\}$.

Definition 2.3.2. Let Γ be a subgroup of \mathbb{R} . A Γ -rational polyhedron is

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

where A is a $d \times n$ matrix with entries in \mathbb{Q} , and $\mathbf{b} \in \Gamma^d$. A polyhedral complex Σ is Γ -rational if every polyhedron in Σ is Γ -rational. We will be interested in the case where $\Gamma = \Gamma_{\text{val}}$ is the value group of a field K.

Definition 2.3.3. Let $P \subset \mathbb{R}^n$ be a polyhedron. The *normal fan* of P is the polyhedral fan \mathcal{N}_P consisting of the cones

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

as F varies over the faces of P where $cl(\cdot)$ denotes the closure in the Euclidean topology on \mathbb{R}^n . The fan \mathcal{N}_P is also called the *inner normal fan* of P.

Definition 2.3.4. Let $S = K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ be the Laurent polynomial ring. Given $f = \sum_{\mathbf{u} \in \mathbb{Z}^n} c_u \mathbf{x}^{\mathbf{u}} \in S$, the Newton polytope of f is the polytope

Newt
$$(f) = \operatorname{conv}(\mathbf{u} : c_u \neq 0) \subset \mathbb{R}^n$$
.

If Newt(f) is 2-dimensional then we call it the *Newton polygon*. This notion of Newton polygon differs from the one used in the proof of Theorem 2.1.4. \Box

Example 2.3.5. Let $S = \mathbb{C}[x^{\pm 1}, y^{\pm 1}]$, and let $f = 7x + 8y - 3xy + 4x^2y - 17xy^2 + x^2y^2$. The Newton polygon of f is shown in Figure 2.5. Let $g = x^{-1} - y^{-1} + 3x - 2y + xy$. The Newton polygon of g is the translation of that for f by the vector (-1, -1). The same Newton polygon arises, up to an automorphism of \mathbb{Z}^2 , from the polynomials in Examples 1.8.1 and 2.2.10. \Box



Figure 2.4: The normal fan of a polyhedron ${\cal P}$



Figure 2.5: The Newton polytope of $7x + 8y - 3xy + 4x^2y - 17xy^2 + x^2y^2$

2.3. POLYHEDRAL GEOMETRY

Let Σ_1 and Σ_2 be two polyhedral complexes with the same support. The common refinement of Σ_1 and Σ_2 is the polyhedral complex $\Sigma_1 \wedge \Sigma_2$ consisting of the polyhedra $\{P \cap Q : P \in \Sigma_1, Q \in \Sigma_2\}$. This operation does not change the support, so we have $\operatorname{supp}(\Sigma_1 \wedge \Sigma_2) = \operatorname{supp}(\Sigma_1) = \operatorname{supp}(\Sigma_2)$.

The Minkowski sum of two subsets $A, B \subset \mathbb{R}^n$ is the set

$$A + B = \{a + b : a \in A, b \in B\}.$$

If A and B are polyhedra in \mathbb{R}^n then A + B is also a polyhedron in \mathbb{R}^n . The same holds for polytopes, cones and supports of polyhedral complexes. Here are two useful facts that relates Minkowski sums to other constructions.

• If P and Q are polyhedra in \mathbb{R}^n then the normal fan of their Minkowski sum is the common refinement of the two normal fans:

$$\mathcal{N}_{P+Q} = \mathcal{N}_P \wedge \mathcal{N}_Q. \tag{2.3}$$

• The Newton polytope of a product of two Laurent polynomials is the Minkowski sum of the two given Newton polytopes:

$$\operatorname{Newt}(f \cdot g) = \operatorname{Newt}(f) + \operatorname{Newt}(g). \tag{2.4}$$

Definition 2.3.6. Let Σ be a polyhedral complex in \mathbb{R}^n , and let σ be a polyhedron in Σ . The *star* of $\sigma \in \Sigma$ is a fan in \mathbb{R}^n , denoted $\operatorname{star}_{\Sigma}(\sigma)$, whose cones are indexed by those $\tau \in \Sigma$ for which σ is a face of τ . Fix $w \in \sigma$. Then the cone of $\operatorname{star}_{\Sigma}(\sigma)$ that is indexed by τ is the Minkowski sum

$$\bar{\tau} = \{ v \in \mathbb{R}^n : \exists \epsilon > 0 \text{ with } w + \epsilon v \in \tau \} + \operatorname{aff}(\sigma) - w.$$

This is independent of the choice of w.

Example 2.3.7. The polyhedral complex Σ shown on the left of Figure 2.6 is a quadratic curve in the tropical plane, as seen in Section 1.3. The affine span of the vertex σ_1 in Σ is just the vertex itself. The star is shown on the right. For σ_2 the affine span is the entire *y*-axis, and this is also the star.

Computationally inclined readers may wonder what software packages are available for computing with polytopes and polyhedra. An excellent general purpose platform is the software polymake due to Evgeny Gavrilov and Michael Joswig. For the specific study of polyhedral complexes and fans arising in tropical geometry, we recommend Anders Jensen's software GFan.



Figure 2.6: The star of polyhedron in a polyhedral complex

2.4 Gröbner Bases

In this section we introduce Gröbner bases over a field K with a valuation val. This is a generalization of the Gröbner basis theory familiar from [CLO07] and other text books, such as [Eis95]. We do not require K to be algebraically closed, but we will assume that the valuation is nontrivial, that the value group Γ_{val} is dense in \mathbb{R}^n , and that there is a splitting $\phi: \Gamma_{\text{val}} \to K^*$ which we denote by $\phi(w) = t^w$. If $\text{val}(a) \ge 0$, so a lies in the valuation ring R of K, we denote by \overline{a} the image of a in the residue field \Bbbk . We begin by considering the case of a homogeneous ideal I of the polynomial ring $S = K[x_0, x_1, \ldots, x_n]$. Our primary goal is to introduce a new concept: the Gröbner complex of I.

For a polynomial $f = \sum_{u \in \mathbb{N}^{n+1}} c_u x^u \in S$, the tropicalization of f is the function $\operatorname{trop}(f) \colon \mathbb{R}^n \to \mathbb{R}$ given by $w \mapsto \min(\operatorname{val}(c_u) + w \cdot u)$. Fix $w \in (\Gamma_{\operatorname{val}})^{n+1}$ and $W = \operatorname{trop}(f)(w) = \min\{\operatorname{val}(c_u) + w \cdot u : c_u \neq 0\}$. We set

$$\operatorname{in}_{w}(f) = \overline{t^{-W} \sum_{u \in \mathbb{N}^{n+1}} c_{u} t^{w \cdot u} x^{u}} \in \mathbb{K}[x_{0}, \dots, x_{n}].$$

This is the *initial form* of f with respect to w, and it is equal to

$$\operatorname{in}_{w}(f) = \sum_{\substack{u \in \mathbb{N}^{n+1}: \\ \operatorname{val}(c_{u})+w \cdot u = W}} \overline{c_{u}t^{-\operatorname{val}(c_{u})}} x^{u} = \overline{t^{-\operatorname{trop}(f)(w)}f(t^{w_{1}}x_{1}, \dots, t^{w_{n}}x_{n})}$$

Example 2.4.1. Let $f = (t + t^2)x_0 + 2t^2x_1 + 3t^4x_2 \in \mathbb{C}\{\{t\}\}[x_0^{\pm 1}, x_1^{\pm 1}, x_2^{\pm 1}].$ If w = (0, 0, 0) then W = 1 and $\operatorname{in}_w(f) = (1 + t)x_0 = x_0$. If w = (4, 2, 0) then W = 4 and $\operatorname{in}_w(f) = 2x_1 + 3x_2$. Also, $\operatorname{in}_{(2,1,0)}(f) = x_0 + 2x_1$.

2.4. GRÖBNER BASES

If I is a homogeneous ideal in $K[x_0, \ldots, x_n]$, then its *initial ideal* is

$$\operatorname{in}_w(I) = \langle \operatorname{in}_w(f) : f \in I \rangle \subset \Bbbk[x_0, \dots, x_n]$$

Note that $\operatorname{in}_w(I)$ is an ideal in $\Bbbk[x_0, \ldots, x_n]$. A set $\mathcal{G} = \{g_1, \ldots, g_s\} \subset I$ is a *Gröbner basis* for I with respect to w if $\operatorname{in}_w(I) = \langle \operatorname{in}_w(g_1), \ldots, \operatorname{in}_w(g_s) \rangle$.

Lemma 2.4.2. Let $I \subset K[x_0, \ldots, x_n]$ be a homogeneous ideal and $w \in (\Gamma_{\text{val}})^{n+1}$. Then $\text{in}_w(I)$ is homogeneous, and we may choose a homogeneous Gröbner basis for I. Further, if $g \in \text{in}_w(I)$ then $g = \text{in}_w(f)$ for some $f \in I$.

Proof. To see that $\operatorname{in}_w(I)$ is homogeneous, note that given $f = \sum_{i \ge 0} f_i \in S$ with each f_i homogeneous of degree i we have $\operatorname{in}_w(f) = \sum_{i \in \tau} \operatorname{in}_w(f_i)$, where the sum is over those i with $\operatorname{min}\{\operatorname{val}(c_u) + w \cdot u : c_u \neq 0\}$ minimal, so $\operatorname{in}_w(I)$ is generated by the homogeneous elements $\operatorname{in}_w(f)$ with f homogeneous. This also shows that we may choose our Gröbner basis to consist of homogeneous polynomials. For the last claim, let $g = \sum a_u x^u \operatorname{in}_w(f_u) \in \operatorname{in}_w(I)$, with $f_u \in I$ for all u. Then $g = \sum a_u \operatorname{in}_w(x^u f_u)$. For each a_u choose a lift $c_u \in K$, and let $W_u = \operatorname{trop}(f_u)(w)$. Let $f = \sum_u c_u t^{-W_u - \operatorname{val}(c_u)} x^u f_u$. Then by construction $\operatorname{trop}(f)(w) = 0$, and $\operatorname{in}_w(f) = \sum_u a_u x^u \operatorname{in}_w(f) = g$. \Box

Example 2.4.3. Fix the field $K = \mathbb{Q}$ with the 2-adic valuation, so $\mathbb{k} = \mathbb{Z}/2\mathbb{Z}$. Let n = 3 and consider the line in \mathbb{P}^3_K defined by the ideal of linear forms

$$I = \langle x_0 + 2x_1 - 3x_2, 3x_1 - 4x_2 + 5x_3 \rangle.$$

If w = (0, 0, 0, 0) then the two generators are a Gröbner basis and $\operatorname{in}_w(I) = \langle x_0 + x_2, x_1 + x_3 \rangle$. This is an ideal over $\Bbbk = \mathbb{Z}/2\mathbb{Z}$. If w = (1, 0, 0, 1) then $\operatorname{in}_w(I) = \langle x_1, x_2 \rangle$, and a Gröbner basis is $\{x_2 - 3x_0 + 10x_3, x_1 - 4x_0 + 15x_3\}$. How many distinct initial ideals can we get as w varies over $\Gamma_{\operatorname{val}}^4 = \mathbb{Z}^4$? \Box

The definitions of $\operatorname{in}_w(f)$ and $\operatorname{in}_w(I)$ extend naturally to the case when fand I are taken from the polynomial ring $\Bbbk[x_0, \ldots, x_n]$ over the residue field \Bbbk , and $w \in (\Gamma_{\operatorname{val}})^{n+1}$. This is obvious if K contains \Bbbk . Otherwise we choose a field K' with a nontrivial valuation containing \Bbbk for which the valuation is trivial on \Bbbk , and the residue field is \Bbbk . For example, one option is to take K' to be a ring of generalized power series with coefficients in \Bbbk and value group $\Gamma_{\operatorname{val}}$. Note that for $I \subset \Bbbk[x_0, \ldots, x_n]$ we have $\operatorname{in}_w(I') = \operatorname{in}_w(I)$ where $I' = IK'[x_0, \ldots, x_n]$. This means that any result that assumes that I is a homogeneous ideal in a polynomial ring with coefficients in a field with a nontrivial valuation with $\Gamma_{\operatorname{val}}$ dense in \mathbb{R} also applies to ideals in $\Bbbk[x_0, \ldots, x_n]$. **Lemma 2.4.4.** Let $I \subseteq K[x_0, \ldots, x_n]$, and fix $w, v \in \Gamma_{val}^{n+1}$. Then there exists $\epsilon > 0$ such that the following holds for all $\epsilon' \in \Gamma_{val}$ with $0 < \epsilon' < \epsilon$:

$$\operatorname{in}_{v}(\operatorname{in}_{w}(I)) = \operatorname{in}_{w+\epsilon'v}(I).$$

Proof. We claim that it suffices to check that for all $f \in K[x_0, \ldots, x_n]$ there is $\epsilon > 0$ such that $\operatorname{in}_v(\operatorname{in}_w(f)) = \operatorname{in}_{w+\epsilon'v}(f)$ for all $\epsilon' \in \Gamma_{\operatorname{val}}$ with $0 < \epsilon' < \epsilon$. To see this, note that $\operatorname{in}_v(\operatorname{in}_w(I))$ has a finite generating set $\{g_1, \ldots, g_s\} \subset \mathbb{k}[x_0, \ldots, x_n]$ with each generator g_i of the form $\operatorname{in}_v(\operatorname{in}_w(f_i))$ for some $f_i \in I$. We choose ϵ to be the minimum of the ϵ_i corresponding to these generating f_i . Then $g_i = \operatorname{in}_v(\operatorname{in}_w(f_i)) = \operatorname{in}_{w+\epsilon'v}(f_i)$, so $\operatorname{in}_v(\operatorname{in}_w(I)) \subseteq \operatorname{in}_{w+\epsilon'v}(I)$, for any $\epsilon' < \epsilon$. Conversely, the ideal $\operatorname{in}_{w+\epsilon'v}(I)$ is finitely generated by $h_1, \ldots, h_r \in \mathbb{k}[x_0, \ldots, x_n]$, with $h_i = \operatorname{in}_{w+\epsilon'v}(f'_i)$ for some $f'_i \in I$. Then $\operatorname{in}_{w+\epsilon'v}(f'_i) = \operatorname{in}_v(\operatorname{in}_w(f'_i)) \in \operatorname{in}_v(\operatorname{in}_w(I))$, and we conclude $\operatorname{in}_{w+\epsilon'v}(I) \subseteq \operatorname{in}_v(\operatorname{in}_w(I))$.

We now prove the lemma for a single polynomial $f = \sum_{u \in \mathbb{N}^{n+1}} c_u x^u$. Then

$$\operatorname{in}_w(f) = \sum_{u \in \mathbb{N}^{n+1}} \overline{c_u t^{w \cdot u - W}} x^u$$

where $W = \operatorname{trop}(f)(w)$. Let $W' = \min(v \cdot u : \operatorname{val}(c_u) + w \cdot u = W)$. Then

$$\operatorname{in}_{v}(\operatorname{in}_{w}(f)) = \sum_{v \cdot u = W'} \overline{c_{u} t^{w \cdot u - W}} x^{u}.$$

For all sufficiently small $\epsilon > 0$, we have $W + \epsilon W' = \operatorname{trop}(f)(w + \epsilon v)$ and $\{u : \operatorname{val}(c_u) + (w + \epsilon' v) \cdot u = W + \epsilon W'\} = \{u : \operatorname{val}(c_u) + w \cdot u = W, v \cdot u = W'\}.$ This implies $\operatorname{in}_{w + \epsilon' v}(f) = \operatorname{in}_v(\operatorname{in}_w(f))$ for $\epsilon' \in \Gamma_{\operatorname{val}}$ with $0 < \epsilon' < \epsilon$. \Box

By a *term* in $S = K[x_0, \ldots, x_n]$ we mean the product $c_u x^u$ of a scalar $c_u \in K$ times a monomial. Given $w \in (\Gamma_{\text{val}})^{n+1}$, we define a ordering \prec_w on the set of all terms in S by setting

$$c_u x^u \prec_w c_{u'} x^{u'}$$

if $\operatorname{val}(c_u) + w \cdot u < \operatorname{val}(c_{u'}) + w \cdot u'$, or if $\operatorname{val}(c_u) + w \cdot u = \operatorname{val}(c_{u'}) + w \cdot u'$ and either $w \cdot u < w \cdot u'$ or $w \cdot u = w \cdot u'$ and $x^u \prec_{lex} x^{u'}$. Here \prec_{lex} is the usual lexicographic order. To emphasize that we are speaking of terms here, rather than monomials, for a polynomial $f = \sum c_u x^u \in S$ we use the notation $\operatorname{ST}(f)$ to denote the *smallest* term $c_u x^u \in S$ of f with respect to \prec_w . For a polynomial $\sum_u c_u x^u$ the *support* of f is the set of x^u with $c_u \neq 0$. The following proposition is a variant of the division algorithm, adapted to the notion of term orderings employed here to work over a field with valuation.

2.4. GRÖBNER BASES

Proposition 2.4.5. Fix a homogeneous ideal I in $S = K[x_0, x_1, \ldots, x_n]$ and $\nu \in \mathbb{N}$. For each $w \in (\Gamma_{val})^{n+1}$ lying outside a finite union of hyperplanes in \mathbb{R}^n , the monomials of degree ν outside of $\operatorname{in}_w(I)_{\nu}$ form a K-basis for $(S/I)_{\nu}$. More precisely, each polynomial $f \in S_{\nu}$ has a unique representation f = j + r where $j \in I$, $\operatorname{ST}(r) \succeq_w \operatorname{ST}(f)$, and r is a sum of terms not in $\operatorname{in}_w(I)$.

Proof. The vector space I_{ν} is a subspace of the $N = \binom{n+\nu}{\nu}$ -dimensional vector space S_{ν} over K, which has as basis the monomials x^{u_1}, \ldots, x^{u_N} in S of degree ν . Let g_1, \ldots, g_s be a basis of I_{ν} . Write $g_i = \sum_{j=1}^N c_{ij} x^{u_j}$, and form the $s \times N$ matrix C with ijth entry equal to c_{ij} . Note that C has rank s, since the g_i are a basis for I_d . Given a set $J = \{x^{u_{i_1}}, \ldots, x^{u_{i_s}}\}$ of s monomials of degree ν in S, we set $W_J = \sum_{i=1}^s w \cdot u_{i_i}$, and we denote by d_J the minor of C indexed by J. Since C has rank s, at least one d_J is nonzero, so there is J with $\operatorname{val}(d_J) < \infty$. Our genericity notion for w is that the minimum over all J of val $(d_J) + W_J$ is achieved once. This can be guaranteed by choosing w outside of the finite number of hyperplanes $w \cdot \mathbf{e}_J + \operatorname{val}(d_J) = w \cdot \mathbf{e}_{J'} + \operatorname{val}(d_{J'})$, where J, J' are collections of s monomials of degree d and $\mathbf{e}_J = \sum_{x^u \in J} u$. Let D be the $s \times s$ submatrix of C consisting of the columns indexed by that J minimizing val $(d_J) + W_J$. Since that minimum value val $(d_J) + W_J$ is not infinite, $d_J \neq 0$, so D is invertible. Set $C' = D^{-1}C$. Then the first s columns of C' consist of the $s \times s$ identity matrix, and every other nonzero entry c'_{ij} has $\operatorname{val}(c'_{ij}) + w \cdot u_j - w \cdot u_i > 0$. This latter observation follows from the facts that the J'th minor of C' is $d_{J'}/d_J$ for any collection J', and $\operatorname{val}(d_J) + W_J < \operatorname{val}(d_{J'}) + W_{J'}$, applied to the collection $J' = J \setminus \{x^{u_i}\} \cup \{x^{u_j}\}$.

The polynomials g'_1, \ldots, g'_s corresponding to the rows of C' are also a K-basis for I_{ν} , and $\operatorname{in}_w(g'_i) = x^{u_i} \in \mathbb{k}[x_0, \ldots, x_n]$. If $g \in I_{\nu}$ is nonzero, then $g = \sum a_i g'_i$ where $a_i \in K$, and $\operatorname{in}_w(g)$ is a sum of those terms $a_i x^{u_i}$ with $\operatorname{val}(a_i) + w \cdot u_i$ minimized. Thus x^{u_1}, \ldots, x^{u_s} are a k-basis for $\operatorname{in}_w(I)$. For $f = \sum_{i=1}^N d_i x^{u_i}$, set $j = \sum_{i=1}^s d_i g'_i$. Then $r = f - j = \sum_{i=r+1}^N (d_i - \sum_{j=1}^s d_j c_{jj}) x^{u_i}$ is supported on the monomials not in $\operatorname{in}_w(I)$. Let $\operatorname{ST}(f) = d_l x^{u_l}$, so $\operatorname{val}(d_i) + w \cdot u_i$ is minimized at i = l. Now $\operatorname{val}(d_i - \sum_{j=1}^s d_j c_{ji}) + w \cdot u_i \ge \operatorname{val}(d_i) + w \cdot u_l$, so $\operatorname{ST}(r) \succeq_w \operatorname{ST}(f)$.

To complete the proof, we observe that the monomials not in $\operatorname{in}_w(I)_{\nu}$ are linearly independent in S/I, as if $g = \sum c_i x^{u_i} \in I$, then $\operatorname{in}_w(g) \in \operatorname{in}_w(I)$, so $\operatorname{in}_w(g)$ would not lie in the span of the monomials not in $\operatorname{in}_w(I)_{\nu}$. Conversely, given $f \in S_{\nu}$, the polynomial r constructed above lies in the span of the monomials not in $\operatorname{in}_w(I)_{\nu}$, and $f - r \in I$. Together this show that the monomials not in $\operatorname{in}_w(I)_{\nu}$ form a K-basis for $(S/I)_{\nu}$. In what follows we use the notations $S_K = K[x_0, \ldots, x_n]$ and $S_{\Bbbk} = \&[x_0, \ldots, x_n]$ for the two polynomial rings that contain a given homogeneous ideal I and its various initial ideals $\operatorname{in}_w(I)$. Following [CLO07, §9] we measure the size of these ideals by their *Hilbert functions*. These are numerical functions $\mathbb{N} \to \mathbb{N}$. For large enough arguments, the Hilbert function is a polynomial (called the *Hilbert polynomial*) whose degree is one less than the Krull dimension of the quotient of the polynomial ring modulo that ideal.

Corollary 2.4.6. For any $w \in \Gamma_{\text{val}}^{n+1}$ and any homogeneous ideal I in S_K , the Hilbert functions of I agrees with that of its initial ideal $\operatorname{in}_w(I) \subset S_k$, i.e.

$$\dim_K (S_K/I)_{\nu} = \dim_k (S_k/\operatorname{in}_w(I))_{\nu} \quad \text{for all } \nu \ge 0$$

Thus the Krull dimensions of the residue rings S_K/I and $S_k/\operatorname{in}_w(I)$ coincide.

Proof. Fix $\nu \geq 0$, and choose $\nu \in (\Gamma_{\text{val}})^{n+1}$ generic in the sense of Proposition 2.4.5 for some basis g_1, \ldots, g_s of $\operatorname{in}_w(I)_{\nu}$. We may assume by Proposition 2.4.5 that the $\operatorname{in}_v(g_i) = m_i$ are all distinct monomials, and $\operatorname{in}_v(g_i)$ is not a term in g_j for $i \neq j$. By Lemma 2.4.2 we may assume that $g_i = \operatorname{in}_w(f_i)$ for some $f_i \in I_{\nu}$. Choose a field extension K' of k with a nontrivial valuation with residue field k. Then $\operatorname{in}_v(\operatorname{in}_w(I))_{\nu}$ has a basis of monomials, and by Proposition 2.4.5 the set of monomials \mathcal{B} of degree ν not in $\operatorname{in}_v(\operatorname{in}_w(I))$ form a K'-basis, and thus a k-basis, for $(S_k/\operatorname{in}_w(I))_{\nu}$.

By Lemma 2.4.4 there is $\epsilon > 0$ for which $\operatorname{in}_{v}(\operatorname{in}_{w}(I)) = \operatorname{in}_{w+\epsilon v}(I)$, and $\operatorname{in}_{w+\epsilon v}(f_{i}) = m_{i}$, so the set of m_{i} is a basis for $\operatorname{in}_{w+\epsilon v}(I)_{\nu}$. Let $f = \sum_{i=1}^{N} c_{u_{i}} x^{u_{i}} \in S_{\nu}$, where $\{x^{u_{i}} : 1 \leq i \leq N\}$ are the monomials of S_{ν} , and $\mathcal{B} = \{x^{u_{r+1}}, \ldots, x^{u_{N}}\}$. Then $g = f - \sum_{i=1}^{r} c_{u_{i}} f_{i} \in I$ is supported on \mathcal{B} , so $\operatorname{in}_{w+\epsilon v}(g) \in \operatorname{in}_{w+\epsilon v}(I) = \operatorname{in}_{v}(\operatorname{in}_{w}(I))$, which would contradict that the monomials of degree ν not in \mathcal{B} span $\operatorname{in}_{v}(\operatorname{in}_{w}(I))_{\nu}$ unless g = 0. From this we conclude that f_{1}, \ldots, f_{s} span I_{ν} . They are linearly independent since the monomial m_{i} does not appear in the support of f_{i} for $i \neq j$.

Consider the $r \times N$ matrix C for which c_{ij} is the coefficient of x^{u_j} in f_i , and let d_J be the determinant of the submatrix whose columns are indexed by $J = \{x^{u_{i_1}}, \ldots, x^{u_{i_s}}\}$. Then the fact that m_1, \ldots, m_r span $\inf_{w+\epsilon v}(I)$ implies that $\operatorname{val}(d_J) + \sum_{x^u \in J} w \cdot u$ is minimized at $J = \{x^{u_1}, \ldots, x^{u_r}\}$, so w is generic for w in the sense of Proposition 2.4.5, so \mathcal{B} forms a K-basis for $(S/I)_{\nu}$.

The last sentence then follows from the fact that Hilbert function of a homogeneous ideal determines the Krull dimension of its residue ring. \Box

80

Example 2.4.7. The Hilbert function of the ideals in Example 2.4.3 equals

$$\dim_{\mathbb{Q}}(\mathbb{Q}[x_0, x_1, x_2, x_3]/I)_{\nu} = \dim_{\mathbb{K}}(\mathbb{K}[x_0, x_1, x_2, x_3]/\operatorname{in}_w(I))_{\nu} = \nu + 1.$$

Here $\mathbb{k} = \mathbb{Z}/2\mathbb{Z}$ is the field with two elements. The Hilbert polynomial $\nu + 1$ shows that the projective varieties have dimension 1 and degree 1. They are straight lines in $\mathbb{P}^3_{\mathbb{Q}}$ and $\mathbb{P}^3_{\mathbb{k}}$. Note that $\mathbb{P}^3_{\mathbb{k}}$ is a finite set with 15 elements. \Box

Corollary 2.4.6 implies that the projective varieties $V(I) \subset \mathbb{P}_K$ and $V(\mathrm{in}_w(I)) \subset \mathbb{P}_k$ always have the same dimension. In typical applications, V(I) is an irreducible variety but $V(\mathrm{in}_w(I))$ can have many irreducible components. Our next result states that every irreducible component of $V(\mathrm{in}_w(I))$ has the same dimension as V(I). We shall phrase this in the algebraic language of primary decomposition. Recall that P is a minimal associated prime of an ideal $I \subset S$ if $I \subseteq P$, and there is no prime ideal P' with $I \subseteq P' \subsetneq P$.

Lemma 2.4.8. If $I \subset S_K$ is a homogeneous prime of dimension d, and $w \in (\Gamma_{val})^{n+1}$, then every minimal associated prime of $in_w(I)$ has dimension d.

Proof. Let $\mathcal{G} = \{g_1, \ldots, g_s\}$ be a Gröbner basis for I, and let $g_i^t \in S_K$ be the polynomial $t^{-\operatorname{trop}(g_i)(w)}g_i(t^{w_0}x_0, \ldots, t^{w_n}x_n)$. Let R' be the local subring of $R \subset K$ containing all the coefficients of all g_i^t and the element $t = t^1$ defined as follows. The ring R' is constructed by taking the algebra over the prime field \mathbb{F} of K generated by these coefficients and t, and localizing at a prime \mathfrak{m}' minimal over t. Note that R' is a Noetherian ring by construction. By the Principal Ideal Theorem [Eis95, Theorem 10.1], the local ring R' has Krull dimension one, and so the maximal ideal $\mathfrak{m}' = \mathfrak{m} \cap R'$. The fraction field K' of R' is a subfield of K, and $\Bbbk' = R'/\mathfrak{m}'$ is a subfield of \Bbbk .

Let $I' = I \cap R'[x_0, \ldots, x_n]$, and let $I'' = I \cap K'[x_0, \ldots, x_n]$. Since $I = I'' \otimes_{K'} K$, we have $\dim(K[x_0, \ldots, x_n]/I) = \dim(K'[x_0, \ldots, x_n]/I'') = d$. In addition, $\dim(R'[x_0, \ldots, x_n]/I') = d + 1$. This follows, for example, from [Eis95, Theorem 13.8] applied to the prime $Q = \langle x_0, \ldots, x_n \rangle + \mathfrak{m}'$ of $R'[x_0, \ldots, x_n]/I'$, since R' is one-dimensional and universally catenary by [Eis95, Corollary 18.10]. Thus the codimension of the prime ideal I' is n + 1 - d.

Let P be a prime ideal of $R'[x_0, \ldots, x_n]$ minimal over $I' + \mathfrak{m}'$. Note that any prime containing I' + t must intersect R' in a prime containing t, so must contain \mathfrak{m}' . Thus P is minimal over $I' + \langle t \rangle$. By the Principal Ideal Theorem [Eis95, Theorem 10.1] applied to the domain $R'[x_0, \ldots, x_n]/I'$ the codimension of P/I' is thus one, so the dimension of P is d+1-1 = d. Since minimal primes of $(I' + \mathfrak{m}')/\mathfrak{m}'$ are of the form P/\mathfrak{m}' for minimal primes of $I' + \mathfrak{m}'$, this shows that all minimal primes of $(I' + \mathfrak{m}')/\mathfrak{m}$ are d-dimensional.

It thus suffices to show that $(I' + \mathfrak{m}')/\mathfrak{m} \otimes_{\mathbb{k}'} \mathbb{k} = \operatorname{in}_w(I)$. The polynomials g_i^t lie in $R'[x_0, \ldots, x_n]$ by construction, so their images in $\mathbb{k}[x_0, \ldots, x_n]$ lie in $\mathbb{k}'[x_0, \ldots, x_n]$, which shows that $\operatorname{in}_w(I) \subseteq (I' + \mathfrak{m}')/\mathfrak{m} \otimes \mathbb{k}$. The other inclusion is automatic, so we conclude that every minimal prime of $\operatorname{in}_w(I)$ is d-dimensional.

Remark 2.4.9. The ring $R = \{x \in K : \operatorname{val}(x) \ge 0\}$ need not be Noetherian. For example, when $K = \mathbb{C}\{\{t\}\}\)$, the ideals $I_n = \{x \in R : \operatorname{val}(x) > 1/n\}\)$ form an increasing chain of ideals in R. This necessitated passing to the Noetherian subring R' of R in the proof of Lemma 2.4.8, as many of the fundamental theorems of dimension theory apply only to Noetherian rings.

We now come to the punchline of this section, which is the construction of a polyhedral complex from a given homogeneous ideal $I \subset K[x_0, \ldots, x_n]$. First we shall define the polyhedra in this complex. Given $w \in (\Gamma_{\text{val}})^{n+1}$, let

$$C_I[w] = \{ w' \in (\Gamma_{\text{val}})^{n+1} : \text{ in}_{w'}(I) = \text{ in}_w(I) \}.$$

Let $\overline{C_I[w]}$ be the closure of $C_I[w]$ in \mathbb{R}^{n+1} in the Euclidean topology. The all-one vector $\mathbf{1} = (1, \ldots, 1)$ lies in $(\Gamma_{\text{val}})^{n+1}$, and, since I is homogeneous, it satisfies $\operatorname{in}_w(I) = \operatorname{in}_{w+\lambda \mathbf{1}}(I)$ for all $w \in (\Gamma_{\text{val}})^{n+1}$ and $\lambda \in \Gamma_{\text{val}}$. This implies $C_I[w] = C_I[w + \lambda \mathbf{1}]$ for all w, λ . For that reason, we shall identify the polyhedron $\overline{C_I[w]}$ with its image in the *n*-dimensional space $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$.

Example 2.4.10. Let n = 2 and $K = \mathbb{Q}$ with the 2-adic valuation, and let I be the principal ideal generated by the homogeneous cubic polynomial

$$f = 2x_0^3 + 4x_1^3 + 2x_2^3 + x_0x_1x_2.$$

The initial ideal for w = (0, 0, 0) equals $in_w(I) = \langle x_0 x_1 x_2 \rangle$. The polyhedron $C_I[w]$ is a compact subset of the plane $\mathbb{R}^3/\mathbb{R}\mathbf{1}$. Namely, it equals the triangle

$$C_{I}[w] = \{(v_{0}, v_{1}, v_{2}) \in \mathbb{R}^{3} / \mathbb{R}\mathbf{1} : v_{0} + v_{1} + v_{2} \le \min(3v_{0} + 1, 3v_{1} + 2, 3v_{2} + 1)\}.$$

Note that the use of the valuation is essential here because $x_0x_1x_2$ would not be an initial monomial of f in the usual Gröbner basis sense of [CLO07]. \Box

2.4. GRÖBNER BASES

Theorem 2.4.11. The polyhedra $\overline{C_I[w]}$ as w varies over $(\Gamma_{val})^{n+1}$ form a (Γ_{val}) -rational polyhedral complex inside the n-dimensional space $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$.

Proof. Fix $\nu \geq 0$, and let $C_I^d[w] = \{w' \in (\Gamma_{val})^{n+1} : in_{w'}(I)_d = in_w(I)_d\}$. We first show that the sets $\overline{C_I^d[w]}$ as w varies over $(\Gamma_{val})^{n+1}$ form an (Γ_{val}) -rational polyhedral complex with lineality space span(1). Let $\dim_K I_d = D$, and fix $N = \binom{n+d}{d}$. We enumerate the monomials of degree d as $\{\mathbf{x}^{\mathbf{u}_1}, \ldots, \mathbf{x}^{\mathbf{u}_N}\}$. The K-vector-space I_d corresponds to a point in the Grassmannian $Gr_K(D, N)$, which we consider in its Plücker embedding into $\mathbb{P}_K^{\binom{N}{D}-1}$. By Corollary 2.4.6 we have $\dim_{\Bbbk} in_w(I)_d = D$ for all $w \in (\Gamma_{val})^{n+1}$, so $in_w(I)$ corresponds to a point in $CGr_{\Bbbk}(D, N)$, which we consider in its Plücker coordinates by sets in $[N]^D := \{J \subset \{1, \ldots, N\} : |J| = D\}$. We use the notation P_J for the Plücker coordinates in $\mathbb{P}_K^{\binom{N}{D}-1}$, and p_J for the coordinates in $\mathbb{P}_{\Bbbk}^{\binom{N}{D}-1}$. Let $U_J = \sum_{j \in J} \mathbf{u}_j$. Let $W = \min\{\operatorname{val}(P_J) + w \cdot U_J : J \in [N]^D\}$. The key observation is that (up to a global scaling factor for all $J \in [N]^D$) we have

$$p_J = \overline{t^{w \cdot U_j - W} P_J}.\tag{2.5}$$

To see this, choose a K-basis g_1, \ldots, g_D for I_d , and fix J with $P_J \neq 0$ and $\operatorname{val}(P_J) + w \cdot U_J = W$. Let G be $D \times N$ matrix with entries in K whose *i*th row records the coefficients of g_i . Multiplying the *j*th column by $t^{w \cdot \mathbf{u}_{J_j}}$ adds $w \cdot U_J$ to P_J . This corresponds to replacing g_i by $\tilde{g}_i(x_0, \ldots, x_n) = g_i(t^{w_0}x_0, \ldots, t^{w_n}x_n)$, which are a basis for the degree d part of the ideal $t^{\mathbf{w}} \cdot I := \langle g(t_0x_0, \ldots, t_nx_n) : g \in I \rangle$. By definition $\operatorname{in}_w(I)_d = \operatorname{in}_0(t^{\mathbf{w}} \cdot I)$, so we may assume that w = 0. Multiplying G on the left by an element $A \in \operatorname{GL}(D, K)$ does not change the property that the rows index a basis for I_d of the desired form, though it adds $\operatorname{val}(\det(A))$ to W. We may thus multiply by the inverse of the $D \times D$ submatrix G_J of G indexed by the columns in J to obtain a choice of g_i with G_J the identity matrix, and thus $P_J = 1$, and W = 0. Note that this also means that the valuation of every entry of G is now nonnegative.

Let \overline{G} be the $D \times N$ matrix with entries in \Bbbk whose (i, j)th entry is \overline{G}_{ij} . The *i*th row of \overline{G} records the coefficients of $\operatorname{in}_0(g_i)$. The submatrix indexed by J is the identity matrix, so the polynomials $\operatorname{in}_0(g_1), \ldots, \operatorname{in}_0(g_D)$ are linearly independent, and thus form a basis for $\operatorname{in}_0(I)$. Thus for any $J' \in [N]^D$ the determinant of the $D \times D$ submatrix of \overline{G} indexed by J' is $p_{J'}$, which thus equals $\overline{P}_{J'}$. This proves Equation 2.5. Thus we have $w' \in C_I^d[w]$ if and only if $\operatorname{val}(P_J) + w' \cdot' U_J = W' := \min\{\operatorname{val}(P_J) + w' \cdot U_J : J \in [N]^D\}$ for all $J \in [N]^D$ with $\operatorname{val}(P_J) + w \cdot U_J = W$, and $\operatorname{val}(P_J) + w' \cdot U_J > W'$ for all $J \in [N]^D$ with $\operatorname{val}(P_J) + w \cdot U_J > W$. This means that $C_I^d[w]$ is the set of relative interior of an $(\Gamma_{\operatorname{val}})$ -rational polyhedron $\overline{C_I^d[w]}$. Since the relative interiors of distinct $\overline{C_I^d[w]}$ do not intersect by the definition of $C_I^d[w]$, the collection of all $\overline{C_I^d[w]}$ form a polyhedral complex.

We next observe that there is a finite set \mathcal{D} of degrees for which if $\operatorname{in}_w(I)_d = \operatorname{in}_{w'}(I)_d$ for all $d \in \mathcal{D}$ then $\operatorname{in}_w(I) = \operatorname{in}_{w'}(I)$. By [Mac01, Corollary 2.2] there are only a finite number of monomial ideals with the same Hilbert function as I. It suffices to take \mathcal{D} to contain the degrees of all minimal generators of these monomial ideals, as these degrees contains a Gröbner basis, and thus a generating set, for $\operatorname{in}_w(I)$ and $\operatorname{in}_{w'}(I)$. Note that $C_I[w] = \bigcap_{d \in \mathcal{D}} C_I^d[w]$, so taking the common refinement of the polyhedral complexes $\{\overline{C_I[w]}: w \in (\Gamma_{\operatorname{val}})^{n+1}\}$ for all $d \in \mathcal{D}$ gives the polyhedral complex $\{\overline{C_I[w]}: w \in (\Gamma_{\operatorname{val}})^{n+1}\}$. Since the intersection of $(\Gamma_{\operatorname{val}})$ -rational polyhedra is $\Gamma_{\operatorname{val}}$ -rational, this complex is again $(\Gamma_{\operatorname{val}})$ -rational.

Definition 2.4.12. The Gröbner complex $\Sigma(I)$ of a homogeneous ideal I in $K[x_0, x_1, \ldots, x_n]$ is the polyhedral complex constructed in Theorem 2.4.11. It consists of the polyhedra $\overline{C_I[w]}$ as w ranges over $(\Gamma_{\text{val}})^{n+1}$.

The support of the Gröbner complex $\Sigma(I)$ is the *n*-dimensional space $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$. From the perspective of the tropical semiring $(\mathbb{R}, \oplus, \odot)$, this space can be regarded as the tropical projective space, as it is obtained from \mathbb{R}^{n+1} by identifying vectors that differ from each other by tropical scalar multiplication. For that reason, the $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$ was denoted \mathbb{TP}^n in some early papers on tropical geometry. In this book, we retain the notation $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$ because we wish to reserve the name tropical projective space notation \mathbb{TP}^n for the natural compactification obtained by including ∞ in the tropical semiring. Points in $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$ can be uniquely represented by vectors of the form $(0, v_1, \ldots, v_n)$, and this also is the convention used for drawing pictures.

Example 2.4.13. Let $f = tx_1^2 + 2x_1x_2 + 3tx_2^2 + 4x_0x_1 + 5x_0x_2 + 6tx_0^2 \in \mathbb{C}\{\{t\}\}[x_0, x_1, x_2]$, and let $I = \langle f \rangle$ be the ideal generated by f. The Gröbner complex of I is a polyhedral complex in the plane $\mathbb{R}^3/\mathbb{R}\mathbf{1}$ shown in Figure 2.7.

The ideal I has 19 distinct initial ideals, corresponding to the various cells of $\Sigma(I)$. There are 6 cells of dimension two, 9 cells of dimension one, and 4 cells of dimension one. The following table lists eight of the 19 initial ideals, namely, those corresponding to the labels in the diagram.



Figure 2.7: The Gröbner complex of a plane curve subdivides $\mathbb{R}^3/\mathbb{R}\mathbf{1}$.

Cell	Initial ideal		ell Initial Ideal
А	$\langle 4x_0x_1\rangle$	E	$\langle 5x_0x_2\rangle$
В	$\langle 4x_0x_1+6x_0^2\rangle$	F	$\langle 3x_2^2 \rangle$
С	$\langle 6x_0^2 \rangle$	G	$\langle 2x_1x_2\rangle$
D	$\langle 4x_0x_1 + 5x_0x_2 + 6x_0^2 \rangle$	H	$\langle x_1^2 \rangle$

In this example we observe that the initial ideal $\operatorname{in}_w(I)$ contains a monomial if and only if the corresponding cell is full-dimensional in the plane $\mathbb{R}^3/\mathbb{R}\mathbf{1}$. \Box

The construction of the Gröbner complex allows us to define the concept of a universal Gröbner basis. By this we mean a finite subset \mathcal{U} of I such that, for all $w \in (\Gamma_{\text{val}})^{n+1}$, the set $\operatorname{in}_w(\mathcal{U}) = \{\operatorname{in}_w(f) : f \in \mathcal{U}\}$ generates the initial ideal $\operatorname{in}_w(I)$. In particular, \mathcal{U} generates I, as seen by taking w = 0.

Corollary 2.4.14. Every homogeneous ideal $I \subset K[x_0, \ldots, x_n]$ has finite universal Gröbner basis.

Proof. The Gröbner complex $\Sigma(I)$ is finite. Pick a representative w for each cell, and compute a Gröbner basis of I with respect to each of these finitely many w. The union of these Gröbner bases is a universal Gröbner basis. \Box

An important special case arises when our homogeneous ideal I is generated by polynomials f whose coefficients all have valuation zero. If this happens then we say that the polynomial f has *constant coefficients* and same for the ideal I. The constant coefficient case is ubiquitous in this book, and it furnishes the bridge to the more familiar notion of term orders used for Gröbner bases, as in [CLO07, Eis95, Stu96]. Indeed, if f is a homogeneous polynomial with constant coefficients and $w \in (\Gamma_{val})^{n+1}$ is sufficiently generic, then $in_w(f)$ is the leading monomial of f with respect to the term order determined by -w. See e.g. [Eis95, §15.1]. This identifies the Gröbner complex with the Gröbner fan of [Stu96, Chapter 2], up to a sign change.

Corollary 2.4.15. Let I be a homogeneous ideal with constant coefficients. Then the negated Gröbner complex $-\Sigma(I)$ is equal to the Gröbner fan of I.

In many of the geometric examples seen later in this book we will examine a projective variety whose defining ideal I has coefficients in the field \mathbb{Q} of rational numbers. Such an ideal I has well-defined Gröbner fan, and it arises as $-\Sigma(I)$ from the inclusion of \mathbb{Q} into any field with non-trivial valuation, such as the Puiseux series $\mathbb{C}\{\{t\}\}\$. On the other hand, we can also consider the *p*-adic Gröbner complex of the same ideal I. The *p*-adic Gröbner complex $\Sigma(I)$ is generally not a fan, as it arises from the *p*-adic valuation on \mathbb{Q} .

$\mathbf{2.5}$ **Tropical Bases**

In the previous section we introduced Gröbner bases and the Gröbner complex for homogeneous ideals in a polynomial ring $K[x_0, x_1, \ldots, x_n]$ over a field K with valuation. We now examine the case when the ambient ring is the Laurent polynomial ring $K[x^{\pm}] = K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. We wish to argue that there is no natural intrisic notion of Gröbner bases for ideals in $K[x^{\pm}]$. However, there is a natural analogue to the notion of a universal Gröbner basis, namely, that of a tropical basis, and this is our subject in this section.

For every polynomial $f \in K[x^{\pm}]$ and $w \in \Gamma_{val}^{n}$ we define the initial form $in_w(f) \in k[x^{\pm}]$ by the same rule as in the previous section. Namely, we set

$$\operatorname{in}_w(f) = \sum_{\substack{u:\operatorname{val}(c_u)+w\cdot u \\ =W}} \overline{t^{-\operatorname{val} c_u} c_u} x^u$$

where $W = \min\{\operatorname{val}(c_u) + w \cdot u : c_u \neq 0\}$. Let I be any ideal in $K[x^{\pm}] = K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. The initial ideal $\operatorname{in}_w(I)$ is the ideal in $\mathbb{k}[x^{\pm}]$ generated by the initial forms $\operatorname{in}_w(f)$ as f runs over I. So far, this is the same as in the polynomial ring, but there is an important distinction that arises when we work with Laurent polynomials. For generic choices of $w = (w_1, \ldots, w_n)$, the initial form $\operatorname{in}_w(f)$ is a unit in $\Bbbk[x^{\pm}]$, and the initial ideal $\operatorname{in}_w(I)$ will be equal to the whole Laurent polynomial ring $\Bbbk[x^{\pm}]$. If this happens then the initial ideal contains no information at all. Tropical geometry is concerned with the study of those special weight vectors $w \in \Gamma_{\operatorname{val}}^n$ for which the initial ideal $\operatorname{in}_w(I)$ is actually proper ideal in $\Bbbk[x^{\pm}]$.

In order to compute and study these initial ideals, it is useful to work with homogeneous polynomials as in Section 2.4. We write I_{proj} for its homogenization in $K[x_0, x_1, \ldots, x_n]$. This is the ideal generated by all polynomials

$$\tilde{f} = x_0^m f\left(\frac{x_1}{x_0}, \dots, \frac{x_n}{x_0}\right),$$

where $f \in I$ and m is the smallest integer that clears the denominator.

The initial ideals $\operatorname{in}_w(I)$ of an ideal $I \subseteq K[x^{\pm}]$ can be computed from the initial ideals of its homogenization I_{proj} as follows. The weight vectors for the homogeneous ideal I_{proj} live in naturally in the quotient space $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$, and we identify this space with \mathbb{R}^n via $w \mapsto (0, w)$.

Proposition 2.5.1. Let I be an ideal in $K[x^{\pm}]$, and fix $w \in (\Gamma_{val})^n$. Then $\operatorname{in}_w(I)$ equals the image of $\operatorname{in}_{(0,w)}(I_{proj})$ in $\Bbbk[x^{\pm}]$ obtained by setting $x_0 = 1$.

Proof. We first observe that for all $f \in I \cap K[x_1, \ldots, x_n]$ we have $\operatorname{in}_{(0,w)}(\tilde{f}) = inn_w(f)$ as a Laurent polynomial. Indeed, let $f = \sum c_u x^u \in K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and $\tilde{f} = \sum c_u x^u x_0^{j_u}$, where $j_u = \max_{c_u \neq 0} |u| - |u|$. Then $\operatorname{trop}(f)(w) = \min(\operatorname{val}(c_u) + w \cdot u) = \min(\operatorname{val}(c_u) + (0, w) \cdot (j_u, u)) = \operatorname{trop}(\tilde{f})((0, w))$. Thus $\operatorname{in}_{(0,w)}(\tilde{f}) = \sum_{\operatorname{val}(c_u) + w \cdot u = W} \overline{c_u t^{-\operatorname{val}(c_u)}} x^u x_0^{j_u}$, so $\operatorname{in}_{(0,w)}(\tilde{f})|_{x_0=1} = \sum_{\operatorname{val}(c_u) + w \cdot u = W} c_u t^{-\operatorname{val}(c_u)} x^u = \operatorname{in}_w(f)$.

Next note that by multiplying by monomials if necessary, we may choose a Gröbner basis $\{f_1, \ldots, f_s\}$ for I consisting of polynomials in $K[x_1, \ldots, x_n]$. The above calculation applied to the f_i shows that $\operatorname{in}_w(I) \subseteq \operatorname{in}_{(0,w)}(I_{\operatorname{proj}})|_{x_0=1}$. For the reverse inclusion, note that if g is a homogeneous polynomial in I_{proj} , then $g = x_0^j \tilde{h}$ for some j, where h(x) = g(1, x), so since by Lemma 2.4.2 we can choose a homogeneous Gröbner basis for I_{proj} this case also follows from the above calculation.

Here are some facts about initial ideals of Laurent polynomial ideals.

Lemma 2.5.2. Let I be an ideal in $K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and fix $w \in (\Gamma_{val})^n$. 1. If $in_w(I) = \langle 1 \rangle$, then there is $f \in I$ with $in_w(f) = 1$.

- 2. If $\operatorname{in}_{\mathbf{v}}(\operatorname{in}_{w}(I)) = \operatorname{in}_{w}(I)$ for $\mathbf{v} \in \mathbb{Z}^{n}$, then $\operatorname{in}_{w}(I)$ is homogeneous with respect to the grading given by $\operatorname{deg}(x_{i}) = v_{i}$.
- 3. If $g \in in_w(I)$, then $g = in_w(h)$ for some $h \in I$.
- 4. If $f, g \in K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, then $in_w(fg) = in_w(f) in_w(g)$.

Proof. If $\operatorname{in}_w(I) = \langle 1 \rangle$, then by Proposition 2.5.1, there is a monomial x^u in $\operatorname{in}_{(0,w)}(I_{\operatorname{proj}})$. By Lemma 2.4.2 there is $f \in I_{\operatorname{proj}}$ with $\operatorname{in}_{(0,w)}(f) = x^u$. Let $\overline{u} = (u_1, \ldots, u_n)$ and $g = x^{-\overline{u}} f(1, x_1, \ldots, x_n)$. Then $g \in I$, and $\operatorname{in}_w(g) = 1$.

Suppose now that $\operatorname{in}_{\mathbf{v}}(\operatorname{in}_{w}(I)) = \operatorname{in}_{w}(I)$. This means that $\operatorname{in}_{w}(I)$ has generating set of the form $\operatorname{in}_{\mathbf{v}}(g)$ for $g \in \operatorname{in}_{w}(I)$. For any $g = \sum a_{u}x^{u} \in \mathbb{k}[x_{1}^{\pm 1}, \ldots, x_{n}^{\pm 1}]$, the initial form $\operatorname{in}_{\mathbf{v}}(g) = \sum_{\mathbf{v} \cdot u = W} a_{u}x^{u}$ for $W = \min_{a_{u} \neq 0} \mathbf{v} \cdot u$. This is homogeneous in the **v**-grading of degree W, so $\operatorname{in}_{w}(I)$ has a **v**-homogeneous generating set.

Part 3 follows directly from Lemma 2.4.2 and Proposition 2.5.1.

For the last part, write $f = \sum c_u x^u$ and $g = \sum d_u x^u$. Then $fg = \sum_w e_v x^v$ for $e_v = \sum_{u+u'=v} c_u d_{u'}$. Let $W_1 = \operatorname{trop}(f)(w)$, and let $W_2 = \operatorname{trop}(g)(w)$. Since $\operatorname{trop}(fg) = \operatorname{trop}(f) + \operatorname{trop}(g)$, $\operatorname{trop}(fg)(w) = W_1 + W_2$, so $\operatorname{in}_w(fg) = \sum_{\operatorname{val}(e_v)+w\cdot v=W_1+W_2} \overline{e_v t^{-\operatorname{val}(e_v)}} x^v$. Now $\operatorname{val}(e_v) + w \cdot v \ge \operatorname{val}(c_u) + w \cdot u + \operatorname{val}(d_{u'}) + w \cdot u'$ for all u, u' with u + u' = v, so $\operatorname{val}(e_v) + w \cdot v = W_1 + W_2$ only if $\operatorname{val}(c_u) + w \cdot u = W_1$ and $\operatorname{val}(d_{u'}) + w \cdot u' = W_2$. So $\operatorname{in}_w(fg) = \operatorname{in}_w(f) \operatorname{in}_w(g)$ as required.

Definition 2.5.3. Let I be an ideal in the Laurent polynomial ring $K[x^{\pm}]$ over a field K with a valuation. A finite generating set \mathcal{T} of I is said to be a *tropical basis* if, for all weight vectors $w \in \Gamma_{\text{val}}^n$, the initial ideal $\text{in}_w(I)$ contains a unit if and only if $\text{in}_w(\mathcal{T}) = \{\text{in}_w(f) : f \in \mathcal{T}\}$ contains a unit.

Theorem 2.5.4. Every ideal I in $K[x^{\pm}]$ has a finite tropical basis.

Proof. Consider the homogenization I_{proj} of I. Its Gröbner complex $\Sigma(I_{\text{proj}})$ is a polyhedral complex in $\mathbb{R}^{n+1}/\mathbb{R}\mathbf{1}$. For each of the finitely many cells $\Gamma^{(i)}$ in that complex, we select one representive vector $(0, w^{(i)}) \in \Gamma_{\text{var}}^{n+1}$. For each index i such that $\inf_{w^{(i)}}(I) = \langle 1 \rangle$, we select a Laurent polynomial $f^{(i)} \in I$ such that $\inf_{w^{(i)}}(f^{(i)}) = 1$. Such a choice is possible by Proposition 2.5.1 and part 1 of Lemma 2.5.2. Now we define \mathcal{T} by taking any finite generating set of I and augmenting it by the Laurent polynomials $f^{(i)}$ constructed above. Then \mathcal{T} is a generating set of I. Consider an arbitrary weight vector $w \in \Gamma_{\text{val}}^n$. There exists an index i such that $\inf_{(0,w)}(I_{\text{proj}}) = \inf_{(0,w^{(i)})}(I_{\text{proj}})$, and this ideal

88

2.5. TROPICAL BASES

contains a monomial if and only if $\operatorname{in}_{w^{(i)}}(f^{(i)})$ is a unit. Hence the initial ideal $\operatorname{in}_w(I)$ contains a unit if and only if the finite set $\operatorname{in}_w(\mathcal{T})$ contains a unit. \Box

Our first example of a tropical basis concerns principal ideals.

Example 2.5.5. If $f \in K[x^{\pm}]$ then $\{f\}$ is a tropical basis for the ideal $I = \langle f \rangle$ it generates. Indeed, suppose that $\operatorname{in}_w(I)$ contains a unit. Then there exists $g \in K[x^{\pm}]$ such that $\operatorname{in}_w(fg) = \operatorname{in}_w(f) \cdot \operatorname{in}_w(g)$ is a unit, and this implies that $\operatorname{in}_w(f)$ is a unit.

The concept of a tropical basis extends naturally to ideals in a polynomial ring. For instance, if J is a homogeneous ideal in $K[x_0, \ldots, x_n]$ then a generating set \mathcal{T} of J is a *tropical basis* of J if, for all $w \in \Gamma_{\text{val}}^n$, the ideal $\operatorname{in}_w(J)$ contains a monomial if and only if $\operatorname{in}_w(\mathcal{T})$ contains a monomial. In this setting, being a Gröbner basis and being a tropical basis are unrelated.

Example 2.5.6. Let I be the homogeneous ideal in $\mathbb{Q}[x, y, z]$ generated by $\mathcal{G} = \{x + y + z, x^2y + xy^2, x^2z + xz^2, y^2z + yz^2\}$. The set \mathcal{G} is a *universal Gröbner basis*, that is, \mathcal{G} is a Gröbner basis of I for all $w \in \Gamma^3_{\text{val}}$. However, \mathcal{G} is not a tropical basis. To see this, we take w = (0, 0, 0). The ideal $\text{in}_w(I) = I$ contains the monomial xyz but $\text{in}_w(\mathcal{G}) = \mathcal{G}$ contains no monomial. \Box

Our next goal is to show that the notion of a tropical basis is invariant under multiplicative coordinate changes in $K[x^{\pm}]$. Along the way, we shall prove a more general lemma that will be used in the proofs of Chapter 3.

Given a morphism $\phi : \mathbb{T}^n \to \mathbb{T}^m$, with associated ring homomorphism $\phi^* : K[x_1^{\pm 1}, \ldots, x_m^{\pm 1}] \to K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, we also denote by ϕ^* the map $\mathbb{Z}^m \to \mathbb{Z}^n$ given by setting $\phi^*(\mathbf{e}_i) = \mathbf{u}$ where $\phi^*(x_i) = x^{\mathbf{u}}$. This gives an induced map, which we also denote by ϕ , by applying $\operatorname{Hom}(-,\mathbb{Z})$ to ϕ^* :

$$\phi \colon \operatorname{Hom}(\mathbb{Z}^n, \mathbb{Z}) \cong \mathbb{Z}^n \to \operatorname{Hom}(\mathbb{Z}^m, \mathbb{Z}) \cong \mathbb{Z}^m.$$

If the homomorphism ϕ^* is given by $\phi^*(x_i) = x^{\mathbf{a}_i}$ for $\mathbf{a}_i \in \mathbb{Z}^n$, let A be the $n \times m$ matrix with *i*th column \mathbf{a}_i . Then the map $\phi \colon \mathbb{Z}^n \to \mathbb{Z}^m$ is given by A^T . We also denote by ϕ the induced map $\phi \colon \mathbb{Z}^n \otimes \mathbb{R} \cong \mathbb{R}^n \to \mathbb{Z}^m \otimes \mathbb{R} \cong \mathbb{R}^m$ induced by tensoring with \mathbb{R} . Note that the restriction of ϕ to Γ^n_{val} has image contained in Γ^m_{val} . Note also that for $y = (y_1, \ldots, y_n) \in \mathbb{T}^n$ we have

$$\operatorname{val}(\phi(y)) = (\operatorname{val}(y^{\mathbf{a}_1}), \dots, \operatorname{val}(y^{\mathbf{a}_m}))$$

= $(\mathbf{a}_1 \cdot \operatorname{val}(y), \dots, \mathbf{a}_m \cdot \operatorname{val}(y))$
= $A^T \operatorname{val}(y) = \phi(\operatorname{val}(y)).$ (2.6)

Lemma 2.5.7. Let $\phi^* \colon K[x_1^{\pm 1}, \dots, x_m^{\pm 1}] \to K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be a monomial map. Let $I \subseteq K[x_1^{\pm 1}, \dots, x_m^{\pm 1}]$ be an ideal, and let $I' = \phi^{*-1}(I)$. Then

$$\phi^*(\operatorname{in}_{\phi(w)}(I')) \subseteq \operatorname{in}_w(I) \quad \text{for all } w \in (\Gamma_{\operatorname{val}})^n.$$

Thus, in particular, if $\operatorname{in}_w(I) \neq \langle 1 \rangle$ then we also have $\operatorname{in}_{\phi(w)}(I') \neq \langle 1 \rangle$.

Proof. Let the monomial map ϕ^* be given by $\phi(x_i) = x^{\mathbf{a}_i}$, where $\mathbf{a}_i \in \mathbb{Z}^n$. Then $\phi(x^u) = x^{Au}$, where A is the $n \times m$ matrix with *i*th column \mathbf{a}_i . Let $f = \sum c_u x^u \in I'$, so $\phi(f) = \sum c_u x^{Au} \in I$. Let $W = \operatorname{trop}(f)(A^T w) = \min_{c_u \neq 0}(\operatorname{val}(c_u) + wAu) = \operatorname{trop}(\phi(f))(w)$. Thus

$$\phi^*(\operatorname{in}_{\phi(w)}(f)) = \phi^*(\sum_{\operatorname{val}(c_u)+wAu=W} \overline{t^{-\operatorname{val}(c_u)}c_u} x^u)$$
$$= \sum_{\operatorname{val}(c_u)+wAu=W} \overline{t^{-\operatorname{val}(c_u)}c_u} x^{Au}$$
$$= \operatorname{in}_w(\phi(f)).$$

Thus $\phi^*(\operatorname{in}_{\phi(w)}(I')) \subseteq \operatorname{in}_w(I)$. If $\operatorname{in}_{\phi(w)}(I') = \langle 1 \rangle$, then $1 = \phi^*(1) \in \phi^*(\operatorname{in}_{\phi(w)}(I')) \subseteq \operatorname{in}_w(I)$, so if $\operatorname{in}_w(I) \neq \langle 1 \rangle$ we also have $\operatorname{in}_{\phi(w)}(I') \neq \langle 1 \rangle$.

Corollary 2.5.8. Let $\phi^* \colon K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}] \to K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ be a monomial automorphism, let $I \in K[x^{\pm 1}, \ldots, x_n^{\pm 1}]$ and let $I' = \phi^{*-1}(I)$. Then $in_w(I) = \langle 1 \rangle$ if and only if $in_{\phi(w)}(I') = \langle 1 \rangle$.

2.6 Exercises

- 1. Show that the residue field of $\Bbbk\{\{t\}\}\$ is isomorphic to \Bbbk .
- 2. Let $K = \mathbb{Q}$ with the *p*-adic valuation. Show that the residue field of K is isomorphic to $\mathbb{Z}/p\mathbb{Z}$.
- 3. Show that if K is an algebraically closed field with a valuation val : $K^* \to \mathbb{R}$, and $\mathbb{k} = R/\mathfrak{m}$ is its residue field, then \mathbb{k} is algebraically closed. Give an example to show that if \mathbb{k} is algebraically closed it does not automatically follow that K is algebraically closed.
- 4. Apply the algorithm implicit in the proof that $\mathbb{C}\{\{t\}\}\$ is algebraically closed to compute (the start of) a solution to the equation $x^2+t+1=0$.

- 5. (a) Show that if $\phi \colon \mathbb{C}^* \to \mathbb{C}^*$ is a homomorphism of algebraic groups, then ϕ has the form $\phi(x) = x^n$ for some $n \in \mathbb{Z}$.
 - (b) Deduce that $\operatorname{Hom}_{alg}(\mathbb{T}^n, \mathbb{C}^*) \cong \mathbb{Z}^n$.
 - (c) Conclude that the group of automorphisms of \mathbb{T}^n as an algebraic group is $\operatorname{GL}(n,\mathbb{Z})$.
- 6. Show that for a polyhedron σ in a polyhedral complex Σ the fan $\operatorname{star}_{\Sigma}(\sigma)$ defined in Definition 2.3.6 is independent of the choice of w.
- 7. Compute all initial ideals of of $I = \langle 7x_0^2 + 8x_0x_1 x_1^2 + x_0x_2 + 3x_2^2 \rangle \subseteq \mathbb{C}[x_0, x_1, x_2]$, and draw the Gröbner complex of I. Repeat for the ideal $I = \langle tx_1^2 + 3x_1x_2 tx_2^2 + 5x_0x_1 x_0x_2 + 2tx_0^2 \rangle \subseteq \mathbb{C}\{\{t\}\}[x_0, x_1, x_2].$
- 8. Let $I = \langle 7 + 8x_1 x_1^2 + x_2 + 3x_2^2 \rangle \subseteq \mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}]$. Draw the set $\{w \in \mathbb{Q}^2 : \operatorname{in}_w(I) \neq \langle 1 \rangle\}$. Repeat for the ideal $I = \langle tx_1^2 + 3x_1x_2 tx_2^2 + 5x_1 x_2 + 2t \rangle \subseteq \mathbb{C}\{\{t\}\}[x_1^{\pm 1}, x_2^{\pm 1}]$.

Chapter 3

Tropical Varieties

In this chapter we introduce the main player of this book, the *tropical variety*. Here we restrict this term to mean the tropicalization of a classical variety. Many research articles in tropical geometry use a more inclusive notion of tropical varieties, which allows for balanced polyhedral complexes that do not necessarily lift to a classical variety. This issue will be discussed later in Chapter 8 but, for now, we always start with a classical variety and we pass to its tropicalization. Throughout this chapter, the underlying field K is assumed to be algebraically closed with a nontrivial valuation val: $K \to \mathbb{R} \cup \{\infty\}$. The first section concerns the case of hypersurfaces, but thereafter we pass to tropical varieties arising from arbitrary subvarieties of the algebraic torus $\mathbb{T}^n = (K^*)^n$. We shall prove a range of key results about tropical varieties, including the Fundamental Theorem (Theorem 3.2.4) and the Structure Theorem (Theorem 3.3.4).

3.1 Tropical Hypersurfaces

Let $K[x^{\pm}] = K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ denote the ring of Laurent polynomials over the field K. Given a Laurent polynomial $f = \sum c_u \mathbf{x}^{\mathbf{u}}$ in $K[x^{\pm}]$, we define its *tropicalization* trop(f) to be the real-valued function on \mathbb{R}^n that is obtained by replacing each coefficient c_u by its valuation and by performing all additions and multiplications in the tropical semiring (R, \oplus, \odot) . Explicitly,

$$\operatorname{trop}(f)(w) = \min(\operatorname{val}(c_u) + \sum_{i=1}^n u_i w_i) = \min(\operatorname{val}(c_u) + \mathbf{u} \cdot w)$$



Figure 3.1: A tropical line

The tropical polynomial trop(f) is a piecewise linear function $\mathbb{R}^n \to \mathbb{R}$.

The classical variety of the Laurent polynomial $f \in K[x^{\pm}]$ is its hypersurface in the algebraic torus \mathbb{T}^n over the algebraically closed field K:

$$V(f) = \{ \mathbf{v} \in \mathbb{T}^n : f(\mathbf{v}) = 0 \}.$$

We now define the tropical hypersurface associated with the same f:

Definition 3.1.1. The tropical hypersurface trop(V(f)) is the set

 $\{w \in \mathbb{R}^n : \text{ the minimum in } \operatorname{trop}(f) \text{ is achieved at least twice } \}.$

This is the locus in \mathbb{R}^n where piecewise linear function $\operatorname{trop}(f)$ fails to be linear. This can be paraphrased as follows in terms of the initial forms

$$\operatorname{in}_w(f) = \sum_{\substack{u:\operatorname{val}(c_u)+w\cdot u\\ =\operatorname{trop}(f)(w)}} \overline{t^{-\operatorname{val} c_u} c_u} x^u.$$

The tropical hypersurface $\operatorname{trop}(V(f))$ is the topological closure in \mathbb{R}^n of the set of weight vectors $w \in (\Gamma_{\operatorname{val}})^n$ for which $\operatorname{in}_w(f)$ is not a unit in $K[x^{\pm}]$. The equivalence of these two definitions is the easy part of Theorem 3.1.3 below.

Example 3.1.2. Let $K = \mathbb{C}\{\{t\}\}\$ be the field of Puiseux series with complex coefficients. We consider Laurent polynomials $f \in K[x^{\pm 1}, y^{\pm 1}]$.



Figure 3.2: A tropical quadric

1. Let f = x + y + 1. Then trop $(f) = \min(x, y, 0)$, so

 $\operatorname{trop}(V(f)) = \{x = y \le 0\} \cup \{x = 0 \le y\} \cup \{y = 0 \le x\}.$

This is the tropical line that is shown in Figure 3.1.

2. Let $f = t^2 x^2 + xy + (t^2 + t^3)y^2 + (1 + t^3)x + t^{-1}y + t^3$. Then $\operatorname{trop}(f) = \min(2+2x, x+y, 2+2y, x, -1+y, 3)$, so $\operatorname{trop}(V(f))$ consists of the three line segments joining the pairs $\{(-1, 0), (-2, 0)\}, \{(-1, 0), (-1, -3)\},$ and $\{(-1, 0), (3, 4)\}$, and the six rays $\{(-2, 0) + \lambda(0, 1) : \lambda \in \mathbb{R}_{\geq 0}\}, \{(-2, 0) + \lambda(-1, -1) : \lambda \in \mathbb{R}_{\geq 0}\}, \{(-1, -3) + \lambda(-1, -1) : \lambda \in \mathbb{R}_{\geq 0}\}, \{(-1, -3) + \lambda(1, 0) : \lambda \in \mathbb{R}_{\geq 0}\}, \{(3, 4) + \lambda(0, 1) : \lambda \in \mathbb{R}_{\geq 0}\},$ and $\{(3, 4) + \lambda(1, 0) : \lambda \in \mathbb{R}_{\geq 0}\}$. This is illustrated in Figure 3.2.

The following theorem was first stated in the early 1990's in a unpublished manuscript by Mikhael Kapranov. It establishes the link between classical hypersurfaces over a field K with valuation and tropical hypersurfaces in \mathbb{R}^n . In the next section, we state and prove the more general "Fundamental Theorem" which works for varieties of arbitrary codimension. Kapranov's Theorem for hypersurfaces will then serve as the base case for the proof.

Theorem 3.1.3 (Kapranov's Theorem). Fix a Laurent polynomial $f = \sum_{u \in \mathbb{Z}^n} c_u x^u$ in $K[x^{\pm}] = K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. The following three sets coincide:

- 1. the tropical hypersurface $\operatorname{trop}(V(f))$ in \mathbb{R}^n ;
- 2. the closure in \mathbb{R}^n of the set $\{w \in (\Gamma_{val})^n : in_w(f) \text{ is not a monomial }\};$
- 3. the closure of the set $\{(\operatorname{val}(v_1), \ldots, \operatorname{val}(v_n)) : \mathbf{v} \in V(f)\}.$

In addition, if $w = \operatorname{val}(\mathbf{v})$ for $\mathbf{v} \in (K^*)^n$ with $f(\mathbf{v}) = 0$ and n > 1 then $\mathcal{U}_w = \{\mathbf{v}' \in V(f) : \operatorname{val}(\mathbf{v}') = w\}$ is an infinite subset of the hypersurface V(f).

Example 3.1.4. Let $K = \mathbb{C}\{\{t\}\}$ and $f = x + y + 1 \in K[x^{\pm 1}, y^{\pm 1}]$. Then $X = \{(z, -1 - z) : z \in K, z \neq 0, -1\}$, and $\operatorname{trop}(V(f))$ is the tropical line in Figure 3.1. Note that $\operatorname{in}_w(f)$ is a monomial unless w is a positive multiple of $(1, 0), (0, 1), \text{ or } (-1, -1), \text{ in which cases } \operatorname{in}_w(f) \text{ is } y + 1, x + 1, \text{ or } x + y$ respectively, or w = (0, 0) with $\operatorname{in}_w(f) = x + y + 1$. We have

$$(\operatorname{val}(z), \operatorname{val}(-1-z)) = \begin{cases} (\operatorname{val}(z), 0) & \text{if } \operatorname{val}(z) > 0; \\ (\operatorname{val}(z), \operatorname{val}(z)) & \text{if } \operatorname{val}(z) < 0; \\ (0, a) & \text{if } \operatorname{val}(z) = 0, \\ z = -1 + \alpha t^a + \tilde{z}, \\ & \text{with } \operatorname{val}(\tilde{z}) > a > 0; \\ (0, 0) & \text{otherwise.} \end{cases}$$

As z varies in $K \setminus \{0, -1\}$ we get all points in trop(V(f)) with rational coordinates, confirming Theorem 3.1.3.

Proof of Theorem 3.1.3. Let $(w_1, \ldots, w_n) \in \operatorname{trop}(V(f))$. Then by definition the minimum $W = \min_{\mathbf{u}:c_u \neq 0}(\operatorname{val}(c_u) + \mathbf{u} \cdot w) = \operatorname{trop}(f)(w)$ is achieved at least twice. This means that $\operatorname{in}_w(f) = \sum_{\mathbf{u}:c_u \neq 0, \operatorname{val}(c_u) + \mathbf{u} \cdot w = W} \overline{t^{-\operatorname{val}(c_u)} c_u} x^u$ is not a monomial, and thus set 1 is contained in set 2. Conversely, if $\operatorname{in}_w(f)$ is not a monomial, then $\min_{\mathbf{u}:c_u \neq 0} \{\operatorname{val}(c_u) + \mathbf{u} \cdot w\}$ is achieved at least twice, so $w \in \operatorname{trop}(V(f))$. This shows the other containment, so the first two sets are equal.

We now prove the inclusion of set 3 in set 1. Since set 1 is closed, it is enough to consider points in 3 of the form $\operatorname{val}(v) := (\operatorname{val}(v_1), \ldots, \operatorname{val}(v_n))$ where $\mathbf{v} = (v_1, \ldots, v_n) \in (K^*)^n$ satisfies $f(\mathbf{v}) = 0$. Let $\mathbf{v} \in (K^*)^n$ satisfy $f(\mathbf{v}) = 0$, so $\sum_{\mathbf{u} \in \mathbb{Z}^n} c_u \mathbf{v}^{\mathbf{u}} = 0$. This means that $\operatorname{val}(\sum_{\mathbf{u} \in \mathbb{Z}^n} c_u \mathbf{v}^{\mathbf{u}}) = \operatorname{val}(0) =$ $\infty > \operatorname{val}(c_u \mathbf{v}^{\mathbf{u}})$ for all \mathbf{u} with $c_u \neq 0$. Lemma 2.1.1 then implies that the minimum of $\operatorname{val}(c_u \mathbf{v}^{\mathbf{u}}) = \operatorname{val}(c_u) + \mathbf{u} \cdot \operatorname{val}(\mathbf{v})$ for \mathbf{u} with $c_u \neq 0$ must be achieved at least twice, where $\operatorname{val}(\mathbf{v}) = (\operatorname{val}(v_1), \ldots, \operatorname{val}(v_n))$. Thus $\operatorname{val}(\mathbf{v}) \in$ $\operatorname{trop}(V(f))$ as required.

96

3.1. TROPICAL HYPERSURFACES

The final inclusion is of the set 1 into the set 3. This is the content of Proposition 3.1.5, which also proves that the set \mathcal{U}_w is infinite.

Proposition 3.1.5. Let $f = \sum_{\mathbf{u} \in \mathbb{Z}^n} c_u \mathbf{x}^{\mathbf{u}} \in K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and let $w \in \mathbb{Z}^n$ Γ_{val}^n . If $\operatorname{in}_w(f)$ is not a monomial, let $\alpha \in (\mathbb{k}^*)^n$ satisfy $\operatorname{in}_w(f)(\alpha) = 0$. Then there exists $y \in (K^*)^n$ with f(y) = 0, $\operatorname{val}(y) = w$, and $\overline{t^{-w}y} = \alpha$. If n > 1then there are infinitely many such y.

Proof. The proof is by induction on n. Suppose first that n = 1. Since the truth of the lemma is not affected by multiplying f by a monomial, we may assume that $f = \sum_{i=0}^{s} c_i x^i = \prod_{j=1}^{s} (a_j x - b_j)$, where $c_0, c_s \neq 0$, and since $\operatorname{in}_w(f)$ is not a monomial we have s > 0. Then $\operatorname{in}_w(f) = \prod_{j=1}^s \operatorname{in}_w(a_j x - b_j)$ by Lemma 2.5.2. Since this is not a monomial, we have $in_w(a_j x - b_j)$ not a monomial for some j, so $val(a_i) + w = val(b_i)$. Set $y = b_i/a_i$. Note that f(y) = 0, and val(y) = w as required.

We now assume that n > 1 and that the proposition holds for all smaller n. The proof now breaks into two cases. The first case is when $in_w(f)|_{x_n=\alpha_n} \neq \infty$ 0. In this case let y_n be one of the infinitely many elements of K^* with $\operatorname{val}(y_n) = w_n, \text{ and } \overline{t^{-w_n}y_n} = \alpha_n. \text{ Let } g(x_1, \dots, x_{n-1}) \in K[x_1^{\pm 1}, \dots, x_{n-1}^{\pm 1}] \text{ be}$ defined by setting $g(x_1, \dots, x_{n-1}) = f(x_1, \dots, x_{n-1}, y_n).$ Then $g = \sum_{\mathbf{u}' \in \mathbb{Z}^{n-1}} d_{u'} x^{u'},$ where $d_{u'} = \sum_{j \in \mathbb{Z}} c_{(u',j)} y_n^j$. Let $W = \operatorname{trop}(f)(w) = \min_{c_u \neq 0} \{ \operatorname{val}(c_u) + w \cdot u \}$. Then

$$\operatorname{in}_{w}(f) = \sum_{u:\operatorname{val}(c_{u})+w\cdot u=W} \overline{t^{-\operatorname{val}c_{u}}c_{u}} x^{u} = \sum_{u:\operatorname{val}(c_{u})+w\cdot u=W} a_{u} x^{u},$$

where $a_u \in \mathbb{k}$ is equal to $\overline{t^{-\operatorname{val}(c_u)}c_u}$. Thus

$$\operatorname{in}_{w}(f)(x_{1},\ldots,x_{n-1},\alpha_{n}) = \sum_{\mathbf{u}'\in\mathbb{Z}^{n-1}} (\sum_{j:\operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} a_{(\mathbf{u}',j)}\alpha_{n}^{j})x^{\mathbf{u}'}.$$

Since this is not the zero polynomial, there is $\mathbf{u}' \in \mathbb{Z}^{n-1}$ with $\sum_{j:\operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} a_{(\mathbf{u}',j)} \alpha_n^j \neq 0.$ Let $w' = (w_1, \ldots, w_{n-1}) \in \Gamma_{\operatorname{val}}^{n-1}$. Note that for such \mathbf{u}' $\operatorname{val}(d_{u'}) + w' \cdot u' \ge \min_{j:\operatorname{val}(c_{(u',j)}) + w \cdot u = W} \operatorname{val}(c_{(u',j)}y_n^j) + w' \cdot u'$ $= \min_{j: \operatorname{val}(c_{(u',j)}) + w \cdot u = W} \operatorname{val}(c_{(u',j)}) + w \cdot (u',j)$ = W

We next note that this inequality is in fact an equality. Indeed

$$\begin{split} \overline{t^{-W+w'\cdot u'}d_{u'}} &= \sum_{j: \operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} \overline{c_{(u',j)}y_n^j t^{-W+w'\cdot u'}} \\ &= \sum_{j: \operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} \overline{c_{(u',j)}y_n^j t^{-\operatorname{val}(c_{(u',j)})-jw_n}} \\ &= \sum_{j: \operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} \overline{c_{(u',j)}t^{-\operatorname{val}(c_{(u',j)})}t^{-jw_n}y_n^j} \\ &= \sum_{j: \operatorname{val}(c_{(\mathbf{u}',j)})+w\cdot u=W} a_{(u',j)}\alpha_n^j \\ &\neq 0, \end{split}$$

so val $(d_{u'}) + w' \cdot u' = W$. Thus for \mathbf{u}' with $\sum_{j: \operatorname{val}(c_{(\mathbf{u}',j)}) + w \cdot u = W} a_{(\mathbf{u}',j)} \alpha_n^j \neq 0$

$$\operatorname{in}_{w'}(g) = \sum_{d_{u'} \neq 0} \overline{t^{-W + w' \cdot u'} d_{u'}} x^{u'}$$

$$= \sum_{d_{u'} \neq 0} \sum_{j \in \mathbb{Z}} \overline{t^{-W + w' \cdot u'} c_{(u',j)} y_n^j} x^{u'}$$

$$= \sum_{d_{u'} \neq 0} \sum_{j \in \mathbb{Z}} \overline{t^{-W + w \cdot u} c_{(u',j)}} \overline{t^{-jw_n} y_n^j} x^{u'}$$

$$= \sum_{d_{u'} \neq 0} (\sum_{j \in \mathbb{Z}} a_{(u',j)} \alpha_n^j) x^{u'}$$

$$= \sum_{u'} (\sum_{j \in \mathbb{Z}} a_{(u',j)} \alpha_n^j) x^{u'}$$

$$= \operatorname{in}_w(f)(x_1, \dots, x_{n-1}, \alpha_n).$$

This means that $\operatorname{in}_{w'}(g)(\alpha_1, \ldots, \alpha_{n-1}) = 0$, so $\operatorname{in}_{w'}(g)$ has a root in the torus, and is thus not a monomial. It then follows by induction that there exist y_1, \ldots, y_{n-1} with $\operatorname{val}(y_i) = w_i$ and $g(y_1, \ldots, y_{n-1}) = 0$. This gives the required infinite number of $y = (y_1, \ldots, y_n)$ with f(y) = 0 and $\operatorname{val}(y) = w$.

The second case is that $\operatorname{in}_w(f)|_{x_n=\alpha_n} = 0$. This means that $\operatorname{in}_w(f) = (x_n - \alpha_n)\tilde{f}$, where $\tilde{f} \in \mathbb{k}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Then for any $w' \in (\operatorname{inval})^n$ with $w'_n = 0$ we have $\operatorname{in}_{w'}(\operatorname{in}_w(f)) = (x_n - \alpha_n)\operatorname{in}_{w'}(\tilde{f})$, so $\operatorname{in}_{w'}(\operatorname{in}_w(f))$ is not a

3.2. THE FUNDAMENTAL THEOREM

monomial. Then by Lemma 2.4.4 we know that $in_{w+\epsilon w'}(f)$ is not a monomial for all such w' and all sufficiently small $\epsilon \in \Gamma_{\text{val}}$. Choose $\mathbf{v} \in \mathbb{Z}^n$ with $gcd(|v_i|) = 1$ for which $in_{w+\epsilon \mathbf{v}}(f)$ is a monomial for all sufficiently small $\epsilon \in \Gamma_{\text{val}}$, and an automorphism $\phi \colon \mathbb{T}_K^n \to \mathbb{T}_K^n$ with $\phi(\mathbf{v}) = \mathbf{e}_1$. Such a \mathbf{v} exists for any polynomial, since by Lemma 2.4.4 it suffices to note that in the constant coefficient case one can always find $\mathbf{v} \in \mathbb{Z}^n$ with $in_{\mathbf{v}}(f)$ a monomial. The existence of the automorphism then follows from Lemma 2.2.9. Note that since ϕ is linear we have $\phi(w + \epsilon \mathbf{v}) = \phi(w) + \epsilon \mathbf{e}_1$. Let $h = \phi^{*-1}(f)$. It suffices to show that we can find $z \in (K^*)^n$ with h(z) = 0 and $\operatorname{val}(z) = \phi(w)$, as then $y = \phi^{-1}(z)$ satisfies f(y) = 0 and $\operatorname{val}(y) = w$. The automorphism $\phi \colon \mathbb{T}_K^n \to \mathbb{T}_K^n$ induces an automorphism, which we also denote by ϕ , of \mathbb{T}_k^n . Let $\tilde{\alpha} = \phi(\alpha) \in \mathbb{T}^n_{\Bbbk}$. Since the property of an initial form being a monomial is invariant under multiplicative automorphisms, we conclude that $in_{\phi(w)+\epsilon e_1}(h)$ is a monomial for all sufficiently small ϵ , so $\operatorname{in}_{\phi(w)}(h)(x_1,\ldots,x_{n-1},\tilde{\alpha}_n) \neq 0$ 0. Thus by the first case the desired z with h(z) = 0 and $val(z) = \phi(w)$ exists. \square

In the rest of this section we examine the polyhedral geometry of tropical hypersurfaces, making the connection to the topics discussed in Section 2.3.

Proposition 3.1.6. Let $f \in K[x^{\pm}]$ be a Laurent polynomial. The tropical hypersurface trop(V(f)) is the support of a pure polyhedral complex of dimension n-1.

Proof.

Proposition 3.1.7. Suppose that all coefficients of the Laurent polynomial f have zero valuation. Then $\operatorname{trop}(V(f))$ is the support of (n-1)-dimensional polyhedral fan. This fan is the (n-1)-dimensional skeleta of the normal fan to the Newton polyhedron of f.

Proof.

3.2 The Fundamental Theorem

The goal of this section is to prove the fundamental theorem of tropical algebraic geometry, which establishes a tight connection between classical varieties and tropical varieties. We must begin by defining the latter objects.

Definition 3.2.1. Let I be an ideal in the Laurent polynomial ring $K[x^{\pm}] = K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ and let X = V(I) be its variety in the algebraic torus \mathbb{T}^n . The *tropicalization* trop(X) of the variety X is the intersection of all tropical hypersurfaces defined by Laurent polynomials in the ideal I. In symbols,

$$\operatorname{trop}(X) = \bigcap_{f \in I} \operatorname{trop}(V(f)) \subseteq \mathbb{R}^n.$$
(3.1)

Note that $\operatorname{trop}(X)$ depends only on the radical ideal \sqrt{I} and not on I itself. By a *tropical variety* in \mathbb{R}^n we mean any subset if the form $\operatorname{trop}(X)$ where X is a subvariety of the torus \mathbb{T}^n over some field K with a valuation.

In this definition, it does not suffice to take the intersection over the tropical hypersurfaces $\operatorname{trop}(V(f))$ where f runs over any generating set of I.

Example 3.2.2. Let n = 3, $K = \mathbb{C}\{\{t\}\}$, and $I = \langle x + y + z, x + 2y \rangle$. Then trop $(V(x + 2y)) = \{(w_1, w_2, w_3) \in \mathbb{R}^3 : w_1 = w_2\}$ and trop $(V(x + y + z)) = \{(w_1, w_2, w_3) \in \mathbb{R}^3 : w_1 = w_2 \le w_3 \text{ or } w_2 = w_3 \le w_1 \text{ or } w_3 = w_1 \le w_2\}$. The intersection of these two tropical hypersurfaces

$$\operatorname{trop}(V(x+y+z)) \cap \operatorname{trop}(V(x+2y)) = \{(w_1, w_2, w_3) \in \mathbb{R}^3 : w_1 = w_2 \le w_3\}.$$

However, we have $z - y \in I$, and $\operatorname{trop}(V(z - y)) = \{(w_1, w_2, w_3) : w_2 = w_3\}$ does not contain some of the points in this intersection, such as (1, 1, 2). Thus we cannot compute a tropical variety by just intersecting the given hypersurfaces, but we usually have to enlarge the basis of the ideal I. \Box

This brings us back to the notion of a tropical basis, as in Section 2.5.

Corollary 3.2.3. Every tropical variety is a finite intersection of tropical hypersurfaces. More precisely, if \mathcal{T} is a tropical basis of the ideal I then

$$\operatorname{trop}(X) = \bigcap_{f \in \mathcal{T}} \operatorname{trop}(V(f)).$$

Proof. Suppose that $w \in \mathbb{R}^n$ is not in trop(X). Then there exists $f \in I$ such that $\operatorname{in}_w(f)$ is a monomial, and thus a unit in $K[x^{\pm}]$. By the definition of tropical basis, then there exists $f \in \mathcal{T}$ such that $\operatorname{in}_w(f)$ is a unit. This means $w \notin \operatorname{trop}(V(f))$.

3.2. THE FUNDAMENTAL THEOREM

In Example 3.2.2, the two given generators are not yet a tropical basis of the ideal I. However, we get a tropical basis if we add one more polynomial:

$$\mathcal{T} = \{x + y + z, x + 2y, y - z\}.$$

The tropical variety is the intersection of these three tropical hypersurfaces:

$$\operatorname{trop}(X) = \operatorname{trop}(V(x+y+z)) \cap \operatorname{trop}(V(x+2y)) \cap \operatorname{trop}(V(y-z)) \\ = \{(w_1, w_2, w_3) \in \mathbb{R}^3 : w_1 = w_2 = w_3\}.$$

We now come to the main result of this section, which is the direct generalization of Kapranov's Theorem from hypersurfaces to arbitrary varieties.

Theorem 3.2.4 (Fundamental Theorem of Tropical Algebraic Geometry). Let I be an ideal in $K[x^{\pm}]$ and X = V(I) its variety in the algebraic torus $\mathbb{T}^n \cong (K^*)^n$. Then the following three subsets of \mathbb{R}^n coincide:

- 1. The tropical variety trop(X) as defined in equation (3.1);
- 2. the closure in \mathbb{R}^n of the set of all vectors $w \in (\Gamma_{\text{val}})^n$ with $\text{in}_w(I) \neq \langle 1 \rangle$;
- 3. the closure in \mathbb{R}^n of the set of coordinatewise valuations of points in X:

$$\operatorname{val}(X) = \{ (\operatorname{val}(u_1), \dots, \operatorname{val}(u_n)) : (u_1, \dots, u_n) \in X \}.$$

The rest of this section is devoted to proving Theorem 3.2.4. We begin with a sequence of lemmas whose purpose is to get prepared for that proof.

Recall from commutative algebra that a minimal associated prime of an ideal I is a prime ideal $P \supset I$ for which there is no prime ideal Q with $P \supseteq Q \supset I$. The variety V(I) has a decomposition as $\cup_{P \text{ minimal }} V(P)$.

Lemma 3.2.5. Let $X \subset \mathbb{T}^n$ be an irreducible variety of dimension d, with prime ideal $I \subset K[x^{\pm}]$, and $w \in \operatorname{trop}(X) \cap \Gamma_{\operatorname{val}}^n$. Then all minimal associated primes of the initial ideal $\operatorname{in}_w(I)$ in $\mathbb{K}[x^{\pm}]$ have the same dimension d.

Proof. Let $I_{\text{proj}} \subseteq K[x_0, x_1, \ldots, x_n]$ be as in Definition 2.2.5. Then I_{proj} is prime of dimension d+1, so by Lemma 2.4.8 all minimal primes of $\operatorname{in}_{(0,w)}(I_{\text{proj}})$ have dimension d+1, and thus by the Principal Ideal Theorem all minimal primes of $\operatorname{in}_{(0,w)}(I_{\text{proj}}) + \langle x_0 - 1 \rangle$ have dimension at least d. Since $\operatorname{in}_{(0,w)}(I_{\text{proj}})$ is homogeneous by Lemma 2.4.2, all minimal primes are homogeneous and containined in $\langle x_0, \ldots, x_n \rangle$, so do not contain $x_0 - 1$. Thus minimal primes of $\operatorname{in}_{(0,w)}(I_{\operatorname{proj}}) + \langle x_0 - 1 \rangle$ have dimension exactly d. By Proposition 2.5.1 we have that $\operatorname{in}_w(I) = \operatorname{in}_{(0,w)}(I_{\operatorname{proj}})|_{x_0=1}$ viewed as an ideal in $\Bbbk[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, so minimal primes of $\operatorname{in}_w(I)$ are the images in $\Bbbk[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ of those primes minimal over $\operatorname{in}_{(0,w)}(I_{\operatorname{proj}}) + \langle x_0 - 1 \rangle$ that do not contain any monomial in x_1, \ldots, x_n , which all thus have dimension d. Since the dimension of $\operatorname{in}_w(I)$ is the minimum of the dimensions of all minimal primes, this implies that $\operatorname{dim}(\operatorname{in}_w(I)) = d$.

The proof Theorem 3.2.4 will proceed by projecting to the hypersurface case. We need to know that we can choose this projection sufficiently nicely, which the following proposition guarantees.

Proposition 3.2.6. Let X be a subvariety of \mathbb{T}^n . Then there exists a projection $\phi \colon \mathbb{T}^n \to \mathbb{T}^m$ for which the image $\phi(X) \subset \mathbb{T}^m$ is closed in the Zariski topology. This projection can be chosen so that the kernel of the associated map $\phi \colon \mathbb{R}^n \to \mathbb{R}^m$ does not nontrivially intersect a finite number of subspaces of \mathbb{R}^n .

Proof.

Much of the power of tropical geometry comes from the fact that the Gröbner definition (Part 2 of Theorem 3.2.4) gives a polyhedral complex structure to the tropical variety, as we now see.

Proposition 3.2.7. Let I be an ideal in $K[x^{\pm 1}]$. The set $\{w \in \mathbb{R}^n : in_w(I) \neq \langle 1 \rangle\}$ is the support of an (Γ_{val}) -rational polyhedral complex.

Proof. Let I_{proj} be as in Proposition 2.5.1. The Gröbner complex $\Sigma(I_{\text{proj}})$ is an (Γ_{val}) -rational polyhedral complex in \mathbb{R}^{n+1} by Theorem 2.4.11. By Proposition 2.5.1 we have $\operatorname{in}_w(I) = \langle 1 \rangle$ if and only if $1 \in \operatorname{in}_{(0,w)}(I_{\text{proj}})|_{x_0=1}$. This occurs if and only if there is an element in $\operatorname{in}_{(0,w)}(I_{\text{proj}})$ that is a polynomial in x_0 times a monomial in x_1, \ldots, x_n , and thus if and only if there is a monomial in $\operatorname{in}_{(0,w)}(I_{\text{proj}})$, since $\operatorname{in}_{(0,w)}(I_{\text{proj}})$ is homogeneous by Lemma 2.4.2. So $\{w \in \mathbb{R}^n : \operatorname{in}_w(I) \neq \langle 1 \rangle\} = \{w : \operatorname{in}_{(0,w)}(I_{\text{proj}}) \text{ does not contain a monomial}\}.$ It is thus a union of slices of polyhedra in $\Sigma(I_{\text{proj}})$, and so a polyhedral complex.

The polyhedral complex structure guaranteed by Proposition 3.2.7 is not unique. We next note the following bound on its dimension. This will be further improved to an equality in Theorem 3.3.7, whose proof uses Theorem 3.2.4.

3.2. THE FUNDAMENTAL THEOREM

Lemma 3.2.8. Let $X \subset \mathbb{T}^n$ have dimension d, with ideal $I \subseteq K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Then every polyhedron in the polyhedral complex Σ whose support is the set $\{w \in (\Gamma_{val})^n : in_w(I) \neq \langle 1 \rangle\}$ has dimension at most d.

Proof. Let $w \in (\Gamma_{val})^n$ lie in the relative interior of a maximal-dimensional polyhedron $P \in \Sigma$. Let the affine span of P = w + L, where L is a subspace of \mathbb{R}^n . By Lemma 2.2.9 and Corollary 2.5.8 we may assume that L is the span of $\mathbf{e}_1, \ldots, \mathbf{e}_k$ for some k. We need to show that $k = \dim(L) \leq d$. Since w lies in the relative interior of P, $\operatorname{in}_{w+\epsilon \mathbf{v}}(I) \neq \langle 1 \rangle$ for all $\mathbf{v} \in \mathbb{Z}^n \cap L$ and $\epsilon \in (\Gamma_{\rm val})$ sufficiently small. Thus Lemma 2.4.4 and Proposition 2.5.1 imply that $\operatorname{in}_{\mathbf{v}}(\operatorname{in}_{w}(I)) = \operatorname{in}_{w}(I)$ for all $\mathbf{v} \in L \cap \mathbb{Z}^{n}$. Choose a generating set \mathcal{G} for $in_w(I)$ with the property that no element is the sum of two other polynomials in $\operatorname{in}_w(I)$ containing fewer monomials. Then $f \in \mathcal{G}$ implies that $\operatorname{in}_v(f) = f$ for all $\mathbf{v} \in L$, as $\operatorname{in}_{\mathbf{v}}(f)$ is otherwise a polynomial in $\operatorname{in}_{w}(I)$ containing fewer monomials. In particular $in_{\mathbf{e}_i}(f) = f$ for $1 \leq i \leq k$, so f = mf, where m is a monomial, and x_1, \ldots, x_k do not appear in f. Since monomials are units in $\mathbb{k}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, this means that $\mathrm{in}_w(I)$ has a generating set where no generator contains x_1, \ldots, x_k , and thus $k \leq \dim(\operatorname{in}_w(I)) \leq \dim(X) = d$ as required.

We now use Theorem 3.1.3 to prove Theorem 3.2.4.

Proof of Theorem 3.2.4. Let $(val(u_1), \ldots, val(u_n))$ lie in set 3, so $u = (u_1, \ldots, u_n) \in X$. Then for any $f \in I$ we have f(u) = 0, so by Proposition 3.1.3 we know $(val(u_1), \ldots, val(u_n)) \in trop(V(f))$, and thus in set 1. This means that $(val(u_1), \ldots, val(u_n))$ lies in set 1. Since set 1 is a closed set by construction, we have set 3 contained in set 1.

Next, let **w** lie in set 1. Then for any $f = \sum c_u x^u \in I$ the minimum of $\{\operatorname{val}(c_u) + \mathbf{u} \cdot \mathbf{w} : c_u \neq 0\}$ is achieved twice. Thus $\operatorname{in}_w(f)$ is not a monomial, so by Lemma 2.5.2 we see that $\operatorname{in}_w(I)$ is not equal to $\langle 1 \rangle$, so **w** lies in set 2.

It thus remains to prove that set 2 is contained in set 3. We first reduce to the case where X is irreducible. Since $\operatorname{in}_w(f^r) = \operatorname{in}_w(f)^r$ for all f, r, $\operatorname{in}_w(I) = \langle 1 \rangle$ if and only if $\operatorname{in}_w(\sqrt{I}) = \langle 1 \rangle$, so we may assume that I is radical. Thus we can write $I = \bigcap_{i=1}^s P_i$, where P_i is prime, and $V(P_1), \ldots, V(P_s)$ are the irreducible components of X. Note that if $w \in \Gamma_{\operatorname{val}}^n$ has $\operatorname{in}_w(I) \neq \langle 1 \rangle$ then there is a $1 \leq j \leq s$ for which $\operatorname{in}_w(P_j) \neq \langle 1 \rangle$. Indeed, if not, by Lemma 2.5.2 there are f_1, \ldots, f_s with $f_i \in P_i$ and $\operatorname{in}_w(f_i) = 1$. Set $f = \prod_{i=1}^s f_i$. Then $\operatorname{in}_w(f) = 1$ and $f \in I$, so $\operatorname{in}_w(I) = \langle 1 \rangle$, contradicting our assumption. Thus if w lies in set 2 for X, then it lies in set 2 for some irreducible component X' of X, and so if we show that $w = \operatorname{val}(x)$ for some $x \in X'$ we will have shown that $w = \operatorname{val}(x)$ for some $x \in X$. The irreducible case is the content of Proposition 3.2.9 below.

Proposition 3.2.9. Let X be an irreducible d-dimensional subvariety of \mathbb{T}^n , with ideal $I = I_X \subseteq K[x^{\pm 1}]$. Fix $w \in \Gamma_{\text{val}}^n$ with $\text{in}_w(I) \neq \langle 1 \rangle$, and $\alpha \in (\mathbb{k}^*)^n \in V(\text{in}_w(I))$. Then there is $y \in X$ with val(y) = w and $\overline{t^{-w}y} = \alpha$. If $\dim(X) > 0$ then there are infinitely many such y.

Remark 3.2.10. In fact the set of $y \in X$ satisfying the conclusion of Proposition 3.2.9 is Zariski dense in X. See [Pay09].

Proof of Proposition 3.2.9. The proof is by induction on n. The base case is n = 1, where it follows from Proposition 3.1.5. Suppose now that n > 1. The case where X is a hypersurface is Proposition 3.1.5, so we may assume that $d = \dim(X) < n - 1$.

By Proposition 3.2.7 there is a polyhedral complex Σ whose support is the closure of $\{w' \in \Gamma_{\text{val}}^n : \operatorname{in}_{w'}(I) \neq \langle 1 \rangle\}$. By Lemma 3.2.8 we know that every polyhedron in this complex has dimension at most d. For each $P \in \Sigma$, let H_P be the affine subspace of \mathbb{R}^n spanned by P and w, with $H_P = w + L_P$ for a subspace L_P of \mathbb{R}^n . Then $\dim(H_P) = \dim(L_P) \leq d+1 < n$. Choose a projection $\phi: \mathbb{T}^n \to \mathbb{T}^{n-1}$ so that the associated map $\phi: \mathbb{R}^n \to \mathbb{R}^{n-1}$ satisfies $\ker(\phi) \cap L_P = \{0\}$ for all $P \in \Sigma$. Let $Y \subset T_{\mathbb{k}}^n$ be the variety $\alpha^{-1}V(\operatorname{in}_w(I))$, which has dimension $\dim(I) = d$ by Lemma 3.2.5. Choose ϕ , as guaranteed by Proposition 3.2.6, so the image $\phi(X)$ of the projection is closed in \mathbb{T}^{n-1} , and so that $\ker(\phi) \cap Y = 1$. If $\operatorname{in}_{w'}(I) \neq \langle 1 \rangle$ for some $w' \in \Gamma_{\operatorname{val}}^n$ then $w' \in H_P$ for some P, so $w' - w \in L_P$. Thus $\phi(w') = \phi(w)$ and $\operatorname{in}_{w'}(I) \neq \langle 1 \rangle$ implies that w = w'. Also, if $\phi(\alpha') = \phi(\alpha)$ for some $\alpha' \in V(\operatorname{in}_w(I)) \subseteq T_{\mathbb{k}}^n$ then $\alpha'/\alpha \in Y \cap \ker(\phi)$, which our choice of ϕ implies is 1, so $\alpha = \alpha'$.

Let $X' = \phi(X) = V(I')$, where $I' = \phi^{*-1}(I)$). By Lemma 2.5.7 $\operatorname{in}_{\phi(w)}(I') \neq \langle 1 \rangle$. Let $\alpha' = \phi(\alpha)$. By the induction assumption there is $z \in \mathbb{T}^{n-1}$ with $z \in X'$, $\operatorname{val}(z) = \phi(w)$, and $\overline{t^{-\phi(w)}z} = \phi(\alpha)$. If $\dim(X') > 0$ then there are infinitely many such z.

Since $\phi(X) = \phi(X)$ we can find $y \in X$ with $\phi(y) = z$. If $\dim(X') > 0$ there are thus infinitely many such y. If $\dim(X') = 0$, then since X is irreducible it must be the single point z, so if $\dim(X) > 0$ all infinitely many points in X map to z, so there are again infinitely many such y. Since $\phi(w) = \operatorname{val}(\phi(x)) = \phi(\operatorname{val}(x))$ by Equation 2.6, by our choice of ϕ we have $\operatorname{val}(y) = w$. It remains to prove that $\alpha' = \overline{t^{-w}y}$ equals α . By our choice of ϕ Figure 3.3: Some examples of balanced fans

it suffices to show that $\alpha' \in V(\operatorname{in}_w(I), \operatorname{so} \operatorname{in}_w(f)(\alpha') = 0$ for all $f \in I$. Fix $f = \sum c_u x^u \in I$. Then $\operatorname{in}_w(f) = \sum \overline{t^{-\operatorname{trop}(f)(w)} c_u t^{-w \cdot u} x^u}$, so

$$in_w(f)(\alpha') = \sum \overline{t^{-\operatorname{trop}(f)(w)} c_u t^{w \cdot u}} \alpha'^u$$
$$= \sum \overline{t^{-\operatorname{trop}(f)(w)} c_u y^u}$$
$$= \overline{t^{-\operatorname{trop}(f)(w)} f(y)}$$
$$= 0.$$

This proves the existence of such a $y \in X$ with $\operatorname{val}(y) = w$ and $\overline{t^{-w}y} = \alpha$, and the existence of infinitely many such y if $\dim(X) > 0$.

3.3 The Structure Theorem

In what follows we explore the question of which polyhedral complexes are tropical varieties. The main result is the Structure Theorem (Theorem 3.3.4) which says that if X is an irreducible subvariety of \mathbb{T}^n of dimension d then $\operatorname{trop}(X)$ is the support of a pure d-dimensional weighted balanced $\Gamma_{\text{val-}}$ rational polyhedral complex that is connected in codimension one.

We first define these concepts. Let $\Sigma \subset \mathbb{R}^n$ be a one-dimensional fan with s rays. Let \mathbf{u}_i be the first lattice point on the *i*th ray of Σ . We give Σ the structure of a *weighted fan* by assigning a weight $m_i \in \mathbb{N}$ to the *i*th ray of Σ . We say that Σ is *balanced* if

$$\sum m_i \mathbf{u}_i = 0.$$

This is sometimes called the *zero-tension condition*; a tug-of-war game with ropes in the directions \mathbf{u}_i and participants of strength m_i would have no winner. See Figure 3.3 for some examples. We now extend this concept to arbitrary weighted polyhedral complexes.

Definition 3.3.1. Let $\Sigma \subseteq \mathbb{R}^n$ be a rational polyhedral fan, pure of dimension d, and fix $m(\sigma) \in \mathbb{N}$ for all cones σ of dimension d.

Given $\tau \in \Sigma$ of dimension d-1, let L be the affine span of τ , which is a (d-1)-dimensional subspace of \mathbb{R}^n . Note that since τ is a rational cone, $L_{\mathbb{Z}} = L \cap \mathbb{Z}^n$ is a free abelian group of rank d-1 with $N_{\tau} = \mathbb{Z}^n/L_{\mathbb{Z}} \cong \mathbb{Z}^{n-d+1}$. For each $\sigma \in \Sigma$ with $\tau \subsetneq \sigma$ the cone $(\sigma + L)/L$ is a one-dimensional cone (ray) in N_{τ} . Let \mathbf{u}_{σ} be the first lattice point on this ray.

The fan Σ is balanced at τ if

$$\sum m(\sigma)\mathbf{u}_{\sigma} = 0$$

The fan Σ is *balanced* if it is balanced at all $\tau \in \Sigma$ with dim $(\tau) = d - 1$.

If Σ is a pure Γ -rational polyhedral complex of dimension d with weights $m(\sigma) \in \mathbb{N}$ on each d-dimensional polyhedron in Σ , then for each $\sigma \in \Sigma$ the fan star_{Σ}(σ) inherits a weighting function m. The polyhedral complex Σ is balanced if the fan star_{Σ}(σ) is balanced for all $\sigma \in \Sigma$ with dim(σ) = d - 1.

We next define what it means for a polyhedral complex to be connected in codimension one.

Definition 3.3.2. Let $\Sigma \subset \mathbb{R}^n$ be a pure *d*-dimensional polyedral complex. The complex Σ is *connected in codimension one* if for any two *d*-dimensional polyhedra $P, P' \in \Sigma$ there is a chain $P = P_1, P_2, \ldots, P_s = P'$ for which P_i and P_{i+1} share a common facet F_i for $1 \leq i \leq s - 1$. Since the P_i are facets of Σ and the F_i are ridges, we call this a facet-ridge path connecting P and P'.

Example 3.3.3. A pure one-dimensional polyhedral complex is connected in codimension one if and only if it is connected. An example of a connected two-dimensional polyhedral complex that is not connected in codimension one is shown in Figure 3.4.

This lets us state the structure theorem, whose proof takes the rest of the chapter.

Theorem 3.3.4 (Structure Theorem for Tropical Varieties). Let X be an irreducible d-dimensional subvariety of \mathbb{T}^n . Then $\operatorname{trop}(X)$ is the support of a balanced weighted Γ_{val} -rational polyhedral complex pure of dimension d. If $\operatorname{char}(K) = 0$ then this complex is connected in codimension-one.

Proof. That $\operatorname{trop}(X)$ is a pure $\Gamma_{\operatorname{val}}$ -rational *d*-dimensional polyhedral complex is Theorem 3.3.7. That it is balanced is Theorem 3.4.8, and in characteristic zero the connectivity result is Theorem 3.5.1.

106

Figure 3.4: A polyhedral complex that is not connected in codimension one



In the remainder of this section we prove the dimension part of this theorem (Theorem 3.3.7). This will use the following proposition, which says that the star of any polyhedron in a polyhedral complex structure on trop(X) is itself a tropical variety.

Proposition 3.3.5. Let $X \subset \mathbb{T}_K^n$, with X = V(I) for $I \subseteq K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ and let Σ be a polyhedral complex whose support is the closure of $\{w \in (\Gamma_{\text{val}})^n : \text{in}_w(I) \neq \langle 1 \rangle\} \subset \mathbb{R}^n$. Fix $w \in \Sigma \cap (\text{im val})^n$, and let σ be the polyhedron of Σ containing w in its relative interior. Then

$$\operatorname{star}_{\Sigma}(\sigma) = \{ v \in (\Gamma_{\operatorname{val}})^n : \operatorname{in}_v(\operatorname{in}_w(I)) \neq \langle 1 \rangle \}$$

Proof. We have

$$\{v \in \mathbb{R}^n : \operatorname{in}_v(\operatorname{in}_w(I)) \neq \langle 1 \rangle\} = \{v \in \mathbb{R}^n : \operatorname{in}_{w+\epsilon v}(I) \neq \langle 1 \rangle \text{ for sufficiently small } \epsilon > 0\}$$
$$= \{v \in \mathbb{R}^n : w + \epsilon v \in \Sigma \text{ for sufficiently small } \epsilon > 0\}$$
$$= \operatorname{star}_{\Sigma}(\sigma),$$

where the first equality follows from Lemma 2.4.4 and Proposition 2.5.1. \Box

Example 3.3.6. Let $X = V(tx^2 + x + y + xy + t) \subset T_K^2$ for $K = \mathbb{C}\{\{t\}\}$. Then trop(X) is shown in Figure 3.5. Let $I = \langle tx^2 + x + y + xy + t \rangle$. The initial ideal $\operatorname{in}_{(1,1)}(I) = \langle x + y + 1 \rangle$, which has tropical variety the standard tropical line. The star of the vertex (1,1) has rays spanned by (1,0), (0,1), and (-1-1), so these coincide. This is also the star of the vertex (-1,0), which is explained by $\operatorname{in}_{(-1,0)}(I) = \langle x^2 + x + xy \rangle = \langle x + 1 + y \rangle$. At the vertex (0,0) the star has rays (1,1), (-1,0), and (0,-1), which is the tropicalization of $V(\operatorname{in}_{(0,0)}(I)) = V(\langle x + y + xy \rangle)$.



Figure 3.5:

Theorem 3.3.7. Let X be an irreducible subvariety of \mathbb{T}^n of dimension d. Then trop(X) is the support of a pure d-dimensional Γ_{val} -rational polyhedral complex.

Proof. That $\operatorname{trop}(X)$ is the support of a Γ_{val} -rational polyhedral complex Σ follows from Proposition 3.2.7 and Theorem 3.2.4. Lemma 3.2.8 shows that the dimension of each polyhedron in Σ is at most d. It thus remains to show that each maximal polyhedron in Σ has dimension precisely d.

Suppose that $\sigma \in \Sigma$ is a maximal polyhedron of dimension $\dim(\sigma) = k$, and fix $w \in \operatorname{relint}(\sigma)$. Let $I = I_X$. By Proposition 3.3.5 we know that $\operatorname{trop}(\operatorname{in}_w(I)) = \operatorname{star}_{\Sigma}(\sigma)$, which is a translate of the affine span of σ , and thus a subspace of \mathbb{R}^n of dimension k. After an appropriate change of coordinates we may assume that this is the subspace L spanned by $\mathbf{e}_1, \ldots, \mathbf{e}_k$. Since $\operatorname{in}_v(\operatorname{in}_w(I)) = \operatorname{in}_{w+\epsilon v}(I) = \operatorname{in}_w(I)$ for all $v \in \Gamma_{\operatorname{val}}^n \cap L$, I is homogeneous with respect to the grading given by $\operatorname{deg}(x_i) = \mathbf{e}_i$ for $1 \leq i \leq k$ and $\operatorname{deg}(x_i) =$ 0 for i > k. This means that there is a generating set for $\operatorname{in}_w(I)$ where every Laurent polynomial uses only the variables x_{k+1}, \ldots, x_n . Let J = $\operatorname{in}_w(I) \cap K[x_{k+1}^{\pm 1}, \ldots, x_n^{\pm 1}]$. If $v' \in \Gamma_{\operatorname{val}}^{n-k} \cap \operatorname{trop}(V(J))$, then $\operatorname{in}_{v'}(J) \neq \langle 1 \rangle$, so $\operatorname{in}_v(\operatorname{in}_w(I)) \neq \langle 1 \rangle$, where $v = (0, v') \in \Gamma_{\operatorname{val}}^n$ has first k coordinates equal to zero. Thus $\operatorname{trop}(V(J)) = \{\mathbf{0}\}$. It then follows from Lemma 3.3.8 that V(J)is finite, so $\dim(\operatorname{in}_w(I)) \leq k$. By Lemma 3.2.8 $\dim(\operatorname{in}_w(I)) = d$, so $d \leq k$ as required. \Box

Lemma 3.3.8. Let X be a subvariety of \mathbb{T}^n . If $\operatorname{trop}(X)$ is a finite set of points, then X is a finite set of points.
3.4. MULTIPLICITIES AND BALANCING

Proof. The proof is by induction on n. When n = 1 all nontrivial varieties are finite sets of points, so there is nothing to prove. Now suppose that n > 1and the lemma is true for all smaller n. If X is a finite set of points then there is nothing to prove, and If X is a hypersurface then Proposition 3.1.6 implies that $\operatorname{trop}(X)$ is not a finite set of points, so we may assume that $0 < \dim(X) < n - 1$. Choose a projection $\pi \colon \mathbb{T}^n \to \mathbb{T}^{n-1}$ with $Y := \overline{\pi(X)} = \pi(X)$ as guaranteed by Proposition 3.2.6. After change of coordinates we may assume that π is projection onto the first n - 1 coordinates.

Suppose first that $\operatorname{trop}(Y)$ is a finite set of points. Then by the induction hypothesis Y is a finite set of points $y_1, \ldots, y_s \in \mathbb{T}^{n-1}$. This means that $X \subseteq \bigcup_{i=1}^s V(x_1 - (y_i)_1, \ldots, x_{n-1} - (y_i)_{n-1})$. Since $\dim(X) > 0$ we must then have $V(x_1 - (y_i)_1, \ldots, x_{n-1} - (y_i)_{n-1}) \subseteq X$ for some *i*, which implies that the line $\{(\operatorname{val}(y_i), \lambda) : \lambda \in \mathbb{R}\} \subset \mathbb{R}^n$ lies in the finite set $\operatorname{trop}(X)$.

We thus conclude that $\operatorname{trop}(Y)$ is not a finite set. Choose $w_1, \ldots, w_s \in \operatorname{trop}(Y)$ distinct with $s > |\operatorname{trop}(X)|$. By Theorem 3.2.4 there is $y_1, \ldots, y_s \in Y$ with $\operatorname{val}(y_i) = w_i$. Choose $x_i \in X$ with $\pi(x_i) = y_i$ for $1 \leq i \leq s$. Then $\pi(\operatorname{val}(x_i)) = \operatorname{val}(\phi(x_i)) = w_i$, so the points $\operatorname{val}(x_i)$ are all distinct, and thus $|\operatorname{trop}(X)| \geq s$, which is a contradiction. We thus conclude that $\dim(X) = 0$ as required.

3.4 Multiplicities and Balancing

In this section we describe place an extra structure on a tropical variety which gives it the structure of a weighted balanced polyhedral complex.

Given a subvariety $X \subset \mathbb{T}^n$ with defining ideal $I \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Proposition 3.2.7 implies that the tropical variety $\operatorname{trop}(X)$ is the support of polyhedral complex Σ . This polyhedral complex has the property that for any polyhedron $\sigma \in \Sigma$ we have $\operatorname{in}_w(I)$ constant for all $w \in \operatorname{relint}(\sigma)$. The choice of Σ is not unique (as the choice of homogenization of the ideal I is not canonical), but we will assume in the first part of this section that a choice has been made.

We first recall the notion of multiplicity of an associated prime from commutative algebra.

Definition 3.4.1. Let $S = \mathbb{k}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. An ideal $Q \subset S$ is primary if for all $f, g \in S$, if $fg \in Q$ then $f \in Q$ or $g^m \in Q$ for some $m \in \mathbb{N}$. If Q is primary then the radical of Q is a prime P of S. Given an ideal $I \subset S$ we can write $I = \bigcap_{i=1}^{s} Q_i$ where Q_i is primary with radical P_i , no Q_i is redundant, and no P_i is repeated. While this decomposition is not unique in general, the set of P_i appearing is determined, and these are called the *associated primes* of *I*. We write $Ass(I) = \{P_1, \ldots, P_s\}$.

For an S-module M let $\ell(M)$ denote the length as an S-module of M. The *multiplicity* of P_i in I is

$$\operatorname{mult}(P_i, I) = \ell((S/Q_i)_{P_i}) = \ell(((I : P_i^{\infty})/I)_{P_i}).$$

See [Eis95, Chapter 3] for more details.

Example 3.4.2. Let $S = \mathbb{k}[x^{\pm 1}]$, and let $f = \sum_{i=0}^{s} c_i x^i$. Then we can write $f = \alpha \prod_{i=1}^{r} (x - \lambda_i)^{m_i}$ for $\alpha, \lambda_i \in \mathbb{k}$, and $m_i \in \mathbb{N}$ with $\sum_{i=1}^{r} m_i = s$. The associated primes of $\langle f \rangle$ are then $\{\langle x - \lambda_i \rangle : 1 \leq i \leq r\}$, and the multiplicity $\operatorname{mult}(\langle x - \lambda_i \rangle, \langle f \rangle)$ of $\langle x - \lambda_i \rangle$ is m_i .

Definition 3.4.3. Let X be a subvariety of \mathbb{T}^n , with defining ideal $I \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Let Σ be a polyhedral complex that is the support of trop(X) with the property that $\operatorname{in}_w(I)$ is constant for $w \in \operatorname{relint}(\sigma)$ for all $\sigma \in \Sigma$. For a polyhedron $\sigma \in \Sigma$ maximal with respect to inclusion, the *multiplicity* $\operatorname{mult}(\sigma)$ is defined by

$$\operatorname{mult}(\sigma) = \sum_{P \in \operatorname{Ass}(\operatorname{in}_w(I))} \operatorname{mult}(P, \operatorname{in}_w(I))$$

for any $w \in \operatorname{relint}(\sigma)$.

Remark 3.4.4. Geometrically, for σ maximal with respect to inclusion in Σ the variety $V(\text{in}_w(I))$ is a union of *d*-dimensional torus orbits. The multiplicity $m(\sigma)$ is the number of such orbits, counted with multiplicity.

Example 3.4.5. Let $f = xy^2 + 4y^2 + 3x^2y - xy + 8y + x^4 - 5x^2 + 4 \in \mathbb{C}[x^{\pm 1}, y^{\pm 1}]$. Then trop(V(f)) is the one-skeleton of the Newton polygon of f, which consists of four rays, generated by $\mathbf{u}_1 = (1,0)$, $\mathbf{u}_2 = (0,1)$, $\mathbf{u}_3 = (-2,-3)$, and $\mathbf{u}_4 = (0,-1)$. This is illustrated in Figure 3.6. The initial ideals and multiplicities are shown in the following table:

w	$\operatorname{in}_w(\langle f angle)$	m(pos(w))
(1,0)	$\langle 4y^2 + 8y + 4 \rangle = \langle (y+1)^2 \rangle$	2
(0,1)	$\langle x^4 - 5x^2 + 4 \rangle = \langle (x-2)(x-1)(x+1)(x+2) \rangle$	
	$= \langle x - 2 \rangle \cap \langle x - 1 \rangle \cap \langle x + 1 \rangle \cap \langle x + 2 \rangle$	4
(-2, -3)	$\langle xy^2 + x^4 \rangle = \langle y^2 + x^3 \rangle$	1
(0, -1)	$\langle xy^2 + 4y^2 \rangle = \langle x + 4 \rangle$	1



Figure 3.6: The tropical variety of $V(xy^2 + 4y^2 + 3x^2y - xy + 8y + x^4 - 5x^2 + 4)$

Note that

$$2\begin{pmatrix}1\\0\end{pmatrix}+4\begin{pmatrix}0\\1\end{pmatrix}+1\begin{pmatrix}-2\\-3\end{pmatrix}+1\begin{pmatrix}0\\-1\end{pmatrix}=\begin{pmatrix}0\\0\end{pmatrix}.$$

We now make more more precise algebraically the geometric content of Remark 3.4.4. After a multiplicitive change of variables we may transport any polyhedron in Σ to one with affine span the span of $\mathbf{e}_1, \ldots, \mathbf{e}_d$.

Lemma 3.4.6. Let $X \subset T^n$ be irreducible of dimension d with ideal $I \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Fix a polyhedral structure Σ on trop(X). Let $w \in \Gamma_{\text{val}}^n$ lies in a polyhedron σ with affine span $\mathbf{e}_1, \ldots, \mathbf{e}_d$. Let $S' = K[x_{d+1}^{\pm 1}, \ldots, x_n^{\pm 1}]$. Then $\text{mult}(\sigma) = \dim_K(S'/(I \cap S'))$.

Proof.

The multiplicity forces the polyhedral complex $\operatorname{trop}(X)$ to be balanced. We first illustrate this in the constant coefficient case when X is a curve in the plane. In this case $\operatorname{trop}(X)$ is a one-dimensional polyhedral fan, which has a unique coarsest polyhedral complex structure.

Proposition 3.4.7. Let $f \in \mathbb{k}[x^{\pm 1}, y^{\pm 1}]$, and let $\mathbf{u}_1, \ldots, \mathbf{u}_s$ be the first lattice points on the rays of trop(V(f)). Let m_i be the multiplicity of the cone $pos(\mathbf{u}_i)$ in trop(V(f)). Then

$$\sum_{i=1}^{s} m_i \mathbf{u}_i = 0.$$

Proof. We may assume that the \mathbf{u}_i are ordered cyclically in a clockwise order. Let $f = \sum_{(a,b)\in\mathbb{Z}^2} c_{ab}x^a y^b$, and let $P = \operatorname{conv}((a,b): c_{ab} \neq 0)$ be its Newton polytope. Then P is a polygon in \mathbb{R}^2 with s facets. Let F_i be the facet of P whose inner normal is \mathbf{u}_i , and let \mathbf{v}_i be its counter-clockwise vertex. Then $\operatorname{in}_{\mathbf{u}_i}(f) = \sum_{(a,b)\in F_i} c_{ab}x^a y^b$. Write $\mathbf{u}_i = (k,l)$. Then we can write $\operatorname{in}_{\mathbf{u}_i}(f) = x^{a_i}y^{b_i}\sum_{j=0}^t c_j(x^{-l}y^k)^j = \alpha_i x^{a_i}y^{b_i}\prod_{j=1}^t (x^{-l}y^k - \lambda_{ij})$. Thus the sum of the multiplicities of associated primes of $\langle \operatorname{in}_w(f) \rangle$ is t, so $t = m_i$. Note that $m_i(-l,k) = \mathbf{v}_{i+1} - \mathbf{v}_i$, where $\mathbf{v}_{s+1} = \mathbf{v}_1$.

Now $(\sum_{i=1}^{s} m_i \mathbf{u}_i)_1 = \sum_{i=1}^{s} m_i (\mathbf{u}_i)_1 = \sum_{i=1}^{s} (\mathbf{v}_{i+1} - \mathbf{v}_i)_2 = 0$, and $(\sum_{i=1}^{s} m_i \mathbf{u}_i)_2 = \sum_{i=1}^{s} m_i (\mathbf{u}_i)_2 = \sum_{i=1}^{s} (\mathbf{v}_i - \mathbf{v}_{i+1})_1 = 0$, so $\sum_{i=1}^{s} m_i \mathbf{u}_i = \mathbf{0}$ as required. \Box

Proposition 3.4.7 is then the base case of the following theorem, which says that all tropical varieties are weighted balanced polyhedral complexes with the weight function mult of Definition 3.4.3.

Theorem 3.4.8. Let X be a subvariety of \mathbb{T}^n . Fix a polyhedral complex Σ with support trop(X) with $in_w(I_X)$ constant for w in the relative interior of a polyhedron in Σ . Then Σ is a weighted balanced polyhedral complex with the weight function mult of Definition 3.4.3.

Proof.

3.5 Connectivity

The polyhedral complex underlying a tropical variety has a strong connectedness property, which we now describe.

Theorem 3.5.1. Fix char(K) = 0. Let X be an irreducible subvariety of \mathbb{T}^n of dimension d. Then trop(X) is the support of a pure d-dimensional polyhedral complex that is connected in codimension one.

The proof of Theorem 3.5.1 is by induction on the dimension d of X. The base case d = 1 is surprisingly nontrivial, and will be proved in Chapter 6. The hypothesis that char(K) = 0 is almost certainly unnecessary and may be removed in a later draft.

Proposition 3.5.2. Let X be a one-dimensional irreducible subvariety of \mathbb{T}^n . Then trop(X) is connected.

3.5. CONNECTIVITY

Definition 3.5.3. Let Σ_1 and Σ_2 be two polyhedral complexes in \mathbb{R}^n , and let $w \in \Sigma_1 \cap \Sigma_2$. The point w lies in the relative interior of a unique polyhedron σ_i in Σ_i for i = 1, 2. The complexes Σ_1, Σ_2 meet transversely at $w \in \Sigma_1 \cap \Sigma_2$ if the affine span of σ_i is $w + L_i$ for $i = 1, 2, L_1 + L_2 = \mathbb{R}^n$, and $\sigma_1 \cap \sigma_2 \neq \{w\}$.

Proposition 3.5.4. Let X, Y be two subvarieties of \mathbb{T}^n . If $\operatorname{trop}(X)$ and $\operatorname{trop}(Y)$ meet transversely at $w \in \Gamma_{\operatorname{val}}^n$ then $w \in \operatorname{trop}(X \cap Y)$.

Proof. Let $\operatorname{trop}(X)$ be the support of the polyhedral complex $\Sigma_1 \subset \mathbb{R}^n$, and let $\operatorname{trop}(Y)$ be the support of the polyhedral complex Σ_2 . Let $I = I_X, J = I_Y \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$. Let $\sigma_i \in \Sigma_i$ be the polyhedron containing w in its relative interior for i = 1, 2, with the affine span of σ_i equal to L_i . The assumption that $\operatorname{trop}(X)$ and $\operatorname{trop}(Y)$ meet transversely at w means that $L_1 + L_2 = \mathbb{R}^n$. After a change of variables on \mathbb{T}^n we may assume that $w = \mathbf{e}_1, L_1$ contains $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_s$ and L_2 contains $\mathbf{e}_1, \mathbf{e}_{s+1}, \mathbf{e}_n$. As in the proof of Theorem 3.3.7 this means that $\operatorname{in}_w(I)$ is homogeneous with respect to a $\mathbb{Z}^{\dim(L_1)}$ -grading, so we can thus find a generating set f_1, \ldots, f_l for $\operatorname{in}_w(I)$ only using x_{s+1}, \ldots, x_n . Similarly there is a generating set g_1, \ldots, g_m for $\operatorname{in}_w(J)$ only using x_2, \ldots, x_s . Since $\operatorname{in}_w(I), \operatorname{in}_w(J) \neq \langle 1 \rangle$, by the Nullstellensatz there are $y_2, \ldots, y_s \in (\mathbb{k}^*)^{s-1}$ and $z_{s+1}, \ldots, z_n \in (\mathbb{k}^*)^{n-s}$ with $f_i(y) = g_j(z) = 0$ for all i, j.

Note that

$$\operatorname{in}_w(I)\operatorname{in}_w(J) \subseteq \operatorname{in}_w(IJ) \subseteq \operatorname{in}_w(I \cap J) \subseteq \operatorname{in}_w(I) \cap \operatorname{in}_w(J).$$

Since $\operatorname{in}_w(I)$ and $\operatorname{in}_w(J)$ have generating sets involving disjoint sets of variables we have $\operatorname{in}_w(I) \operatorname{in}_w(J) = \operatorname{in}_w(I) \cap \operatorname{in}_w(J)$. This follows, for example, from considering any algorithm to compute the intersection of these two ideals. From this we conclude that $\operatorname{in}_w(I \cap J) = \operatorname{in}_w(I) \cap \operatorname{in}_w(J)$. Since $\mathbf{e}_1 \in L_1 \cap L_2$ and $w = \mathbf{e}_1$ lies in the relative interior of σ_1 and σ_2 , the same argument applies to $w + \epsilon \mathbf{e}_1$ for all sufficiently small ϵ , so $\operatorname{in}_{\mathbf{e}_1}(\operatorname{in}_w(I \cap J)) = \operatorname{in}_w(I \cap J)$, and thus by Lemma 2.5.2 $\operatorname{in}_w(I \cap J)$ is homogeneous with respect to the grading given by $\operatorname{deg}(x_1) = 1$ and $\operatorname{deg}(x_i) = 0$ for i > 0.

We next note that in addition $\operatorname{in}_w(I+J) = \operatorname{in}_w(I) + \operatorname{in}_w(J)$. Since $I, J \subseteq I + J$, the inclusion \supseteq is automatic. Fix $f \in I, g \in J$. We want to show that $\operatorname{in}_w(f+g) \in \operatorname{in}_w(I) + \operatorname{in}_w(J)$. Suppose not. Since $\operatorname{in}_w(\cdot)$ commutes with multiplying by a monomial, we may assume that both f and g are polynomials rather than Laurent polynomials. Write $f = f_0(x_2, \ldots, x_n) + x_1 f_1(x_1, \ldots, x_n)$, and $g = g_0(x_2, \ldots, x_n) + x_1 g_1(x_1, \ldots, x_n)$.

Write $f + g = x_1^a h_0(x_2, \ldots, x_n) + x_1^{a+1} h_1(x_1, \ldots, x_n)$. We may assume that $f \in I, g \in J$ have been chosen so that $a \ge 0$ is minimal over all such polynomial counterexamples. This implies in particular that at least one of $f_0, g_0 \ne 0$. If a = 0, then $\operatorname{in}_w(f + g) = h_0(x_2, \ldots, x_n) = f_0 + g_0 \in \operatorname{in}_w(I) + \operatorname{in}_w(J)$. If $a \ge 1$ then $f_0 = -g_0 \in \operatorname{in}_w(I) \cap \operatorname{in}_w(J) = \operatorname{in}_w(I \cap J)$. Since f_0 is homogeneous with respect to the grading given by $\operatorname{deg}(x_1) = 1$, $\operatorname{deg}(x_i) = 0$ for i > 1, by Lemma 2.5.2 there is $p \in I \cap J$ with $\operatorname{in}_w(p) = f_0 = -g_0$. Then $f' = (f - p)/x_1, g' = (g + p)/x_1$ satisfy $f' \in I, g' \in J, f', g'$ are both polynomials, and $f' + g' = (f + g)/x_1$. If $\operatorname{in}_w(f' + g') \in \operatorname{in}_w(I) + \operatorname{in}_w(J)$ then the same is true for $\operatorname{in}_w(f + g)$, so f', g' give another polynomial counterexample. This contradicts the minimality of a, so we conclude that $\operatorname{in}_w(I + J) = \operatorname{in}_w(I) + \operatorname{in}_w(J)$.

Now for any $t \in \mathbb{k}^*$ the vector $(t, y_2, \dots, y_s, z_{s+1}, \dots, z_n) \in V(\operatorname{in}_w(I)) \cap V(\operatorname{in}_w(J)) = V(\operatorname{in}_w(I) + \operatorname{in}_w(J)) = V(\operatorname{in}_w(I+J))$, so $\operatorname{in}_w(I+J) \neq \langle 1 \rangle$, and thus $w \in \operatorname{trop}(V(I+J)) = \operatorname{trop}(X \cap Y)$.

Corollary 3.5.5. Let X, Y be two subvarieties of \mathbb{T}^n . If $\operatorname{trop}(X)$ and $\operatorname{trop}(Y)$ meet transversely at every point of their intersection then $\operatorname{trop}(X \cap Y) = \operatorname{trop}(X) \cap \operatorname{trop}(Y)$.

Proof. If $w \in \operatorname{trop}(X \cap Y)$ then there is $y \in X \cap Y$ with $\operatorname{val}(y) = w$ by Theorem 3.2.4, so $w = \operatorname{val}(y) \in \operatorname{trop}(X)$ and $w \in \operatorname{trop}(Y)$, and thus $\operatorname{trop}(X \cap Y) \subseteq \operatorname{trop}(X) \cap \operatorname{trop}(Y)$. The reverse inclusion follows from Proposition 3.5.4.

Remark 3.5.6. The hypothesis that $\operatorname{trop}(X)$ and $\operatorname{trop}(Y)$ meet transversely is essential in Corollary 3.5.5. For example, let $X = \{1\} \subset K^*$ and let $Y = \{2\} \subset K^*$. Then $\operatorname{trop}(X) = \operatorname{trop}(Y) = \operatorname{trop}(X) \cap \operatorname{trop}(Y) = \{0\}$, but $X \cap Y = \emptyset$, so $\operatorname{trop}(X \cap Y) = \emptyset$.

Proof of Theorem 3.5.1. That trop(X) is the support of a pure *d*-dimensional polyhedral complex is the content of Theorem 3.3.7, so we need only show that this complex is connected in codimension one. The proof is by induction on $d = \dim(X)$, with the base case d = 1 being Proposition 3.5.2, since connected in codimension-one means connected for a one-dimensional polyhedral complex.

We may thus suppose that $d = \dim(X) \ge 2$, and that the theorem is true for all smaller dimensions. We may also assume that X is not contained in any translate of proper subtorus \mathbb{T}' of \mathbb{T}^n , and that $X \neq \mathbb{T}^n$. Thus for any

3.6. EXERCISES

facet σ of Σ there is a facet σ' whose affine span is not equal to that of σ , as if not Σ would be contained in a *d*-dimensional proper affine subspace L of \mathbb{R}^n , so X would be contained in $\{y \in \mathbb{T}^n : \operatorname{val}(y) \in L\}$, which is a translate of a subtorus. Thus for any pair σ, σ' of facets of Σ with the same affine span there is a facet $\sigma'' \in \Sigma$ whose affine span is not equal to that of σ or σ' . It thus suffices to show that σ and σ' are connected by a facet-ridge path when their affine spans are different.

Fix polyhedra $\sigma, \sigma' \in \Sigma$ with different affine spans. Pick $w \in \operatorname{relint}(\sigma) \cap \Gamma_{\operatorname{val}}^n$ and $w' \in \operatorname{relint} \sigma' \cap \Gamma_{\operatorname{val}}^n$ such that the line joining them does not pass through any vertices of Σ . Pick an affine hyperplane H with rational normal vector that contains w, w' but not any vertices of Σ . To see that this is possible note that after translating by -w the line through w and w' passes through the origin, so is a subspace of \mathbb{R}^n not containing any of the translated vertices. After quotienting by this subspace, we just need to choose a hyperplane through the origin in \mathbb{R}^{n-1} not containing the finitely many images of these translated vertices. Write $H = \{\mathbf{u} \in \mathbb{R}^n : \sum_{i=1}^n a_i u_i = b\}$. For any $y \in (K^*)^n$ with $\operatorname{val}(y) = \mathbf{0}$ we set $f_y = (\prod_{i:a_i>0}(y_i x_i)^{a_i}) - (\prod_{i:a_i<0}(y_i x_i)^{a_i})t^b$. Note that $H = \operatorname{trop}(V(f_y))$ for all such choices of y. Note also that H intersects Σ transversely, so $\operatorname{trop}(X \cap V(f_y)) = \operatorname{trop}(X) \cap H$ by Proposition 3.5.4.

Since X is irreducible and $X \not\subseteq V(f_y)$ the variety $X \cap V(f_y)$ has dimension d-1. For generic choices of y the intersection $X \cap V(f_y)$ is irreducible by [Har77, Theorem III.10.8]. In such a case, fix polyhedra σ_H, σ'_H in the polyhedral complex structure on trop $(X \cap H)$ contained in $\sigma \cap H$ and $\sigma' \cap H$ respectively. By induction we have a facet-ridge path connecting σ_H to σ'_H . Since H does not contain any vertices of Σ , it does not contain any (d-2)-dimensional polyhedra of Σ , so every ridge and facet of this path is contained the intersection of a unique ridge of facet of Σ with H, and thus give a facet-ridge path between σ and σ' . This shows that Σ is connected in codimensionone.

3.6 Exercises

116

Chapter 4

The Tropical Rain Forest

In this chapter we explore some of the diversity of the tropical rain forest. We first revisit the plane curves of Section 1.3, and show how these can be used to give a proof of the classical Bézout and Bernstein theorems. We then investigate surfaces in three dimensions, and see the tropical shadow of such classical phenomena as the 27 lines on the cubic surface. Linear spaces are the simplest possible classical varieties, and their tropical counterparts are similarly fundamental. Similarly, the Grassmannian, which is the most fundamental moduli space in algebraic geometry, has a beautiful tropicalization. The case Gr(2, n) parameterizing two-planes in an *n*-dimensional space tropicalizes to the space of phylogenetic trees from computational biology. Finally, we study the tropicalization of a generic intersection. We show how the tropicalization of a variety defined by generic equations is determined purely by the Newton polytopes of the equations.

4.1 Plane Curves

In this section we revisit the plane curves of Section 1.3. We first describe explicitly the algorithm to draw a plane curve given its polynomial, and then discuss how we can see the intersection theory of plane curves tropically.

4.2 Surfaces

In this section we consider the tropicalization of surfaces in T^3 . This allows us to see tropically the shadows of classical theorems for surfaces, such as the existence of 27 lines on a cubic surface.

4.3 Linear Spaces

We first consider the class of varieties $X \subset T^n$ whose equations are all linear. Let $S = K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, and let $I = \langle f_1, \ldots, f_r \rangle \subset S$ be an ideal in Sminimally generated by linear forms f_1, \ldots, f_r . Write $f_i = a_{i0} + a_{i1}x_1 + \ldots a_{in}x_n$ for $1 \leq i \leq r$. We assume here that all coefficients a_{ij} live in subfield of K with valuation zero. The reader will not lose any generality by assuming that K is the field $\mathbb{C}\{\{t\}\}$ of Puiseux series with complex coefficients (see 2.1.2), and all coefficients live in \mathbb{C} . We now describe the tropical variety of $X = V(I) \subset T^n$.

Example 4.3.1. Let $I = \langle x_1 + x_2 + x_3 + x_4 + 1, x_1 + 2x_2 + 3x_3 \rangle \subset K[x_1^{\pm 1}, x_2^{\pm 1}, x_3^{\pm 1}, x_4^{\pm 1}]$. Then X = V(I) is a two-dimensional subvariety of $T \cong (K^*)^4$.

It is easier to work with $I_{\text{proj}} = \langle a_{i0}x_0 + a_{i1}x_1 + \cdots + a_{in}x_n : 1 \leq i \leq r \rangle \subseteq K[x_0, \ldots, x_n]$. The support $\operatorname{supp}(\ell)$ of a linear form $\ell = \sum a_i x_i \in I_{\text{proj}}$ is $\{i : a_i \neq 0\}$. A non-empty subset C of $\{0, 1, \ldots, n\}$ is said to be *circuit* of X if $C = \operatorname{supp}(\ell)$ for some linear form ℓ in the ideal I, and C is inclusion-minimal with this property. Note that C uniquely determines the linear form ℓ up to scaling. The number of circuits of a d-dimensional linear X in T^n is is at most $\binom{n+1}{d+1}$, and this bound is attained for generic subspaces X. Our first result says that the circuits give a tropical basis for I.

Proposition 4.3.2. Let $X \subset T^n$ be a constant coefficient linear space. The set of linear polynomials in I = I(X) whose support is a circuit of X forms a tropical basis for I. This means a vector $w \in \mathbb{R}^n$ lies in trop(X) if and only if, for any circuit C of the subspace X, the minimum of the numbers w_i , as i ranges over C, is attained at least twice.

Proof. The only-if direction is immediate from the definition of a tropical variety because every circuit is the support of a linear form ℓ that lies in the ideal I. For the if direction suppose that w is not in trop(X). Compute the reduced Gröbner basis of I_{proj} with respect to a term order that refines w. Since we are in the constant coefficient case in a polynomial ring and I_{proj} is homogeneous this is the "classical" (non-valuation) Gröbner basis of [CLO07]. Computing such a Gröbner basis is Gaussian elimination, so we the elements of that reduced Gröbner basis are linear forms that are supported

on circuits. Moreover, the initial ideal $\operatorname{in}_w(I_{\operatorname{proj}})$ is generated by the leading forms of these linear forms, so is prime. In addition these leading forms form a Gröbner basis for $\operatorname{in}_w(I_{\operatorname{proj}})$. Our hypothesis states that $w \notin \operatorname{trop}(X)$, so $\operatorname{in}_w(I_{\operatorname{proj}})$ contains a monomial by 2.5.1. Since $\operatorname{in}_w(I_{\operatorname{proj}})$ is prime, this implies that some variable x_i lies in $\operatorname{in}_w(I_{\operatorname{proj}})$. There must thus be an element f of the reduced Gröbner basis with this as a leading term. In fact the entire leading form must be x_i , as otherwise the remainder on division of x_i by $\operatorname{in}_w(I_{\operatorname{proj}})$ would not be zero. This means that the minimum of w_i for i in the corresponding circuit $C = \operatorname{supp}(f)$ is attained only once. \Box

Example 4.3.3. Let *I* be as in in Example 4.3.1. The circuits are $\{\{1, 2, 3\}, \{0, 2, 3, 4\}, \{0, 1, 3, 4\}, \{0, 1, 2, 4\}\}$, which correspond to the linear forms $\{x_1 + 2x_2 + 3x_3, x_2 + 2x_3 - x_4 - 1, x_1 - x_3 + 2x_4 + 2, 2x_1 + x_2 + 3x_4 + 3\}$. Note that the circuits do not all have the same size here. Proposition 4.3.2 says that

$$\operatorname{trop}(X) = \operatorname{trop}(V(x_1 + 2x_2 + 3_3)) \cap \operatorname{trop}(V(x_2 + 2x_3 - x_4 - 1))$$
$$\cap \operatorname{trop}(V(x_1 - x_3 + 2x_4 + 2)) \cap \operatorname{trop}(V(2x_1 + x_2 + 3x_4 + 3)).$$

We now give a combinatorial description of the tropical variety of a linear space X. A key ingredient will be the *lattice of flats* of X. Let A be the $r \times n + 1$ matrix with entries the coefficients a_{ij} of the defining polynomials f_i , and let B be a $(n + 1 - r) \times n + 1$ matrix whose rows are a basis for ker(A). Thus X is equal to the intersection of the row space of B with the torus T. Let $\mathcal{B} = {\mathbf{b}_0, \ldots, \mathbf{b}_n} \subset K^{n-r+1}$ be the columns of the matrix B. While \mathcal{B} depends on the choice of the matrix B, it is determined up to the action of GL(n - r + 1, K).

The lattice of flats $\mathcal{L}(B)$ of the linear space row(B) has elements the subspaces (flats) of K^{n+1-r} spanned by subsets of \mathcal{B} . We make $\mathcal{L}(B)$ into a poset (partially ordered set) by setting $S_1 \leq S_2$ if $S_1 \subseteq S_2$ for two subspaces S_1, S_2 of K^{n-r} spanned by subsets of \mathcal{B} . The poset $\mathcal{L}(B)$ is actually a lattice of rank n + 1 - r. This means that every maximal chain in $\mathcal{L}(B)$ has length n + 1 - r. See, for example, [Sta97, Chapter 3] for more on lattices.

Example 4.3.4. We continue Example 4.3.1. In this case the matrices A and B are

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 0 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & -2 & 1 & 0 & 1 \\ 1 & -2 & 1 & 0 & 0 \\ 0 & -1 & -1 & 1 & 1 \end{pmatrix}.$$



Figure 4.1: The lattice of flats for the linear space of Example 4.3.1

Figure 4.2:

We thus have $\mathbf{b}_0 = (0, 1, 0)$, $\mathbf{b}_1 = (-2, -2, -1)$, $\mathbf{b}_2 = (1, 1, -1)$, $\mathbf{b}_3 = (0, 0, 1)$, and $\mathbf{b}_4 = (1, 0, 1)$. There are fifteen subspaces of K^3 spanned by subsets of $\mathcal{B} = \{(-2, -2, -1), (1, 1, -1), (0, 0, 1), (1, 0, 1), (0, 1, 0)\}$. These are $\{0\} \cup \{\operatorname{span}(\mathbf{b}_i) : 1 \le i < j \le 5\} \cup \{\operatorname{span}(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}, \operatorname{span}(\mathbf{b}_1, \mathbf{b}_4), \operatorname{span}(\mathbf{b}_1, \mathbf{b}_5), \operatorname{span}(\mathbf{b}_2, \mathbf{b}_4), \operatorname{span}(\mathbf{b}_2, \mathbf{b}_5), \operatorname{span}(\mathbf{b}_3, \mathbf{b}_4), \operatorname{span}(\mathbf{b}_3, \mathbf{b}_5), \operatorname{span}(\mathbf{b}_4, \mathbf{b}_5)\}$. This gives the lattice shown in Figure 4.1.

There is a simplicial complex, called the *order complex*, associated to any poset. This has vertices the elements of the poset, and simplices all proper chains, which are totally ordered subsets of the poset not using the bottom or top elements (0 or \mathbb{R}^n in our case). The order complex of $\mathcal{L}(B)$ is pure of dimension n-r. In the case of our poset there is a nice geometric realization of this simplicial complex, which we now describe.

Definition 4.3.5. Let \mathbf{e}_i be the *i*th standard basis vector on \mathbb{R}^n . Given a subset $\sigma \subset [n]$ we set $\mathbf{e}_{\sigma} = \sum_{i \in \sigma} \mathbf{e}_i$. If V is a subspace of \mathbb{R}^{n-r} spanned by some of the \mathbf{b}_i , set $\sigma(V)$ to be $\{i : \mathbf{b}_i \in V\}$. Map the order complex of the lattice of flats of \mathcal{B} into \mathbb{R}^n by sending a subspace V to $\operatorname{pos}(\mathbf{e}_{\sigma(V)}) + \operatorname{span}(\mathbf{1})$, and a simplex $\{V_1, \ldots, V_s\}$ to $\operatorname{pos}(\mathbf{e}_{\sigma(V_i)} : 1 \leq i \leq s) + \operatorname{span}(\mathbf{1})$, where $\mathbf{1} = (1, \ldots, 1) \in \mathbb{R}^n$.

Example 4.3.6. We continue Example 4.3.1. The fan $\delta(\mathcal{B})$ has 13 rays, corresponding to the five rays spanned by the \mathbf{b}_i and the eight planes spanned by them. There is a two-dimensional cone for every inclusion of a ray into a plane, of which there are 17 in total. The intersection of this fan with the 3-sphere is a graph, which is illustrated in Figure 4.2.

We next show that the tropical variety $\operatorname{trop}(V(I))$ is equal to $\Delta(\mathcal{B})$.

Theorem 4.3.7. Let I be a linear ideal in S. The tropical variety of $X = V(I) \subset T$ is equal to $\Delta(\mathcal{B})$.

Proof. We first show that $\operatorname{trop}(X) \subseteq \Delta(\mathcal{B})$. Suppose $\mathbf{v} \notin \Delta(\mathcal{B})$. Let $V^j = \{i : v_i \geq j\}$. We denote by $\operatorname{span}(V^j)$ the subspace spanned by those \mathbf{b}_i with $i \in V^j$. Let $l = \min\{j : \text{ there exists } \mathbf{b}_i \in \operatorname{span}(V^j) \setminus V^j\}$. If no such l existed, then the subspaces $\operatorname{span}(V^j)$ as j varies would form a chain in the lattice of flats $\mathcal{L}(B)$, and \mathbf{v} would live in the corresponding cone of $\Delta(\mathcal{B})$. Let $F = \operatorname{span}(V^l)$. Pick $\mathbf{b}_k \in F \setminus V^l$. Then $v_k < l$ by the definition of V^l . Since $\{\mathbf{b}_i : i \in V^l\}$ spans F, we can write $\mathbf{b}_k = \sum_{i \in V^l} \lambda_i \mathbf{b}_i$ for $\lambda \in K$. This means that $\mathbf{e}_k - \sum \lambda_i \mathbf{e}_i \in \ker(B) = \operatorname{row}(A)$. Thus $f = x_k - \sum_{i \in V^l} \lambda_i x_i \in I$. Now $\operatorname{in}_{\mathbf{v}}(f) = x_k$, so $\operatorname{in}_{\mathbf{v}}(I) = \langle 1 \rangle$, and so $\mathbf{v} \notin \operatorname{trop}(X)$.

We next show that $\Delta(\mathcal{B}) \subseteq \operatorname{trop}(X)$ by exhibiting for each $\mathbf{v} \in \Delta(\mathcal{B})$ an element $y \in V_K(I)$ with $\operatorname{val}(y) = \mathbf{v}$. Given $\mathbf{v} \in \Delta(\mathcal{B})$, let $V_1 \subset V_2 \subset$ $\cdots \subset V_{n-r} = K^{n-r}$ be the chain of flats labelling a maximal cone of $\Delta(\mathcal{B})$ containing **v**, so dim $(V_i) = i$. Pick $\mathbf{b}_{i_1} \in V_1$, and $\mathbf{b}_{i_j} \in V_j \setminus V_{j-1}$ for $2 \leq i_j \leq i_j \leq i_j \leq j_j \leq$ $j \leq n-r$. Note that $v_{i_j} \geq v_{i_{j+1}}$. After renumbering if necessary we may assume that $i_j = j$, and that the matrix B has been chosen so that the first $(n-r) \times (n-r)$ square submatrix is the identity, which is possible as the \mathbf{b}_{i_i} are linearly independent by construction. This implies that the last $r \times r$ submatrix of A must be invertible, so we may assume that it is the identity matrix (since performing row operations on A corresponds to choosing a different generating set for I). Set $y_i = t^{v_i}$ for $1 \le i \le n-r$. For $n-r+1 \le i \le n$, set $y_i = \sum_{j=1}^{n-r} -a_{(i-n+r)j}t^{v_j}$. Then Ay = 0 by construction, so $y \in V_K(I)$. The valuation $\operatorname{val}(y_i) = v_i$ for $1 \leq i \leq n - r$ by construction. For $n - r + 1 \le i \le n$ we have $\operatorname{val}(y_i) = \min\{v_i : a_{(i-n+r)i} \ne i \le n\}$ $0, 1 \le j \ne n - r$ = v_s for $s = \max\{j : a_{(i-n+r)j} \ne 0, 1 \le j \le n - r\}$. Now if $\operatorname{val}(y_i) = v_i$, then $\mathbf{b}_i \in V_i \setminus V_{i-1}$, so by the choice of \mathbf{v} we have $v_i = v_i$, and thus $\operatorname{val}(y_i) = v_i$. So y is the desired element of $V_K(I)$ with $\operatorname{val}(y) = \mathbf{v}$, and so $\Delta(\mathcal{B}) \subseteq \operatorname{trop}(X)$.

Example 4.3.8. Let $I = \langle x_3 - x_1 + 1, x_4 - x_2 + 1, x_5 - x_2 + x_1 \rangle \subset K[x_1^{\pm 1}, \dots, x_5^{\pm 1}]$. Then the matrix A is



Figure 4.3: The tropical variety of Example 4.3.8

and B is

The lattice of flats has elements $\{0, 1, 2, 3, 4, 5\}$ at the lowest level, and then $\{05, 14, 23, 013, 024, 125, 345\}$ at the next level.

The complex $\Delta(B)$ is thus a two-dimensional fan in $\mathbb{R}^5 = \mathbb{R}^6/1$ with 13 rays and 18 two-dimensional cones. The intersection with the 4-sphere is a graph with 13 vertices and 18 edges. This is the well-studied Peterson graph with three edges (unnecessarily) subdivided. See Figure 4.3.

Every initial ideal $\operatorname{in}_w(I) \subset \mathbb{k}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ is generated by linear forms, so is prime, and thus all the multiplicities are one. Note that the dimension of $\Delta(\mathcal{B})$ is n - r by construction, so $\dim(\operatorname{trop}(X)) = \dim(X)$ as expected. We leave it as an exercise to verify the balancing condition and the rest of the conditions guaranteed by Theorem 3.3.4.

4.4 Grassmannians

The Grassmannian G(r, n) is one of the fundamental moduli spaces in algebraic geometry. It is a smooth projective variety of dimension r(n - r)for which each point corresponds to an r-dimensional subspaces of a fixed *n*-dimensional vector space. The Grassmannian G(r, n) naturally embeds

4.4. GRASSMANNIANS

into $\mathbb{P}^{\binom{n}{r}-1}$. In this section we discuss tropicalizing the subvariety $G^0(r,n) = G(r,n) \cap T^{\binom{n}{r}-1} \subset T^{\binom{n}{r}-1} \subset \mathbb{P}^{\binom{n}{r}-1}$.

We first review the construction of the Grassmannian. Given a vector space V over a field K, we can represent an r-dimensional subspace by the row-space of an $r \times n$ matrix of rank r. An issue with this representation is that different matrices can have the same row-space; specifically, if A and B are two $r \times n$ matrices with the same row-space, then one can be obtained from the other by row operations, so there is an element $G \in GL(r, K)$ with A = GB. We solve this problem by mapping these matrices to the length $\binom{n}{r}$ vector of their $r \times r$ minors. This has coordinates indexed by all subsets I of size r of $[n] = \{1, \ldots, n\}$, where the coordinate indexed by I is the determinant of the $r \times r$ submatrix with columns indexed by I. If A = GBfor some $G \in GL(r, K)$ then the Ith minor of A is $\deg(G)$ times the Ith minor of B, so these represent the same point of $\mathbb{P}^{\binom{n}{r}-1}$.

Example 4.4.1. Let $U \subset \mathbb{C}^4$ be the rowspace of the matrix

$$A = \left(\begin{array}{rrrr} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 \end{array} \right).$$

Note that U is also the rowspace of the matrix

$$B = \left(\begin{array}{rrrr} 3 & 2 & 1 & 0 \\ 0 & 1 & 2 & 3 \end{array}\right).$$

The 2×2 minors of a 2×4 matrix are indexed by the sets $\{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}$. The vector of minors of *A*, listed in this order, is (1,2,3,1,2,1), while the one for *B* is (3,6,9,3,6,3). Note that

$$(1:2:3:1:2:1) = (3:6:9:3:6:3) \in \mathbb{P}^5.$$

Miraculously the set of all such minor-vectors forms a variety. We denote by $K[p_I] = K[p_I : I \subset [n], |I| = r]$ the coordinate ring of $\mathbb{P}^{\binom{n}{r}-1}$. The Plücker ideal $I_{r,n}$ is the set of all polynomials vanishing on such a minor-vector, so is the set of all polynomial relations among the minors of an $r \times n$ matrix. It is generated by the Plücker relations, which are defined as follows. Fix a subset $I \subset [n]$ of size r - 1, and a subset $J \subset [n]$ of size r + 1. For $I \subset [n]$ of size r and $j \in J$ the sign sgn(I, j) of the pair $\{I, j\}$ is the sum of the number of $i \in I$ with i > j and the number of $i \in J$ with i < J. Then the Plücker relation $p_{I,J}$ is

$$p_{I,J} = \sum_{j \in J} (-1)^{\operatorname{sgn}(I,j)} p_{I \cup j} p_{J \setminus j},$$

where $p_{I\cup j} = 0$ if $j \in I$. If $I = I' \cup \{i\}$ and $J = I' \cup \{j, k, l\}$ with i < j < k < lthen $p_{I,J} = p_{I'ij}p_{I'kl} - p_{I'ik}p_{I'jl} + p_{I'il}p_{I'jk}$. Such Plücker relations (for any order of i, j, k, l with the signs appropriately permuted) are called three-term Plücker relations.

The Plücker ideal $I_{r,n}$ is

$$I_{r,n} = \langle p_{I,J} : I, J \subseteq [n], |I| = r - 1, |J| = r + 1 \rangle.$$

The Grassmannian G(r, n) is the variety $V(I_{r,n}) \subset \mathbb{P}^{\binom{n}{r}-1}$.

Example 4.4.2. Consider the case r = 2, n = 4. We denote the six variables of $K[p_I]$ by $p_{12}, p_{13}, p_{14}, p_{23}, p_{24}, p_{34}$. The Plücker relation $p_{1,234}$ is $p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23}$. This is equal, up to sign, to $p_{2,134}, p_{3,124}$, and $p_{4,123}$. All other Plücker relations, such as $p_{1,123}$ are zero, so

$$I_{2,4} = \langle p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23} \rangle.$$

Note that the points $(1:2:3:1:2:1) \in \mathbb{P}^5$ of Example 4.4.1 lies in $V(I_{2,4})$.

The torus $T^{\binom{n}{r}-1}$ of $\mathbb{P}^{\binom{n}{r}-1}$ is the set of all points in $\mathbb{P}^{\binom{n}{r}-1}$ with all coordinates nonzero. This is the quotient of the torus $(K^*)^{\binom{n}{r}}$ by the diagonal action of K^* . We now consider the tropicalization of the intersection $G^0(r,n) = G(r,n) \cap T^{\binom{n}{r}-1}$. By Theorem 3.3.4 trop $(G^0(r,n))$ is a pure r(n-r)-dimensional rational polyhedral fan in $\mathbb{R}^{\binom{n}{r}-1} \cong R^{\binom{n}{r}}/\mathbb{R}\mathbf{1}$.

The Plücker ideal is homogeneous with respect to a \mathbb{Z}^n -grading given by $\deg(p_I) = \sum_{i \in I} \mathbf{e}_i \in \mathbb{Z}^n$. This means that lift of $\operatorname{trop}(G^0(r, n))$ to $\mathbb{R}^{\binom{n}{r}}$ has an *n*-dimensional lineality space, given by the row-space of the $(n \times \binom{n}{r})$ -grading matrix. Explicitly, this is $V = \operatorname{span}(\sum_{i \in I} \mathbf{e}_I : 1 \leq i \leq n) \subseteq \mathbb{R}^{\binom{n}{r}}$. Since $\mathbf{1} \in V$, this descends to give an (n-1)-dimensional lineality space for $\operatorname{trop}(G^0(2, n))$. Geometrically this comes from the (n-1)-dimensional torus action on \mathbb{P}^{n-1} , where we view G(r, n) as parameterizing (r-1)-dimensional planes in \mathbb{P}^{n-1} .



Figure 4.4: The combinatorics of $trop(G^0(2,5))$.

Example 4.4.3. The tropical variety $\operatorname{trop}(G^0(2, 4))$ is the image in $\mathbb{R}^6/\mathbf{1}$ of the tropical hypersurface defined by the polynomial $p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23}$. The tropical hypersurface in \mathbb{R}^6 has lineality space $V = \operatorname{span}((1, 1, 1, 0, 0, 0), (1, 0, 0, 1, 1, 0), (0, 1, 0, 1, 0, 1), (0, 0, 1, 0, 1, 1))$. There are three cones: $V + \operatorname{pos}((1, 0, 0, 0, 0, 1)), V + \operatorname{pos}((0, 1, 0, 0, 1, 0))$ and $V + \operatorname{pos}((0, 0, 1, 1, 0, 0, 0))$. We can identify \mathbb{R}^6/V with \mathbb{R}^2 by sending $\mathbf{e}_{12}, \mathbf{e}_{34}$ to $(1, 0), \mathbf{e}_{13}, \mathbf{e}_{24}$ to (0, 1),and $\mathbf{e}_{14}, \mathbf{e}_{23}$ to (-1, -1). The image of $\operatorname{trop}(G^0(2, 4))$ is then the standard tropical line.

Example 4.4.4. The tropical variety $\operatorname{trop}(G^0(2,5))$ is a six-dimensional fan in \mathbb{R}^9 with a four-dimensional lineality space. The quotient of the $\operatorname{trop}(G^0(2,5))$ by the lineality space is a two-dimensional fan in \mathbb{R}^5 with 10 rays and 15 two-dimensional cones that has an action of the symmetry group S_5 on it. When this is intersected with the 4-sphere we get a graph with 10 vertices and 15-edges with an S_5 -action. This is again the Peterson graph; see Figure 4.4. Compare also Example 4.3.8.

When r = 2 the tropicalization $\operatorname{trop}(G^0(2, n))$ has a particularly nice description. There is a unique coarsest fan structure on $\operatorname{trop}(G^0(2, n))/L$, which is known as the *space of phylogenetic trees*.

When r > 2 the tropical variety $\operatorname{trop}(G^0(r, n))$ is not as nice as when r = 2, as the following examples illustrate. Example 4.4.5 shows that $G^0(r, n)$ can depend on the characteristic of K, while Example 4.4.6 shows that the simplicial complex structure on $\operatorname{trop}(G^0(r, n))$ is not determined by its edges.



Figure 4.5: The tropical variety $\operatorname{trop}(G^0(2,5))$ as a space of phylogenetic trees

Example 4.4.5. This will have the content of [SS04, Section 7].

Example 4.4.6. This will contain the example of G(3,6) (see [SS04, Theorem 5.4]) showing that the tropical Grassmannian is not always a flag complex.

We will then follow with some of the content of [SS04, Section 6].

4.5 Complete Intersections

In this section we will describe how to compute combinatorially the tropical variety of the intersection of generic hypersurfaces.

4.6 Exercises

Chapter 5 Linear Algebra

In classical linear algebra over a field K, there are many equivalent ways of representing a d-dimensional linear subspace of an n-dimensional vector space K^n . For instance, V is the span of d linearly independent vectors, or it is the solution set of n-d independent linear equations. Both of these notions translate naturally to the tropical semiring $(\mathbb{R}, \oplus, \odot)$, but, it turns out that they give rise to different concepts in tropical geometry. The image of a tropically linear map is called a *tropical polytope*, and this is the basic object in tropical convexity theory. The solution set of a finite system of tropical linear equations is a *linear prevariety*, and this is an important object in combinatorial applications of the tropical semiring. From a geometric perspective, it it is more natural to consider *tropicalized linear spaces* and tropical linear spaces. They are parametrized by the tropical Grassmannian and the tropical Dressian. The former arise from classical linear spaces over a field K with a valuation, while the latter are more general polyhedral complexes that share the same good properties. Each of these different definitions has interesting geometric features, and one goal of this chapter is to explain their differences. Our point of departure is the tropical eigenvalue problem.

5.1 Eigenvalues and Eigenvectors

Let A be an $n \times n$ -matrix with entries in the tropical semiring $(\mathbb{R}, \oplus, \odot)$. An *eigenvalue* of A is a real number number λ such that

$$A \odot v = \lambda \odot v \tag{5.1}$$

for some $v \in \mathbb{R}^n$. We say that v an *eigenvector* of the tropical matrix A. The arithmetic operations in the equation (5.1) are tropical. For instance, for n = 2, with $A = (a_{ij})$, the left hand side of (5.1) equals

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \odot \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} a_{11} \odot v_1 \oplus a_{12} \odot v_2 \\ a_{21} \odot v_1 \oplus a_{22} \odot v_2 \end{pmatrix} = \begin{pmatrix} \min(a_{11} + v_1, a_{12} + v_2) \\ \min(a_{21} + v_1, a_{22} + v_2) \end{pmatrix}$$

while the right hand side of (5.1) is equal to

$$\lambda \odot \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} \lambda \odot v_1 \\ \lambda \odot v_2 \end{pmatrix} = \begin{pmatrix} \lambda + v_1 \\ \lambda + v_2 \end{pmatrix}.$$

We represent the matrix $A = (a_{ij})$ by a weighted directed graph G(A)with n nodes labeled 1, 2, ..., n. There is an edge from node i to node j if and only if $a_{ij} < \infty$, and we assign the length a_{ij} to each such edge (i, j). The normalized length of a directed path $i_0, i_1, ..., i_k$ in G is the sum (in classical arithmetic) of the lengths of the edges divided by the length of the path: $(a_{i_0i_1} + a_{i_1i_2} + \cdots + a_{i_{k-1}i_k})/k$. If $i_k = i_0$ then the path is a directed cycle and we refer to this quantity as the normalized length of the cycle. Recall that a directed graph is strongly connected if there is a directed path from any vertex to any other vertex.

Theorem 5.1.1. Let A be a tropical $n \times n$ -matrix whose graph G(A) is strongly connected. Then A has precisely one eigenvalue $\lambda(A)$. That eigenvalue equals the minimal normalized length of any directed cycle in G(A).

Proof. Let $\lambda = \lambda(A)$ be the minimum of the normalized lengths over all directed cycles in G(A). We first prove that $\lambda(A)$ is the only possibility for an eigenvalue. Suppose that $z \in \mathbb{R}^n$ is any eigenvector of A, and let γ be the corresponding eigenvalue. For any cycle $(i_1, i_2, \ldots, i_k, i_1)$ in G(A) we have

$$a_{i_1i_2} + z_{i_2} \ge \gamma + z_{i_1}, \ a_{i_2i_3} + z_{i_3} \ge \gamma + z_{i_2},$$

$$a_{i_3i_4} + z_{i_4} \ge \gamma + z_{i_3}, \dots, \ a_{i_ki_1} + z_{i_1} \ge \gamma + z_{i_k}.$$

Adding the left hand sides and the right hand sides, we find that the normalized length of the cycle is greater than or equal to γ . In particular, we have $\lambda(A) \geq \gamma$. For the reverse inequality, start with any index i_1 . Since z is an eigenvector with eigenvalue γ , there exists i_2 such that $a_{i_1i_2} + z_{i_2} = \gamma + z_{i_1}$. Likewise, there exists i_3 such that $a_{i_2i_3} + z_{i_3} = \gamma + z_{i_2}$. We continue in this manner until we reach an index i_l which was already in the sequence, say, $i_k = i_l$ for k < l. By adding the equations along this cycle, we find that

$$(a_{i_k,i_{k+1}} + z_{i_{k+1}}) + (a_{i_{k+1},i_{k+2}} + z_{i_{k+2}}) + \dots + (a_{i_{l-1},i_l} + z_{i_l})$$

= $(\gamma + z_{i_k}) + (\gamma + z_{i_{k+1}}) + \dots + (\gamma + z_{i_l}).$

We conclude that the normalized length of the cycle $(i_k, i_{k+1}, \ldots, i_l = i_k)$ in G(A) is equal to γ . In particular, $\gamma \ge \lambda(A)$. This proves that $\gamma = \lambda(A)$.

It remains to prove the existence of an eigenvector. Let B be the matrix obtained from A by (classically) subtracting λ from every entry in A. Then all cycles in the weighted graph G(B) have non-negative length, and there exists a cycle of length zero. Using tropical matrix operations we compute

$$B^* = B \oplus B^2 \oplus B^3 \oplus \dots \oplus B^n$$

The entry B_{ij}^* in row *i* and column *j* of the matrix B^* is the length of a shortest path from node *i* to node *j* in the directed graph G(B). Since the graph is strongly connected, we have $B_{ij}^* < \infty$. Moreover,

$$(\mathrm{Id} \oplus B) \odot B^* = B^* \tag{5.2}$$

Here $\mathrm{Id} = B^0$ is the tropical identity matrix whose diagonal entries are 0 and whose off-diagonal entries are ∞ . Fix any node *j* that lies on a zero length cycle of G(B), and let $x = B_{j}^*$ denote the *j*th column vector of the matrix B^* . We have $x_j = B_{jj}^* = 0$. This property together with (5.2) implies

$$x = (\mathrm{Id} \oplus B) \odot x = x \oplus B \odot x = B \odot x,$$

and we conclude that x is an eigenvector with eigenvalue λ of our matrix A:

$$A \odot x = (\lambda \odot B) \odot x = \lambda \odot (B \odot x) = \lambda \odot x.$$

This completes the proof of Theorem 5.1.1.

It appears that the computation of the eigenvalue λ of a tropical $n \times n$ matrix requires inspecting all cycles in G(A). However, this is not the case. There is an efficient algorithm, first proposed in [Kar78], for computing $\lambda(A)$ from the matrix $A = (a_{ij})$ based on linear programming. The idea is to set up the following linear program with n + 1 decision variables v_1, \ldots, v_n and λ :

Maximize
$$\gamma$$
 subject to $a_{ij} + v_j \ge \gamma + v_i$ for all $1 \le i, j \le n$. (5.3)

Proposition 5.1.2 (Karp 1978). The unique eigenvalue $\lambda(A)$ of the matrix $A = (a_{ij})$ coincides with the optimal value γ^* of the linear program (5.3).

Proof. The dual linear program to (5.3) takes the form

Minimize
$$\sum_{i=1}^{n} \sum_{i,j=1}^{n} a_{ij} x_{ij}$$
 subject to $x_{ij} \ge 0$ for $1 \le i, j \le n$,
 $\sum_{i,j=1}^{n} x_{ij} = 1$ and $\sum_{j=1}^{n} x_{ij} = \sum_{k=1}^{n} x_{ki}$ for all $1 \le i \le n$.

Here the x_{ij} are the decision variables, and the problem is to find a probability distribution (x_{ij}) on the edges of G(A) that represents a flow in the directed graph. The vertices of the polyhedron defined by these constraints are the uniform probability distributions on the directed cycles in G(A). Hence the objective function value of the dual linear program equals the minimum of the normalized lengths over all directed cycles in G(A). By strong duality, the primal linear program (5.3) has the same optimal value $\gamma^* = \lambda(A)$. \Box

We next determine the *eigenspace* of the matrix A, which is the set

$$\operatorname{Eig}(A) = \{ x \in \mathbb{R}^n : A \odot x = \lambda(A) \odot x \}.$$

Clearly, $\operatorname{Eig}(A)$ is closed under tropical scalar multiplication, that is, if $x \in \operatorname{Eig}(A)$ and $c \in \mathbb{R}$ then $c \odot x$ is also in $\operatorname{Eig}(A)$. We can thus identify $\operatorname{Eig}(A)$ with its image in $\mathbb{R}^n/\mathbb{R}(1, 1, \ldots, 1)$. We use the short-hand **1** for the vector $(1, \ldots, 1) \in \mathbb{R}^n$, and write $\mathbb{R}^n/\mathbf{1}$ for this quotient.

Every eigenvector of the matrix A is also an eigenvector of the matrix $B = (-\lambda(A)) \odot A$ and vice versa. Hence the eigenspace Eig(A) is equal to

$$\operatorname{Eig}(B) = \{ x \in \mathbb{R}^n : B \odot x = x \}.$$

Theorem 5.1.3. Let B_0^* be the submatrix of B^* given by the columns whose diagonal entry B_{jj}^* is zero. The image of this matrix (with respect to tropical multiplication of vectors on the right) is equal the desired eigenspace:

$$\operatorname{Eig}(A) = \operatorname{Eig}(B) = \operatorname{Image}(B_0^*).$$

Example 5.1.4. Before proving Theorem 5.1.3, we present three examples of 4×4 -matrices and their eigenspaces in $\mathbb{R}^4/\mathbf{1}$. Each point in $\mathbb{R}^4/\mathbf{1}$ is represented by a vector in \mathbb{R}^4 with last coordinate zero, and we here write

"Image" for the operator that computes the image in $\mathbb{R}^4/1$ of a matrix with four rows.

If
$$A = \begin{pmatrix} 3 & 1 & 4 & 5 \\ 5 & 2 & 4 & 2 \\ 4 & 1 & 6 & 3 \\ 2 & 6 & 3 & 6 \end{pmatrix}$$
 then $\lambda(A) = 5/3$ and $\operatorname{Eig}(A) = \operatorname{Image}\begin{pmatrix} -1/3 \\ 1/3 \\ -1/3 \\ 0 \end{pmatrix}$

If
$$A = \begin{pmatrix} 1 & 4 & 4 & 6 \\ 1 & 1 & 1 & 2 \\ 4 & 2 & 1 & 3 \\ 6 & 3 & 6 & 4 \end{pmatrix}$$
 then $\lambda(A) = 1$ and $\operatorname{Eig}(A) = \operatorname{Image}\begin{pmatrix} -2 & 1 & 1 \\ -2 & -2 & -2 \\ -1 & -1 & -2 \\ 0 & 0 & 0 \end{pmatrix}$
If $A = \begin{pmatrix} 4 & 5 & 3 & 3 \\ 3 & 5 & 4 & 6 \\ 6 & 1 & 5 & 3 \\ 5 & 5 & 2 & 5 \end{pmatrix}$ then $\lambda(A) = 9/4$ and $\operatorname{Eig}(A) = \operatorname{Image}\begin{pmatrix} 3/4 \\ 3/2 \\ 1/4 \\ 0 \end{pmatrix}$
If $A = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}$ then $\lambda(A) = 0$ and $\operatorname{Eig}(A) = \operatorname{Image}(A)$.

For this last example matrix, we have $A = B = B^* = B_0^*$, and the eigenspace $\operatorname{Eig}(A)$ is a certain 3-dimensional convex polytope in $\mathbb{R}^4/\mathbf{1}$ which is known as the *standard polytrope*. We will discuss such polyt(r)opes in Section 5.4. \Box

Proof of Theorem 5.1.3. We saw in the proof of Theorem 5.1.1 that every column vector x of B_0^* satisfies $B \odot x = x$. Since tropical linear combinations of such eigenvectors x are again eigenvectors, we have $\text{Image}(B_0^*) \subseteq \text{Eig}(B)$.

To prove the reverse inclusion, consider any vector $z \in \text{Eig}(B)$. Then $B^* \odot z = z$. Let \tilde{z} be the vector obtained from z by erasing all coordinates j such that $B_{ij}^* > 0$. We claim that $z = B_0^* \odot \tilde{z}$. This will show $z \in \text{Image}(B_0^*)$.

Consider any index $i \in \{1, \ldots, n\}$. We have $z_i = \min\{B_{ij}^* + z_j : j = 1, \ldots, n\}$. If $z_i = B_{ij}^* + z_j$ and $z_j = B_{jk}^* + z_k$ then $B_{ij}^* + B_{jk}^* + z_k = z_i \leq B_{ik}^* + z_k$, and the triangle inequality $B_{ij}^* + B_{jk}^* \geq B_{ik}^*$ implies that $z_i = B_{ik}^* + z_k$. Continuing in this manner, we eventually obtain the equality $z_i = B_{il}^* + z_l$ for some index l which lies in a cycle of length 0, that is, $B_{ll}^* = 0$. This equality can be rewritten as $z_i = ((B_0^*) \odot z)_i$, and the proof is complete. \Box

In classical linear algebra, the eigenvalues of a square matrix are the roots of its characteristic polynomial, and we seek to extend this to tropical linear algebra. The *characteristic polynomial* of our $n \times n$ -matrix A equals

$$f_A(t) = \det(A \oplus t \odot \operatorname{Id}),$$

where "det" denotes the tropical determinant. We have the following result.

Corollary 5.1.5. The eigenvalue $\lambda(A)$ of a tropical $n \times n$ -matrix A is the smallest root of its characteristic polynomial $f_A(t)$.

Proof. Consider the expansion of the characteristic polynomial:

 $f_A(t) = t^n \oplus c_1 \odot t^{n-1} \oplus c_2 \odot t^{n-2} \oplus \cdots \oplus c_{n-1} \odot t \oplus c_n.$

The coefficient c_i is the minimum over the lengths of all cycles on *i* nodes in G(A). The smallest root of the polynomial $f_A(t)$ equals

$$\min\{c_1, c_2/2, c_3/3, \ldots, c_n/n\}.$$

This minimum is the smallest normalized cycle length $\lambda(A)$.

Our discussion raises the question of how the tropical eigenvalue problem is related to the classical eigenvalue problem for a matrix over a field Kwith a valuation. Let M be an $n \times n$ -matrix with entries in K and let $A = \operatorname{val}(M)$ be its tropicalization. If the entries in M are general enough then the characteristic polynomial $f_A(t)$ of A coincides with the tropicalization of the classical characteristic polynomial of M. Assuming this to be the case, let us consider an arbitrary solution (μ, v) of the eigenvalue equation for M:

$$M \cdot v = \mu \cdot v$$

This equation does not tropicalize, i.e., there will be cancellations of lowest terms in the matrix-vector product $M \cdot v$, unless μ is an eigenvalue of minimal order $\lambda(A)$. Furthermore, the eigenvector v must satisfy the non-trivial combinatorial constraint imposed by Theorem 5.1.3, namely, the order of v must lie in the image of the matrix B_0^* . Here is an example to show this.

Example 5.1.6. Let n = 3, $K = \overline{\mathbb{Q}(\epsilon)}$, and consider the matrix

$$M = \begin{pmatrix} \epsilon & 1 & \epsilon \\ 1 & \epsilon & -\epsilon^2 \\ \epsilon & \epsilon^2 & \epsilon \end{pmatrix}$$

 \square

5.1. EIGENVALUES AND EIGENVECTORS

This matrix has three distinct eigenvalues μ in K, and we list each of them with a generator v for the corresponding one-dimensional eigenspace:

$$\begin{array}{ccc}
\mu & v \\
\epsilon & (\epsilon^2, -\epsilon, 1)^T \\
\sqrt{1 + \epsilon^2 - \epsilon^4} + \epsilon & (\epsilon - \sqrt{1 + \epsilon^2 - \epsilon^4}, \epsilon \sqrt{1 + \epsilon^2 - \epsilon^4} - 1, \epsilon(\epsilon^2 - 1))^T \\
-\sqrt{1 + \epsilon^2 - \epsilon^4} + \epsilon & (\epsilon^3 - \sqrt{1 + \epsilon^2 - \epsilon^4}, \epsilon^4 - 1, \epsilon^2 \sqrt{1 + \epsilon^2 - \epsilon^4} - \epsilon)^T
\end{array}$$

The tropicalization of the matrix M equals $A = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 2 \\ 1 & 2 & 1 \end{pmatrix}$, and the tropical

characteristic polynomial $f_A(z) = z^3 \oplus 1 \odot z^2 \oplus 0 \odot z \oplus 1$ factors as follows as a polynomial function:

$$f_A(z) = (z \oplus 0)^2 \odot (z \oplus 1).$$

This reflects the fact that M has two eigenvalues of order 0 and one eigenvalue of order 1. By Theorem 5.1.1, $\lambda(A) = 0$ is the only eigenvalue of the matrix A. The eigenspace Eig(A) is computed using Theorem 5.1.3. We have

$$B^* = A^* = A \oplus A^2 \oplus A^3 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}.$$

Hence $\operatorname{Eig}(A)$ is spanned by the column vector $(0, 0, 1)^T$. Equivalently, the eigenspace of A consists of the vectors $(a, a, a+1)^T$ for all $a \in \mathbb{R}$. Each of these arises as the order of an eigenvector of the classical matrix M. For instance, the last two eigenvectors v listed above both have order $(0, 0, 1)^T$. \Box

We close this section with the remark that the geometric interpretation of the determinant as a coplanarity criterion is the same in both classical and tropical geometry. An $n \times n$ -matrix is said to be *tropically singular* if it lies on the tropical hypersurface defined by the determinant, which is here regarded as polynomial of degree n in n^2 unknowns having n! terms.

Proposition 5.1.7. Let A be a real $n \times n$ -matrix. Then A is tropically singular if and only if the rows of A lie on a tropical hyperplane in $\mathbb{R}^n/\mathbf{1}$

Proof. Suppose that A is tropically singular. By the Fundamental Theorem for hypersurfaces, there exists a singular $n \times n$ -matrix U with entries in a field

K with valuation such that val(U) = A. Pick a non-zero vector in K^n that lies in the kernel of U and consider the classical hyperplane H perpendicular to that vector. Then the rows of A lie in the tropical hyperplane trop(H).

Conversely, suppose that the rows of A lie in a tropical hyperplane H. We wish to show that A is tropically singular. Both statements are invariant under tropically multiplying $A = (a_{ij})$ by a diagonal matrix on the left or on the right, so we may assume that A is non-negative and it has a zero in each row and each column of A. We may further assume that H is the hyperplane defined by the tropical linear form by $0 \odot x_1 \oplus \cdots \oplus 0 \odot x_n$. Then each row of $A = (a_{ij})$ actually contains two zero entries. Consider the bipartite graph on $[n] \times [n]$ with an edge (i, j) whenever $a_{ij} = 0$. This graph is connected and it has $\geq 2n$ edges. Hence it contains a cycle. A combinatorial argument shows that such a bipartite graph must contain a matching. This means that the tropical determinant of A is zero. Moreover, the bipartite graph must contain a cycle, and from this we conclude that A is tropically singular. \Box

5.2 Matroids

Matroid theory is a branch of discrete mathematics that is fundamental for tropical linear algebra. Matroids aim to characterize the combinatorial structure of dependence relations among vectors in a linear space over a field K. Tropical linear algebra can be viewed as a extension to the situation when the field K comes with a valuation. In matroid theory, one distinguishes between matroids and realizable matroids, and the extension here will be the distinction between tropical linear spaces and tropicalized linear spaces.

Definition 5.2.1. A tropicalized linear space over K is a the tropical variety of the form $\operatorname{trop}(X)$ where X is any linear subspace of K^n . The points in $\operatorname{trop}(X)$ are vectors $w \in \mathbb{R}^n$ such that $\operatorname{in}_w(f)$ is not a monomial for any f in the ideal $I_X \subset K[x_1, \ldots, x_n]$ of polynomials that vanish on X.

In this section we shall restrict ourselves to the constant coefficient case, and we take our field to be the complex numbers \mathbb{C} . Our aim is to explain the distinction between tropicalized linear spaces and tropical linear spaces when the field is \mathbb{C} . The general case will be discussed in Section 5.5.

Let X be a linear subspace of dimension d in \mathbb{C}^n . The corresponding ideal I_X in $\mathbb{C}[x_1, \ldots, x_n]$ is generated by the linear forms that vanish on X. A non-empty subset C of $\{1, \ldots, n\}$ is said to be *circuit* of X if $C = \{i : i \}$ x_i appears in ℓ for some linear form ℓ in the ideal I_X , and C is inclusionminimal with this property. Note that C uniquely determines the linear form ℓ up to scaling. The number of circuits of a d-dimensional linear subspace of K^n is at most $\binom{n}{d+1}$, and this bound is attained for generic subspaces X. Our first lemma says that the circuits determine the tropical variety trop(X).

Proposition 5.2.2. A vector $w \in \mathbb{R}^n$ lies in the tropicalized linear space trop(X) if and only if, for any circuit C of the subspace $X \subset \mathbb{C}^n$, the minimum of the numbers w_i , as i ranges over C, is attained at least twice.

Proof. The only-if direction is obvious because every circuit is the support of a linear form ℓ that lies in the ideal I_X . For the if direction suppose that w is not in trop(X). Compute the reduced Gröbner basis of I_X with respect to a term order that refines w. By Gaussian elimination, the elements of that reduced Gröbner basis are linear forms that are supported on circuits. Moreover, the initial ideal $in_w(I)$ is generated by the leading forms of these linear forms. Our hypothesis states that some monomial lies in $in_w(I)$. Since $in_w(I)$ is prime, this implies that some variable x_i lies in $in_w(I)$, and from this we conclude that x_i actually appears in the reduced Gröbner basis. \Box

There are many different but equivalent axiom systems for defining a matroid. One of them is the following axiom system for the circuits.

Definition 5.2.3. A matroid is a pair M = (E, C) where E is a finite set and C is a collection of subsets of E, called the *circuits* of M, that satisfies:

- (C1) No proper subset of a circuit is a circuit.
- (C2) If C_1, C_2 are circuits, $C_1 \neq C_2$, and $e \in C_1 \cap C_2$ then $(C_1 \cup C_2) \setminus \{e\}$ contains a circuit.

Clearly, the circuits of a linear subspace $X \subset \mathbb{C}^n$ satisfy (C1) and (C2). A matroid M that arises from such a subspace is said to be *realizable* over \mathbb{C} . Matroids provide a convenient language for linear algebra. Here are some basic definitions. An *independent set* of M is a subset of E that contains no circuit. A *basis* of M is a maximal independent set. All bases of M have the same cardinality d. That number is called the *rank* of M. A *flat* of a matroid M is a set F such that $\#(C \setminus F) \neq 1$ for any circuit C. The set of all flats is ordered by inclusion is a poset known as the *geometric lattice* of M. Each of these objects comes with its own axiom system for matroids. For example: **Definition 5.2.4.** A matroid is a pair $M = (E, \rho)$ where E is a finite set and ρ is a function $E \to \mathbb{N}$, called the *rank function* of M, that satisfies:

- (R1) $\rho(A) \leq |A|$ for all subsets A of E.
- (R2) If A and B are subset of E with $A \subseteq B$ then $\rho(A) \leq \rho(B)$.
- (R3) $\rho(A \cup B) + \rho(A \cap B) \le \rho(A) + \rho(B)$ for any two subsets A and B of E.

The rank of the matroid M is defined to be the rank of E, and we write $\rho(M) := \rho(E)$. Starting with the axiom system (R1)–(R3), the other descriptions of matroids are derived as follows. A subset A of E is independent if $\rho(A) = |A|$, and dependent otherwise. A basis is a maximal independent set, and a circuit is a minimal dependent set. A flat is a subset $A \subseteq E$ such that $\rho(A) < \rho(A \cup \{e\})$ for all $e \in E \setminus A$. In light of Proposition 5.2.2, it makes sense to associate a tropical linear space trop(M) with any matroid M.

Definition 5.2.5. Let M be a matroid on a finite set E, which we identify with $\{1, 2, ..., n\}$. The corresponding *tropical linear space* $\operatorname{trop}(M)$ is the set of all points $w \in \mathbb{R}^n$ such that, for any circuit C of M, the minimum of the numbers w_i is attained at least twice as i ranges over C. Since $\operatorname{trop}(M)$ is invariant under tropical scalar multiplication, we shall from now on regard it as a subset of the tropical projective space $\mathbb{TP}^{n-1} = \mathbb{R}^n / \mathbb{R}(1, 1, ..., 1)$.

We next describe a natural fan structure on the tropical linear space trop(M). Any flat F of the matroid M is represented by its incidence vector $e_F = \sum_{i \in F} e_i$. We regard e_F as an element in \mathbb{TP}^{n-1} . For any chain of proper flats $\sigma = \{ \emptyset \subset F_1 \subset \cdots \subset F_k \subset E \}$, we consider the closed polyhedral cone spanned by their incidence vectors:

 $\operatorname{pos}(\sigma) := \{\lambda_1 e_{F_1} + \lambda_2 e_{F_2} + \dots + \lambda_k e_{F_k} : \lambda_1, \lambda_2, \dots, \lambda_k \ge 0\}.$

Since $e_{F_1}, e_{F_2}, \ldots, e_{F_k}$ are linearly independent, $pos(\sigma)$ is a k-dimensional simplicial cone in \mathbb{TP}^{n-1} , i.e., it is the cone over a (k-1)-dimensional simplex.

Theorem 5.2.6. The collection of cones $pos(\sigma)$, where σ runs over all chains of flats of the matroid M, is a pure simplicial fan of dimension $\rho(M) - 1$ in \mathbb{R}^n . The support of this fan is precisely the tropical linear space trop(M).

Proof. We first show that $pos(\sigma) \subset trop(M)$ for any chain of proper flats $\sigma = \{F_1 \subset \cdots \subset F_k\}$. Let $w = \lambda_1 e_{F_1} + \cdots + \lambda_k e_{F_k}$ where $\lambda_1, \ldots, \lambda_k \ge 0$.

5.2. MATROIDS

Consider any circuit C of M, and let i be the largest index i such that the set $(F_i \cap C) \setminus F_{i-1}$ is non-empty. We claim that this set has at least two elements. If not then it is a singleton, and its unique element would be dependent on F_{i-1} , and this contradicts the hypothesis that F_{i-1} is a flat. Hence $(F_i \cap C) \setminus F_{i-1}$ has cardinality at least two. This means that the minimum of the coordinates w_i where $i \in C$ is attained at least twice. Since this holds for any circuit C, we conclude that $w \in \operatorname{trop}(M)$ as desired.

We next show that every vector $w \in \operatorname{trop}(M)$ lies in the relative interior of a cone $\operatorname{pos}(\sigma)$ where σ is a unique chain of flats of M. After tropical scalar multiplication, we may assume that $w \in \mathbb{R}^n$ is non-negative and its support is a proper subset of E. Then there exists a unique chain $F_1 \subset F_2 \subset$ $\cdots \subset F_k$ of proper subsets of E such that the image of w in \mathbb{TP}^{n-1} lies in the relative interior of $\operatorname{pos}(e_{F_1}, e_{F_2}, \ldots, e_{F_k})$. Equivalently, the function $j \mapsto w_j$ is constant on $F_i \setminus F_{i-1}$ and its value decreases as i increases.

We claim that each F_i is a flat. Suppose that F_i is not a flat. By the characterization of flats in terms of circuits, there exists a circuit C such that $C \setminus F_i = \{e\}$ is a singleton. We have $w_e = \min\{w_i : i \in C\}$, and that minimum is uniquely attained. This is a contradiction to our hypothesis that w lies in the tropical linear space trop(M). We conclude that the cones $pos(\sigma)$ where σ runs over all chains of proper flats is a simplicial fan in \mathbb{TP}^{n-1} .

Each chain of flats can be extended to a maximal chain, and each maximal chain of flats involves precisely $\rho(M) - 1$ proper flats. Hence the fan is a pure fan of dimension $\rho(M) - 1$, and the proof is complete.

We have shown that $\operatorname{trop}(M)$ has the structure of a fan over a simplicial complex Δ_M of dimension $\rho(M)-2$, and we sometimes identify $\operatorname{trop}(M)$ with Δ_M . The simplicial complex Δ_M is the order complex of the geometric lattice of M. It is known that the order complex Δ_M has excellent combinatorial and topological properties. For instance, Δ_M is shellable, and hence its homology is free abelian and concentrated in the top dimension. The rank of that top homology group is denoted $\mu(M)$ and is known as the *Möbius number* of the matroid. It coincides with the Euler characteristic of Δ_M , i.e. $\mu(M)$ is the absolute value of the alternating sum of the number of flats of rank i in M.

There is another fan structure on the tropical linear space $\operatorname{trop}(M)$, which is much coarser than the order complex, and which comes from the Gröbner fan of I_X when M is realized by a classical linear space X. That fan structure is known as the *Bergman fan*, and it can be defined as follows. Given $w \in \operatorname{trop}(M)$ consider the *initial matroid* M_w whose circuits are the (nonsingleton) sets $\{j \in C : w_j = \min_{i \in C}(w_i)\}$, where C runs over all circuits of M. The bases of M_w are those bases B of M whose weight $\sum_{i \in B} w_i$ is maximal. Two vectors w and w' lie in the same (relatively open) cone of the Bergman fan on trop(M) if and only if their initial matroids coincide, that is, $M_w = M_{w'}$. For more details see [AK06, FS05]. We now present some natural examples of matroids and we discuss their tropical linear spaces.

Example 5.2.7 (Uniform matroids). Suppose that X is a generic linear subspace of dimension d in \mathbb{C}^n . The corresponding matroid is the uniform matroid $M = U_{d,n}$, whose circuits are all the subsets of cardinality n-d+1 in $\{1, 2, \ldots, n\}$. The tropical linear space $\operatorname{trop}(U_{d,n})$ is the union of all orthants spanned by any d-1 of the basis vectors e_1, \ldots, e_n in \mathbb{TP}^{n-1} . The associated simplicial complex $\Delta_{U_{d,n}}$ is the (d-2)-skeleton of the (n-1)-simplex. \Box

Example 5.2.8 (Graphic matroids). Let G be a connected directed graph with d + 1 vertices and n edges. The associated d-dimensional linear space X has the parametric representation $x_{ij} = t_i - t_j$ for all directed edges (i, j). Its matroid M_G is called a graphic matroid. Its circuits are precisely the circuits of the graph. The associated tropical linear space trop(G) is the set of all edge weights on G such the minimum along each cycle is attained at least twice. An important special case arises when G the complete graph K_{d+1} . Here the elements of the tropical linear space trop (K_{d+1}) are the tuples $w \in \mathbb{R}^{\binom{d+1}{2}}$ such that $w_{ij} \geq \min\{w_{ik}, w_{jk}\}$ for all i, j, k. Up to a global sign change, these are precisely the ultrametrics on a set with d + 1 elements. \Box

Example 5.2.9 (The Fano plane). Let n = 7 and d = 3. The following seven triples are the circuits of a rank 3 matroid M on the set $\{1, 2, ..., 7\}$:

$$124, 235, 346, 457, 457, 561, 672, 713.$$
 (5.4)

Here Δ_M is a one-dimensional simplicial complex, namely, it is a bipartite graph with 14 vertices and 21 edges. The vertices are the seven points $i \in$ $\{1, 2, \ldots, 7\}$ and the seven triples in (5.4). There is an edge from i to each triple that contains it. This matroid cannot be realized in characteristic zero, so trop(M) is a tropical linear space that is not a *tropicalized linear space*, by which we mean the tropicalization of any classical linear space over \mathbb{C} . \Box

5.3 Tropical Convexity

We now introduce the notions of convexity and convex polytopes in the setting of tropical geometry. Combinatorial types of tropical polytopes are shown to be in bijection with regular triangulations of products of two simplices. This section is based on the article [DSS05]. We note that convexity over arbitrary idempotent semirings, including the min-plus algebra had been introduced considerably earlier in various contexts of applied mathematics, notably in the works of Cohen, Gaubert and Quadrat [CGQ04] and Litvinov, Maslov and Shpiz [LMS01].

A subset S of \mathbb{R}^n is called *tropically convex* if the set S contains the point $a \odot x \oplus b \odot y$ for all $x, y \in S$ and all $a, b \in \mathbb{R}$. The *tropical convex hull* of a given subset $V \subset \mathbb{R}^n$ is the smallest tropically convex subset of \mathbb{R}^n which contains V. We shall see in Proposition 5.3.5 that the tropical convex hull of V coincides with the set of all tropical linear combinations

 $a_1 \odot v_1 \oplus \cdots \oplus a_r \odot v_r$, where $v_1, \ldots, v_r \in V$ and $a_1, \ldots, a_r \in \mathbb{R}$. (5.5)

Any tropically convex subset S of \mathbb{R}^n is closed under tropical scalar multiplication, $\mathbb{R} \odot S \subseteq S$. In other words, if $x \in S$ then $x + \lambda(1, \ldots, 1) \in S$ for all $\lambda \in \mathbb{R}$. We will therefore identify the tropically convex set S with its image in the (n-1)-dimensional tropical projective space \mathbb{TP}^{n-1} . A tropical polytope is the tropical convex hull of a finite subset V in \mathbb{TP}^{n-1} . In Theorem 5.1.3 we have seen that $\operatorname{Eig}(A)$ is a tropical polytope for any $n \times n$ -matrix A:

Remark 5.3.1. The eigenspace of a square matrix is a tropical polytope.

We shall see that every tropical polytope is a finite union of convex polytopes in the usual sense: the tropical convex hull of $V = \{v_1, \ldots, v_r\} \subset \mathbb{R}^n$ has a natural polyhedral cell decomposition, called the *tropical complex* generated by V. One of our goals in Section 5.3 is to prove the following result:

Theorem 5.3.2. The combinatorial types of tropical complexes generated by configurations of r points in \mathbb{TP}^{n-1} are in natural bijection with the regular polyhedral subdivisions of the product of two simplices $\Delta_{n-1} \times \Delta_{r-1}$.

This implies a remarkable duality between tropical (n-1)-polytopes with r vertices and tropical (r-1)-polytopes with n vertices. Another consequence of Theorem 5.3.2 is a formula for the f-vector of a generic tropical complex.

Figure 5.1: Tropical convex sets and tropical line segments in \mathbb{TP}^2 .

We begin with pictures of tropical convex sets in the tropical plane \mathbb{TP}^2 . A point $(x_1, x_2, x_3) \in \mathbb{TP}^2$ is represented by drawing the point with coordinates $(x_2 - x_1, x_3 - x_1)$ in the plane of the paper. The triangle on the left hand side in Figure 5.1 is tropically convex, but it is not a tropical polytope because it is not the tropical convex hull of finitely many points. The thick edges indicate two tropical line segments. The picture on the right hand side is a *tropical triangle*, namely, it is the tropical convex hull of the three points (0,0,1), (0,2,0) and (0,-1,-2) in the tropical plane \mathbb{TP}^2 . The thick edges represent the tropical segments connecting any two of these three points.

Tropical convex sets enjoy many of the features of ordinary convex sets:

Theorem 5.3.3. The intersection of two tropically convex sets in \mathbb{R}^n or in \mathbb{TP}^{n-1} is tropically convex. The projection of a tropically convex set onto a coordinate hyperplane is tropically convex. The ordinary hyperplane $\{x_i - x_j = l\}$ is tropically convex, and the projection map from this hyperplane to \mathbb{R}^{n-1} given by eliminating x_i is an isomorphism of tropical semimodules. Tropically convex sets are contractible spaces. The Cartesian product of two tropically convex sets is tropically convex.

Proof. We prove the statements in the order given. If S and T are tropically convex, then for any two points $x, y \in S \cap T$, both S and T contain the tropical line segment between x and y, and consequently so does $S \cap T$. Therefore $S \cap T$ is tropically convex by definition.

Suppose S is a tropically convex set in \mathbb{R}^n . We claim that the image of S under the coordinate projection $\phi : \mathbb{R}^n \to \mathbb{R}^{n-1}, (x_1, x_2, \dots, x_n) \mapsto$ (x_2, \dots, x_n) is a tropically convex subset of \mathbb{R}^{n-1} . If $x, y \in S$ then we have

$$\phi(c \odot x \oplus d \odot y) = c \odot \phi(x) \oplus d \odot \phi(y).$$

This means that ϕ is a homomorphism of tropical semimodules. Therefore, if S contains the tropical line segment between x and y, then $\phi(S)$ contains the tropical line segment between $\phi(x)$ and $\phi(y)$ and hence is tropically convex. The same holds for the induced map $\phi : \mathbb{TP}^{n-1} \to \mathbb{TP}^{n-2}$.

Most ordinary hyperplanes in \mathbb{R}^n are not tropically convex, but we are claiming that hyperplanes of the special form $x_i - x_j = k$ are tropically convex. If x and y lie in that hyperplane then $x_i - y_i = x_j - y_j$. This last

5.3. TROPICAL CONVEXITY

equation implies the following identity for any real numbers $c, d \in \mathbb{R}$: $(c \odot x \oplus d \odot y)_i - (c \odot x \oplus d \odot y)_j = \min(x_i + c, y_i + d) - \min(x_j + c, y_j + d) = k.$ Thus the tropical segment between x and y is in the hyperplane $\{x_i - x_j = k\}$.

Consider the map from $\{x_i - x_j = k\}$ to \mathbb{R}^{n-1} given by deleting the *i*-th coordinate. This map is injective: if two points differ in the x_i coordinate they must also differ in the x_j coordinate. It is surjective because we can recover the *i*-th coordinate by setting $x_i = x_j + k$. Hence this map is an isomorphism of \mathbb{R} -vector spaces and it is also of $(\mathbb{R}, \oplus, \odot)$ -semimodules.

Let S be a tropically convex set in \mathbb{R}^n or \mathbb{TP}^{n-1} . Consider the family of hyperplanes $H_l = \{x_1 - x_2 = l\}$ for $l \in \mathbb{R}$. We know that the intersection $S \cap H_l$ is tropically convex, and isomorphic to its (convex) image under the map deleting the first coordinate. This image is contractible by induction on the dimension n of the ambient space. Therefore, $S \cap H_l$ is contractible. The result then follows from the topological result that if S is connected, which all tropically convex sets obviously are, and if $S \cap H_l$ is contractible for each l, then S itself is also contractible.

Suppose that $S \subset \mathbb{R}^n$ and $T \subset \mathbb{R}^m$ are tropically convex. Our last assertion states that $S \times T$ is a tropically convex subset of \mathbb{R}^{n+m} . Take any (x, y) and (x', y') in $S \times T$ and $c, d \in \mathbb{R}$. Then

$$c \odot (x, y) \oplus d \odot (x', y') = (c \odot x \oplus d \odot x', c \odot y \oplus d \odot y')$$

lies in $S \times T$ since S and T are tropically convex.

We next give a more precise description of tropical line segments.

Proposition 5.3.4. The tropical line segment between two points x and y in \mathbb{TP}^{n-1} is the concatenation of at most n-1 ordinary line segments. The slope of each line segment is a zero-one vector.

Proof. After relabeling coordinates of $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$, we may assume $y_1 - x_1 \leq y_2 - x_2 \leq \cdots \leq y_n - x_n$. The following points lie in the given order on the tropical segment between x and y:

Figure 5.2: Tropical polytopes: the first two live in \mathbb{TP}^2 , the last in \mathbb{TP}^3 .

Between any two consecutive points, the tropical line segment agrees with the ordinary line segment, which has slope $(0, 0, \ldots, 0, 1, 1, \ldots, 1)$. Hence the tropical line segment between x and y is the concatenation of at most n-1 ordinary line segments, one for each strict inequality $y_i - x_i < y_{i+1} - x_{i+1}$. \Box

Proposition 5.3.4 shows an important feature of tropical convexity: segments use a limited set of slopes. We next characterize tropical convex hulls.

Proposition 5.3.5. The smallest tropically convex subset of \mathbb{TP}^{n-1} which contains a given set V coincides with the set of all tropical linear combinations (5.5). We denote this set by $\operatorname{tconv}(V)$.

Proof. Let $x = \bigoplus_{i=1}^{r} a_i \odot v_i$ be the point in (5.5). If $r \leq 2$ then x is clearly in the tropical convex hull of V. If r > 2 then we write $x = a_1 \odot v_1 \oplus$ $(\bigoplus_{i=2}^{r} a_i \odot v_i)$. The parenthesized vector lies the tropical convex hull, by induction on r, and hence so does x. For the converse, consider any two tropical linear combinations $x = \bigoplus_{i=1}^{r} c_i \odot v_i$ and $y = \bigoplus_{j=1}^{r} d_i \odot v_i$. By the distributive law, $a \odot x \oplus b \odot y$ is also a tropical linear combination of $v_1, \ldots, v_r \in V$. Hence the set of all tropical linear combinations of V is tropically convex, so it contains the tropical convex hull of V.

If V is a finite subset of \mathbb{TP}^{n-1} then $\operatorname{tconv}(V)$ is a tropical polytope. In Figure 5.2 we see three small examples of tropical polytopes. The first and second are tropical convex hulls of three points in \mathbb{TP}^2 . The third tropical polytope lies in \mathbb{TP}^3 and is the union of three squares. One of the basic results in the usual theory of convex polytopes is Carathéodory's theorem. This important theorem holds also in the tropical setting.

Proposition 5.3.6 (Tropical Carathéodory's Theorem). If x is in the tropical convex hull of a set of r points v_1, \ldots, v_r in \mathbb{TP}^{n-1} , with r > n, then x is in the tropical convex hull of at most n of them.

Proof. Let $x = \bigoplus_{i=1}^{r} a_i \odot v_i$ and suppose r > n. For each coordinate $j \in \{1, \ldots, n\}$, there exists an index $i \in \{1, \ldots, r\}$ such that $x_j = c_i + v_{ij}$. Take a subset I of $\{1, \ldots, r\}$ composed of one such i for each j. Then we also have $x = \bigoplus_{i \in I} a_i \odot v_i$, where $\#(I) \le n$.

5.3. TROPICAL CONVEXITY

Recall that the *tropical hyperplane* defined by a tropical linear form $a_1 \odot x_1 \oplus a_2 \odot x_2 \oplus \cdots \oplus a_n \odot x_n$ consists of all points $x = (x_1, x_2, \ldots, x_n)$ in \mathbb{TP}^{n-1} such that the following holds (in ordinary arithmetic):

 $a_i + x_i = a_j + x_j = \min\{a_k + x_k : k = 1, \dots, n\}$ for some indices $i \neq j$. (5.6)

Just like in ordinary geometry, hyperplanes are convex sets:

Proposition 5.3.7. Tropical hyperplanes in \mathbb{TP}^{n-1} are tropically convex.

Proof. Let H be the hyperplane defined by (5.6). Let x and y be in H and consider any linear combination $z = c \odot x \oplus d \odot y$. Let i be an index which minimizes $a_i + z_i$. We need to show that this minimum is attained twice. By definition, z_i is equal to either $c + x_i$ or $d + y_i$, and, after permuting x and y, we may assume $z_i = c + x_i \leq d + y_i$. Since, for all k, $a_i + z_i \leq a_k + z_k$ and $z_k \leq c + x_k$, it follows that $a_i + x_i \leq a_k + x_k$ for all k, so that $a_i + x_i$ achieves the minimum of $\{a_1 + x_1, \ldots, a_n + x_n\}$. Since x is in H, there exists some index $j \neq i$ for which $a_i + x_i = a_j + x_j$. But now $a_j + z_j \leq a_j + c + x_j = c + a_i + x_i = a_i + z_i$. Since $a_i + z_i$ is the minimum of all $a_j + z_j$, the two must be equal, and this minimum is obtained at least twice as desired.

Proposition 5.3.7 implies that if V is a subset of \mathbb{TP}^{n-1} which happens to lie in a tropical hyperplane H, then its tropical convex hull $\operatorname{tconv}(V)$ will lie in H as well. The same holds for tropical linear spaces of higher codimension because these are always finite intersections of tropical hyperplanes:

Corollary 5.3.8. Tropical linear spaces in \mathbb{TP}^{n-1} are tropically convex.

We now concentrate on the combinatorial structure of tropical polytopes. Let $V = \{v_1, v_2, \ldots, v_r\}$ be a fixed finite subset of tropical projective space \mathbb{TP}^{n-1} . Here $v_i = (v_{i1}, v_{i2}, \ldots, v_{in})$. Our objective is to study the tropical polytope $P = \operatorname{tconv}(V)$. We begin by describing the natural cell decomposition of \mathbb{TP}^{n-1} induced by the fixed finite subset V. Consider any point x in \mathbb{TP}^{n-1} . The type of x relative to V is the ordered n-tuple (S_1, \ldots, S_n) of subsets $S_j \subseteq \{1, 2, \ldots, r\}$ which is defined as follows: An index i is in S_j if

$$v_{ij} - x_j = \min(v_{i1} - x_1, v_{i2} - x_2, \dots, v_{in} - x_n).$$

Equivalently, if we set $\lambda_i = \min\{\lambda \in \mathbb{R} : \lambda \odot v_i \oplus x = x\}$ then S_j is the set of all indices i such that $\lambda_i \odot v_i$ and x have the same j-th coordinate. We say that an n-tuple of indices $S = (S_1, \ldots, S_n)$ is a type if it arises in this manner. Note that every i must be in some S_j .

Example 5.3.9. Let r = n = 3, $v_1 = (0, 0, 2)$, $v_2 = (0, 2, 0)$ and $v_3 = (0, 1, -2)$. There are 30 possible types as x ranges over the plane \mathbb{TP}^2 . The corresponding cell decomposition has six convex regions (one bounded, five unbounded), 15 edges (6 bounded, 9 unbounded) and 6 vertices. For instance, the point x = (0, 1, -1) has type $(x) = (\{2\}, \{1\}, \{3\})$ and its cell is a pentagon. The point x' = (0, 0, 0) has type $(x') = (\{1, 2\}, \{1\}, \{2, 3\})$ and its cell is one of the six vertices. The point x'' = (0, 0, -3) has type $(x'') = \{\{1, 2, 3\}, \{1\}, \emptyset\}$ and its cell is an unbounded edge.

Our first application of types is the following separation theorem.

Proposition 5.3.10 (Tropical Farkas Lemma). For all $x \in \mathbb{TP}^{n-1}$, exactly one of the following is true:

(i) the point x is in the tropical polytope P = tconv(V), or

(ii) there exists a tropical hyperplane which separates x from P.

This means: if the hyperplane is given by (5.6) and $a_k + x_k = \min(a_1 + x_1, \ldots, a_n + x_n)$ then $a_k + y_k > \min(a_1 + y_1, \ldots, a_n + y_n)$ for all $y \in P$.

Proof. Consider any point $x \in \mathbb{TP}^{n-1}$, with $\text{type}(x) = (S_1, \ldots, S_n)$, and let $\lambda_i = \min\{\lambda \in \mathbb{R} : \lambda \odot v_i \oplus x = x\}$ as before. We define

$$\pi_V(x) = \lambda_1 \odot v_1 \oplus \lambda_2 \odot v_2 \oplus \cdots \oplus \lambda_r \odot v_r.$$
 (5.7)

There are two cases: either $\pi_V(x) = x$ or $\pi_V(x) \neq x$. The first case implies (i). Since (i) and (ii) clearly cannot occur at the same time, it suffices to prove that the second case implies (ii). Suppose that $\pi_V(x) \neq x$. Then S_k is empty for some index $k \in \{1, \ldots, n\}$. This means that $v_{ik} + \lambda_i - x_k > 0$ for $i = 1, 2, \ldots, r$. Let $\varepsilon > 0$ be smaller than any of these r positive reals. We now choose our separating tropical hyperplane (5.6) as follows:

$$a_k := -x_k - \varepsilon$$
 and $a_j := -x_j$ for $j \in \{1, \dots, n\} \setminus \{k\}.$ (5.8)

This certainly satisfies $a_k + x_k = \min(a_1 + x_1, \dots, a_n + x_n)$. Now, consider any point $y = \bigoplus_{i=1}^r c_i \odot v_i$ in tconv(V). Pick any m such that $y_k = c_m + v_{mk}$. By definition of the λ_i , we have $x_k \leq \lambda_m + v_{mk}$ for all k, and there exists some j with $x_j = \lambda_m + v_{mj}$. These equations and inequalities imply

$$a_k + y_k = a_k + c_m + v_{mk} = c_m + v_{mk} - x_k - \varepsilon > c_m - \lambda_m$$

= $c_m + v_{mj} - x_j \ge y_j - x_j = a_j + y_j \ge \min(a_1 + y_1, \dots, a_n + y_n)$

Therefore, the hyperplane defined by (5.8) separates x from P as desired. \Box
Figure 5.3: The region $X_{(2,1,3)}$ in the tropical convex hull of v_1 , v_2 and v_3 .

The construction in (5.7) defines a map $\pi_V : \mathbb{TP}^{n-1} \to P$ whose restriction to P is the identity. This map is the tropical version of the *nearest point* map onto a closed convex set. If $S = (S_1, \ldots, S_n)$ and $T = (T_1, \ldots, T_n)$ are *n*-tuples of subsets of $\{1, 2, \ldots, r\}$, then we write $S \subseteq T$ if $S_j \subseteq T_j$ for $j = 1, \ldots, n$. We also consider the set of all points whose type contains S:

$$X_S := \left\{ x \in \mathbb{TP}^{n-1} : S \subseteq \operatorname{type}(x) \right\}.$$

Lemma 5.3.11. The set X_S is a closed convex polyhedron. More precisely,

$$X_{S} = \{ x \in \mathbb{TP}^{n-1} : x_{k} - x_{j} \le v_{ik} - v_{ij} \text{ for } 1 \le j, k \le n \text{ and } i \in S_{j} \}.$$
(5.9)

Proof. Let $x \in \mathbb{TP}^{n-1}$ and T = type(x). First, suppose x is in X_S . Then $S \subseteq T$. For every i, j, k such that $i \in S_j$, we also have $i \in T_j$, and so by definition we have $v_{ij} - x_j \leq v_{ik} - x_k$, or $x_k - x_j \leq v_{ik} - v_{ij}$. Hence x lies in the set on the right hand side of (5.9). For the proof of the reverse inclusion, suppose that x lies in the right hand side of (5.9). Then, for all i, j with $i \in S_j$, and for all k, we have $v_{ij} - x_j \leq v_{ik} - x_k$. This means that $v_{ij} - x_j = \min(v_{i1} - x_1, \ldots, v_{in} - x_n)$ and hence $i \in T_j$. Consequently, for all j, we have $S_j \subset T_j$, and so $x \in X_S$.

As an example for Lemma 5.3.11, we consider the region $X_{(2,1,3)}$ in the tropical convex hull of $v_1 = (0,0,2)$, $v_2 = (0,2,0)$, and $v_3 = (0,1,-2)$. This region is defined by six linear inequalities, one of which is redundant, as depicted in Figure 5.3. Lemma 5.3.11 has the following immediate corollaries.

Corollary 5.3.12. The intersection $X_S \cap X_T$ equals the polyhedron $X_{S \cup T}$.

Proof. The inequalities defining $X_{S\cup T}$ are the inequalities defining X_S and X_T , and points satisfying these inequalities are precisely those in $X_S \cap X_T$. \Box

Corollary 5.3.13. The polyhedron X_S is bounded if and only if $S_j \neq \emptyset$ for all j = 1, 2, ..., n.

Proof. Suppose $S_j \neq \emptyset$ for all j = 1, ..., n. Then for every j and k, we can find $i \in S_j$ and $m \in S_k$, which via Lemma 5.3.11 yield the inequalities $v_{mk} - v_{mj} \leq x_k - x_j \leq v_{ik} - v_{ij}$. This implies that each $x_k - x_j$ is bounded on X_S , which means that X_S is a bounded subset of \mathbb{TP}^{n-1} . Conversely,

suppose some S_j is empty. Then the only inequalities involving x_j are of the form $x_j - x_k \leq c_{jk}$. Consequently, if x is in S_j , so is $x - ke_j$ for k > 0, where e_j is the *j*-th basis vector. Therefore, in this case, X_S is unbounded.

Corollary 5.3.14. Suppose $S = (S_1, \ldots, S_n)$ with $S_1 \cup \cdots \cup S_n = \{1, \ldots, r\}$. If $S \subseteq T$ then X_T is a face of X_S , and all faces of X_S are of this form.

Proof. For the first part, it suffices to prove that the statement is true when T covers S in the poset of containment, i.e. when $T_j = S_j \cup \{i\}$ for some $j \in \{1, \ldots, n\}$ and $i \notin S_j$, and $T_k = S_k$ for $k \neq j$. We have the inequality presentation of X_S given by Lemma 5.3.11. The inequality presentation of X_T consists of the inequalities defining X_S together with the inequalities

$$\{x_k - x_j \le v_{ik} - v_{ij} : k \in \{1, \dots, n\}\}.$$
(5.10)

By assumption, i is in some S_m . We claim that X_T is the face of S given by

$$x_m - x_j = v_{im} - v_{ij}. (5.11)$$

Since X_S satisfies the inequality $x_j - x_m \leq v_{ij} - v_{im}$, (5.11) defines a face F of S. The inequality $x_m - x_j \leq v_{im} - v_{ij}$ is in (5.10), so (5.11) is valid on X_T and $X_T \subseteq F$. However, any point in F, being in X_S , satisfies $x_k - x_m \leq v_{ik} - v_{im}$ for $1 \leq k \leq n$. Adding (5.11) to these inequalities proves that the inequalities (5.10) are valid on F, and hence $F \subseteq X_T$. So $X_T = F$ as desired.

By the discussion in the proof of the first part, prescribing equality in the facet-defining inequality $x_k - x_j \leq v_{ik} - v_{ij}$ yields X_T , where $T_k = S_k \cup \{i\}$ and $T_j = S_j$ for $j \neq k$. Therefore, all facets of X_S can be obtained as regions X_T , and it follows recursively that all faces of X_S are of this form. \Box

Corollary 5.3.15. Let $S = (S_1, \ldots, S_n)$ be an n-tuple of indices satisfying $S_1 \cup \cdots \cup S_n = \{1, \ldots, r\}$. Then X_S is equal to X_T for some type T.

Proof. Let x be a point in the relative interior of X_S , and let T = type(x). Since $x \in X_S$, T contains S, and by Lemma 5.3.14, X_T is a face of X_S . However, since x is in the relative interior of X_S , the only face of X_S containing x is X_S itself, so we must have $X_S = X_T$ as desired.

Theorem 5.3.16. The collection of convex polyhedra X_S , where S ranges over all types, is a cell decomposition C_V of \mathbb{TP}^{n-1} . The tropical polytope $P = \operatorname{tconv}(V)$ equals the union of all bounded cells X_S in this decomposition. Proof. Since each point has a type, it is clear that the union of the X_S is equal to \mathbb{TP}^{n-1} . By Corollary 5.3.14, the faces of X_S are equal to X_U for $S \subseteq U$, and by Corollary 5.3.15, $X_U = X_W$ for some type W, and hence X_U is in our collection. The only thing remaining to check to show that this collection defines a cell decomposition is that $X_S \cap X_T$ is a face of both X_S and X_T , but $X_S \cap X_T = X_{S \cup T}$ by Corollary 5.3.12, and $X_{S \cup T}$ is a face of X_S and X_T by Corollary 5.3.14.

For the second assertion consider any point $x \in \mathbb{TP}^{n-1}$ and let S = type(x). We have seen in the proof of the Tropical Farkas Lemma (Proposition 5.3.10) that x lies in P if and only if no S_j is empty. By Corollary 5.3.13, this is equivalent to the polyhedron X_S being bounded.

The set of bounded cells X_S is referred to as the tropical complex generated by V. Theorem 5.3.16 states that this provides a polyhedral decomposition of the polytope P = tconv(V). Different sets V may have the same tropical polytope as their convex hull, but generate different tropical complexes; the decomposition of a tropical polytope depends on the chosen V.

Here is a nice geometric construction of the cell decomposition \mathcal{C}_V of \mathbb{TP}^{n-1} induced by $V = \{v_1, \ldots, v_r\}$. Let \mathcal{F} be the fan in \mathbb{TP}^{n-1} defined by the tropical hyperplane (5.6) with $a_1 = \cdots = a_n = 0$. Two vectors x and y lie in the same relatively open cone of the fan \mathcal{F} if and only if

$$\{ j : x_j = \min(x_1, \dots, x_n) \} \\= \{ j : y_j = \min(y_1, \dots, y_n) \}.$$

If we translate the negative of \mathcal{F} by the vector v_i then we get a new fan which we denote by $v_i - \mathcal{F}$. Two vectors x and y lie in the same relatively open cone of the fan $v_i - \mathcal{F}$ if and only

$$\{j: x_j - v_{ij} = \max(x_1 - v_{i1}, \dots, x_n - v_{in})\}\$$

=
$$\{j: y_j - v_{ij} = \max(y_1 - v_{i1}, \dots, y_n - v_{in})\}.$$

Proposition 5.3.17. The cell decomposition C_V is the common refinement of the r fans $v_i - \mathcal{F}$.

Proof. We need to show that the cells of this common refinement are precisely the convex polyhedra X_S . Take a point x, with T = type(x) and define $S_x = (S_{x1}, \ldots, S_{xn})$ by letting $i \in S_{xj}$ whenever

$$x_j - v_{ij} = \max(x_1 - v_{i1}, \dots, x_n - v_{in}).$$
 (5.12)

Figure 5.4: Tropical complex expressed as the bounded cells in the refinement of the fans $v_1 - \mathcal{F}$, $v_2 - \mathcal{F}$ and $v_3 - \mathcal{F}$. Cells are labeled with their types.

Two points x and y are in the relative interior of the same cell of the common refinement if and only if they are in the same relatively open cone of each fan; this is tantamount to saying that $S_x = S_y$. However, we claim that $S_x = T$. Indeed, taking the negative of both sides of (5.12) yields exactly the condition for i being in T_j , by the definition of type. Consequently, the condition for two points having the same type is the same as the condition for the two points being in the relative interior of the same chamber of the common refinement of the fans $v_1 - \mathcal{F}, v_2 - \mathcal{F}, \ldots, v_r - \mathcal{F}$.

An instance of this construction is shown for our running example, where $v_1 = (0, 0, 2), v_2 = (0, 2, 0)$, and $v_3 = (0, 1, -2)$, in Figure 5.4.

The next few results provide additional information about the polyhedron X_S . Let G_S denote the undirected graph with vertices $\{1, \ldots, n\}$, where $\{j, k\}$ is an edge if and only if $S_j \cap S_k \neq \emptyset$.

Proposition 5.3.18. The dimension d of the polyhedron X_S is one less than the number of connected components of G_S , and X_S is affinely and tropically isomorphic to some polyhedron X_T in \mathbb{TP}^d .

Proof. The proof is by induction on n. Suppose we have $i \in S_j \cap S_k$. Then X_S satisfies the linear equation $x_k - x_j = c$ where $c = v_{ik} - v_{ij}$. Eliminating the variable x_k (projecting onto \mathbb{TP}^{n-2}), X_S is affinely and tropically isomorphic to X_T where the type T is defined by $T_r = S_r$ for $r \neq j$ and $T_j = S_j \cup S_k$. The region X_T exists in the cell decomposition of \mathbb{TP}^{n-2} induced by the vectors w_1, \ldots, w_n with $w_{ir} = v_{ir}$ for $r \neq j$, and $w_{ij} = \max(v_{ij}, v_{ik} - c)$. The graph G_T is obtained from the graph G_S by contracting the edge $\{j, k\}$, and thus has the same number of connected components.

This induction on n reduces us to the case where all of the S_j are pairwise disjoint. We must show that X_S has dimension n-1. Suppose not. Then X_S lies in \mathbb{TP}^{n-1} but has dimension less than n-1. Therefore, one of the inequalities in (5.9) holds with equality, say $x_k - x_j = v_{ik} - v_{ij}$ for all $x \in X_S$. The inequality " \leq " implies $i \in S_j$ and the inequality " \geq " implies $i \in S_k$. Hence S_j and S_k are not disjoint, a contradiction.

The following proposition can be regarded as a converse to Lemma 5.3.11.

Proposition 5.3.19. Let R be any polytope in \mathbb{TP}^{n-1} defined by inequalities of the form $x_k - x_j \leq c_{jk}$. Then R arises as a cell X_S in the decomposition \mathcal{C}_V of \mathbb{TP}^{n-1} defined by some set $V = \{v_1, \ldots, v_n\}$.

Proof. Define the vectors v_i to have coordinates $v_{ij} = c_{ij}$ for $i \neq j$, and $v_{ii} = 0$. (If c_{ij} did not appear in the given inequality presentation then simply take it to be a very large positive number.) Then by Lemma 5.3.11, the polytope in \mathbb{TP}^{n-1} defined by the inequalities $x_k - x_j \leq c_{jk}$ is precisely the unique cell of type $(1, 2, \ldots, n)$ in the tropical conver hull of $\{v_1, \ldots, v_n\}$. \Box

The region X_S is a polytope both in the ordinary sense and in the tropical sense. Such polytopes have been called *polytropes* in [JK08].

Proposition 5.3.20. Every bounded cell X_S in the tropical complex generated by V is itself a tropical polytope, equal to the tropical convex hull of its vertices. The number of vertices of the polytrope X_S is at most $\binom{2n-2}{n-1}$, and this bound is tight for all positive integers n.

Proof. By Proposition 5.3.18, if X_S has dimension d, it is affinely and tropically isomorphic to a region in the convex hull of a set of points in \mathbb{TP}^d , so it suffices to consider the full-dimensional case.

The inequality presentation of Lemma 5.3.11 demonstrates that X_S is tropically convex for all S, since if two points satisfy an inequality of that form, so does any tropical linear combination thereof. Therefore, it suffices to show that X_S is contained in the tropical convex hull of its vertices.

The proof is by induction on the dimension of X_S . All proper faces of X_S are polytopes X_T of lower dimension, and by induction are contained in the tropical convex hull of their vertices. These vertices are a subset of the vertices of X_S , and so this face is in the tropical convex hull. Take any point $x = (x_1, \ldots, x_n)$ in the interior of X_S . We can travel in any direction from x while remaining in X_S . Let us travel in the $(1, 0, \ldots, 0)$ direction until we hit the boundary, to obtain points $y_1 = (x_1 + b, x_2, \ldots, x_n)$ and $y_2 =$ $(x_1 - c, x_2, \ldots, x_n)$ in the boundary of X_S . These points are in the tropical convex hull by the induction hypothesis, which means that $x = y_1 \oplus c \odot y_2$ is also, completing the proof of the first assertion.

For the second assertion, we consider the convex hull of all differences of unit vectors, $e_i - e_j$. This is a lattice polytope of dimension n - 1 and normalized volume $\binom{2n-2}{n-1}$. To see this, we observe that this polytope is tiled by n copies of the convex hull of the origin and the $\binom{n}{2}$ vectors $e_i - e_j$ with i < j. The other n - 1 copies are gotten by cyclic permutation of the coordinates. This latter polytope was studied by Gel'fand, Graev and Postnikov, who showed in [GGP97, Theorem 2.3 (2)] that the normalized volume of this polytope equals the Catalan number $\frac{1}{n} \binom{2n-2}{n-1}$. We conclude that every complete fan whose rays are among the vectors $e_i - e_j$ has at most $\binom{2n-2}{n-1}$ maximal cones. This applies in particular to the normal fan of X_S , hence X_S has at most $\binom{2n-2}{n-1}$ vertices. Since the configuration $\{e_i - e_j\}$ is unimodular, the bound is tight whenever the fan is simplicial and uses all the rays $e_i - e_j$.

Example 5.3.21. The upper bound on the number of vertices of a polytrope in \mathbb{TP}^3 is $\binom{2\cdot4-2}{4-1} = 20$. Figure 8 in [JK08] shows the five distinct combinatorial types of such extremal three-dimensional polytropes.

Proposition 5.3.22. If P and Q are tropical polytopes in \mathbb{TP}^{n-1} then $P \cap Q$ is also a tropical polytope.

Proof. Since P and Q are both tropically convex, $P \cap Q$ must also be. Consequently, if we can find a finite set of points in $P \cap Q$ whose convex hull contains all of $P \cap Q$, we will be done. By Theorem 5.3.16, P and Q are the finite unions of bounded cells $\{X_S\}$ and $\{X_T\}$ respectively, so $P \cap Q$ is the finite union of the cells $X_S \cap X_T$. Consider any $X_S \cap X_T$. Using Lemma 5.3.11 to obtain the inequality representations of X_S and X_T , we see that this region is of the form dictated by Proposition 5.3.19, and therefore obtainable as a cell X_W in some tropical complex. By Proposition 5.3.20, X_W is itself a tropical polytope, and we can find a finite set whose convex hull is equal to $X_S \cap X_T$. Taking the union of these sets over all choices of S and T then gives us the desired set of points whose convex hull contains all of $P \cap Q$. \Box

Proposition 5.3.23. Let $P \subset \mathbb{TP}^{n-1}$ be a tropical polytope. Then there exists a unique minimal set V such that $P = \operatorname{tconv}(V)$.

Proof. Suppose that P has two minimal generating sets, $V = \{v_1, \ldots, v_m\}$ and $W = \{w_1, \ldots, w_r\}$. Write each element of W as $w_i = \bigoplus_{j=1}^m c_{ij} \odot v_j$. We claim that $V \subseteq W$. Consider $v_1 \in V$ and write

$$v_1 = \bigoplus_{i=1}^r d_i \odot w_i = \bigoplus_{j=1}^m f_j \odot v_j \quad \text{where } f_j = \min_i (d_i + c_{ij}). \quad (5.13)$$

5.3. TROPICAL CONVEXITY

If the term $f_1 \odot v_1$ does not minimize any coordinate in the right-hand side of (5.13), then v_1 is a linear combination of v_2, \ldots, v_m , contradicting the minimality of V. However, if $f_1 \odot v_1$ minimizes any coordinate in this expression, it must minimize all of them, since $(v_1)_j - (v_1)_k = (f_1 \odot v_1)_j - (f_1 \odot v_1)_k$. In this case we get $v_1 = f_1 \odot v_1$, or $f_1 = 0$. Pick any i for which $f_1 = d_i + c_{i1}$; we claim that $w_i = c_{i1} \odot v_1$. Indeed, if any other term in $w_i = \bigoplus_{j=1}^m c_{ij} \odot v_j$ contributed nontrivially to w_i , that term would also contribute to the expression on the right-hand side of (5.13), which is a contradiction. Consequently, $V \subseteq W$, which means V = W since both sets are minimal by hypothesis. \Box

Every set $V = \{v_1, \ldots, v_r\}$ of r points in \mathbb{TP}^{n-1} specifies a tropical polytope $P = \operatorname{tconv}(V)$ equipped with a cell decomposition into the tropical complex generated by V. Each cell of this tropical complex is labelled by its type, which is an *n*-vector of finite subsets of $\{1, \ldots, r\}$. Two configurations V and W have the same *combinatorial type* if the types occurring in their tropical complexes are identical; note that by Lemma 5.3.14, this implies that the face posets of these polyhedral complexes are isomorphic.

With the definition in the previous paragraph, the statement of Theorem 5.3.2 has now finally been made precise. We will prove this correspondence between tropical complexes and subdivisions of products of simplices by constructing the polyhedral complex C_P in a higher-dimensional space.

Let $W = \mathbb{R}^{r+n}/\mathbb{R}(1, \ldots, 1, -1, \ldots, -1)$. The natural coordinates on W are denoted $(y, z) = (y_1, \ldots, y_r, z_1, \ldots, z_n)$. As before, we fix an ordered subset $V = \{v_1, \ldots, v_r\}$ of \mathbb{TP}^{n-1} . This defines the unbounded polyhedron

$$\mathcal{P}_{V} = \{ (y, z) \in W : y_{i} + z_{j} \le v_{ij} \text{ for } 1 \le i \le r \text{ and } 1 \le j \le n \}.$$
(5.14)

Lemma 5.3.24. There is a piecewise-linear isomorphism between the tropical complex generated by V and the complex of bounded faces of the (r + n - 1)-dimensional polyhedron \mathcal{P}_V . The image of a cell X_S of \mathcal{C}_P under this isomorphism is the bounded face $\{y_i+z_j = v_{ij} : i \in S_j\}$ of \mathcal{P}_V . That bounded face maps isomorphically to X_S via projection onto the z-coordinates.

Proof. Let F be a bounded face of \mathcal{P}_V , and define S_j via $i \in S_j$ if $y_i + z_j = v_{ij}$ is valid on all of F. If some y_i or z_j appears in no equality, then we can subtract arbitrary positive multiples of that basis vector to obtain elements of F, contradicting the assumption that F is bounded. Therefore, each i must appear in some S_j , and each S_j must be nonempty.

Since every y_i appears in some equality, given a specific z in the projection of F onto the z-coordinates, there exists a unique y for which $(y, z) \in F$, so this projection is an affine isomorphism from F to its image. We need to show that this image is equal to X_S . Let z be a point in the image of this projection, coming from a point (y, z) in the relative interior of F. We claim that $z \in X_S$. Indeed, looking at the *j*th coordinate of z, we find

$$-y_i + v_{ij} \ge z_j \quad \text{for all } i, \tag{5.15}$$

$$-y_i + v_{ij} = z_j \quad \text{for } i \in S_j. \tag{5.16}$$

The defining inequalities of X_S are $x_j - x_k \leq v_{ij} - v_{ik}$ with $i \in S_j$. Subtracting the inequality $-y_i + v_{ik} \geq z_k$ from the equality in (5.16) yields that this inequality is valid on z as well. Therefore, $z \in X_S$. Similar reasoning shows that S = type(z). We note that the relations (5.15) and (5.16) can be rewritten in terms of the tropical product of a row vector and a matrix:

$$z = (-y) \odot V = \bigoplus_{i=1}^{r} (-y_i) \odot v_i.$$
(5.17)

Conversely, suppose $z \in X_S$. We define $y = V \odot (-z)$. This means that

$$y_i = \min(v_{i1} - z_1, v_{i2} - z_2, \dots, v_{in} - z_n).$$
 (5.18)

We claim that $(y, z) \in F$. Indeed, we certainly have $y_i + z_j \leq v_{ij}$ for all i and j, so $(y, z) \in \mathcal{P}_V$. Furthermore, when $i \in S_j$, we know that $v_{ij} - z_j$ achieves the minimum in the right-hand side of (5.18), so that $v_{ij} - z_j = y_i$ and $y_i + z_j = v_{ij}$ is satisfied. Consequently, $(y, z) \in F$ as desired.

It follows that the two complexes are isomorphic: if F is a face corresponding to X_S and G is a face corresponding to X_T , where S and T are both types, then X_S is a face of X_T if and only if $T \subseteq S$. However, by the discussion above, this is equivalent to saying that the equalities G satisfies (which correspond to T) are a subset of the equalities F satisfies (which correspond to S); this is true if and only if F is a face of G. So X_S is a face of X_T if and only if F is a face of G, which establishes the assertion. \Box

The boundary complex of the polyhedron \mathcal{P}_V is polar to the regular subdivision of the product of simplices $\Delta_{r-1} \times \Delta_{n-1}$ defined by the weights v_{ij} . We denote this regular polyhedral subdivision by $(\partial \mathcal{P}_V)^*$. Explicitly, a subset of vertices (e_i, e_j) of $\Delta_{r-1} \times \Delta_{n-1}$ forms a cell of $(\partial \mathcal{P}_V)^*$ if and only if the equations $y_i + z_j = v_{ij}$ indexed by these vertices specify a face of the polyhedron \mathcal{P}_V . We now present the proof of the result stated earlier. Proof of Theorem 5.3.2: The poset of bounded faces of \mathcal{P}_V is antiisomorphic to the poset of interior cells of the subdivision $(\partial \mathcal{P}_V)^*$ of $\Delta_{r-1} \times \Delta_{n-1}$. Since every full-dimensional cell of $(\partial \mathcal{P}_V)^*$ is interior, the subdivision is uniquely determined by its interior cells. Hence, the combinatorial type of \mathcal{P}_V is determined by the lists of facets containing each bounded face of \mathcal{P}_V . These lists are precisely the types of regions in \mathcal{C}_P by Lemma 5.3.24. This completes the proof.

Theorem 5.3.2, which establishes a bijection between the tropical complexes generated by r points in \mathbb{TP}^{n-1} and the regular subdivisions of a product of simplices $\Delta_{r-1} \times \Delta_{n-1}$, has many striking consequences. One of them is the identification of the row span and column span of a matrix:

Theorem 5.3.25. Given any matrix $M \in \mathbb{R}^{r \times n}$, the tropical complex generated by its column vectors is isomorphic to the tropical complex generated by its row vectors. This isomorphism is gotten by restricting the piecewise linear maps $\mathbb{R}^n \to \mathbb{R}^r$, $z \mapsto M \odot (-z)$ and $\mathbb{R}^r \to \mathbb{R}^n$, $y \mapsto (-y) \odot M$.

Proof. By Theorem 5.3.2, the matrix M corresponds via the polyhedron \mathcal{P}_M to a regular subdivision of $\Delta_{r-1} \times \Delta_{n-1}$, and the complex of interior faces of this regular subdivision is combinatorially isomorphic to both the tropical complex generated by its row vectors, which are r points in \mathbb{TP}^{n-1} , and the tropical complex generated by its column vectors, which are n points in \mathbb{TP}^{r-1} . Furthermore, Lemma 5.3.24 tells us that the cell in \mathcal{P}_M is affinely isomorphic to its corresponding cell in both tropical complexes. Finally, in the proof of Lemma 5.3.24, we showed that the point (y, z) in a bounded face F of \mathcal{P}_M satisfies $y = M \odot (-z)$ and $z = (-y) \odot M$. This point projects to y and z, and so the piecewise-linear isomorphism mapping these two complexes to each other is defined by the stated maps.

The common tropical complex of these two tropical polytopes is given by the complex of bounded faces of the common polyhedron \mathcal{P}_M , which lives in a space of dimension r + n - 1; the tropical polytopes are unfoldings of this complex into dimensions r - 1 and n - 1. Theorem 5.3.25 also gives a natural bijection between the combinatorial types of tropical convex hulls of r points in \mathbb{TP}^{n-1} and the combinatorial types of tropical convex hulls of n points in \mathbb{TP}^{r-1} , incidentally proving that there are the same number of each.

Figure 5.5 shows the dual of the convex hull of $\{(0, 0, 2), (0, 2, 0), (0, 1, -2)\},\$

Figure 5.5: A demonstration of tropical polytope duality.

also a tropical triangle (here r = n = 3). For instance, we compute:

$$\begin{pmatrix} 0 & 0 & 2 \\ 0 & 2 & 0 \\ 0 & 1 & -2 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ -2 \end{pmatrix} = \begin{pmatrix} 0 \\ -2 \\ -4 \end{pmatrix}$$

This point is the image of the point (0, 0, 2) under this duality map. Note that duality does not preserve the generating set; the polytope on the right is the convex hull of points $\{F, D, B\}$, while the polytope on the left is the convex hull of points $\{F, A, C\}$. This is necessary, of course, since in general a polytope with r vertices is mapped to a polytope with n vertices, and rneed not equal n as it does in our example.

We now discuss the generic case when the subdivision $(\partial \mathcal{P}_V)^*$ is a regular triangulation of $\Delta_{r-1} \times \Delta_{n-1}$.

Proposition 5.3.26. For a configuration V of r points in \mathbb{TP}^{n-1} with $r \ge n$ the following are equivalent:

- 1. The regular subdivision $(\partial \mathcal{P}_V)^*$ is a triangulation of $\Delta_{r-1} \times \Delta_{n-1}$.
- 2. No k of the points in V have projections onto a k-dimensional coordinate subspace which lie in a tropical hyperplane, for any $2 \le k \le n$.
- 3. No $k \times k$ -submatrix of the $r \times n$ -matrix (v_{ij}) is tropically singular, i.e. has vanishing tropical determinant, for any $2 \le k \le n$.

Proof. The last equivalence follows from Proposition 5.1.7. We shall prove that (1) and (3) are equivalent. The tropical determinant of a k by k matrix M is the tropical polynomial $\bigoplus_{\sigma \in S_k} (\bigoplus_{i=1}^k M_{i\sigma(i)})$. The matrix M is tropically singular if the minimum $\min_{\sigma \in S_k} (\sum_{i=1}^k M_{i\sigma(i)})$ is achieved twice.

The regular subdivision $(\partial \mathcal{P}_V)^*$ is a triangulation if and only if the polyhedron \mathcal{P}_V is simple, which is to say if and only if no r + n of the facets $y_i + z_j \leq v_{ij}$ meet at a single vertex. For each vertex v, consider the bipartite graph G_v consisting of vertices y_1, \ldots, y_n and z_1, \ldots, z_j with an edge connecting y_i and z_j if v lies on the corresponding facet. This graph is connected, since each y_i and z_j appears in some such inequality, and thus it will have a cycle if and only if it has at least r + n edges. Consequently, \mathcal{P}_V is not simple if and only there exists some G_v with a cycle.

If there is a cycle, without loss of generality it is $y_1, z_1, y_2, z_2, \ldots, y_k, z_k$. Consider the submatrix M of (v_{ij}) given by $1 \leq i, j \leq k$. We have $y_1 + z_1 = M_{11}, y_2 + z_2 = M_{22}$, and so on, and also $z_1 + y_2 = M_{12}, \ldots, z_k + y_1 = M_{k1}$. Adding up all of these equalities yields $y_1 + \cdots + y_k + z_1 + \cdots + z_k = M_{11} + \cdots + M_{kk} = M_{12} + \cdots + M_{k1}$. But consider any permutation σ in the symmetric group S_k . Since $M_{i\sigma(i)} = v_{i\sigma(i)} \geq y_i + z_{\sigma(i)}$, we have $\sum M_{i\sigma(i)} \geq x_1 + \cdots + x_k + y_1 + \cdots + y_k$. Consequently, the permutations equal to the identity and to $(12 \cdots k)$ simultaneously minimize the determinant of the minor M. This logic is reversible, proving the equivalence of (1) and (3). \Box

If the r points of V are in general position, the tropical complex they generate is a generic tropical complex. Such a tropical complex is dual to the co-complex of interior faces in a regular triangulation of $\Delta_{r-1} \times \Delta_{n-1}$.

Corollary 5.3.27. All tropical complexes generated by r points in general position in \mathbb{TP}^{n-1} have the same f-vector. Specifically, the number of faces of dimension k is equal to the multinomial coefficient

$$\binom{r+n-k-2}{r-k-1,n-k-1,k} = \frac{(r+n-k-2)!}{(r-k-1)! \cdot (n-k-1)! \cdot k!}.$$

Proof. By Proposition 5.3.26, these objects are in bijection with regular triangulations of $P = \Delta_{r-1} \times \Delta_{n-1}$. The polytope P is unimodular, which means that all simplices formed by vertices of P are unimodular. This property implies that all triangulations of P have the same f-vector. The number of faces of dimension k of the tropical complex generated by given r points is equal to the number of interior faces of codimension k in the corresponding triangulation. Since all triangulations of all products of simplices have the same f-vector, they also have the same interior f-vector, which can be computed by taking the f-vector and subtracting off the f-vectors of the induced triangulations on the proper faces of P. These proper faces are products of simplices and hence equidecomposable, so all of these induced triangulations have f-vectors independent of the original triangulation as well.

To compute this number, we therefore need only compute it for one tropical complex. Let the vectors v_i , $1 \le i \le r$, be given by $v_i = (i, 2i, \dots, ni)$. By Theorem 5.3.11, to count the faces of dimension k in this tropical complex, we enumerate the existing types with k degrees of freedom. Consider any index i. We claim that for any x in the tropical convex hull of $\{v_i\}$, the set $\{j \mid i \in S_i\}$ is an interval I_i , and that if i < m, the intervals I_m and I_i Figure 5.6: The 35 symmetry classes of tropical quadrangles in \mathbb{TP}^2 .

meet in at most one point, which in that case is the largest element of I_m and the smallest element of I_i .

Suppose we have $i \in S_j$ and $m \in S_l$ with i < m. Then we have by definition $v_{ij} - x_j \leq v_{il} - x_l$ and $v_{ml} - x_l \leq v_{mj} - x_j$. Adding these inequalities yields $v_{ij} + v_{ml} \leq v_{il} + v_{mj}$, or $ij + ml \leq il + mj$. Since i < m, it follows that we must have $l \leq j$. Therefore, we can never have $i \in S_j$ and $m \in S_l$ with i < m and j < l. The claim follows immediately, since the I_i cover [1, n].

The number of degrees of freedom of an interval set (I_1, \ldots, I_r) is easily seen to be the number of *i* for which I_i and I_{i+1} are disjoint. Given this, it follows from a simple combinatorial counting argument that the number of interval sets with *k* degrees of freedom is the multinomial coefficient given above. Finally, a representative for every interval set is given by $x_j = x_{j+1} - c_j$, where if S_j and S_{j+1} have an element *i* in common (they can have at most one), $c_j = i$, and if not then $c_j = (\min(S_j) + \max(S_{j+1}))/2$. Therefore, each interval set is in fact a valid type, and our enumeration is complete.

Corollary 5.3.28. The number of combinatorially distinct generic tropical complexes generated by r points in \mathbb{TP}^{n-1} equals the number of distinct regular triangulations of $\Delta_{r-1} \times \Delta_{n-1}$. The number of respective symmetry classes under the natural action of the product of symmetric groups $G = S_r \times S_n$ on both spaces is also the same.

The symmetries in the group G correspond to a natural action on $\Delta_{r-1} \times \Delta_{n-1}$ given by permuting the vertices of the two component simplices; the symmetries in the symmetric group S_r correspond to permuting the points in a tropical polytope, while those in the symmetric group S_n correspond to permuting the coordinates. (These are dual by Corollary 5.3.25.) The number of symmetry classes of regular triangulations of the polytope $\Delta_{r-1} \times \Delta_{n-1}$ is computable via Jörg Rambau's TOPCOM [Ram02] for small r and n:

	2	3
2	5	35
3	35	7,955
4	530	
5	13,631	

For example, the (2,3) entry of the table divulges that there are 35 symmetry classes of regular triangulations of $\Delta_2 \times \Delta_3$. These correspond to the 35 combinatorial types of four-point configurations in \mathbb{TP}^2 , or to the 35 combinatorial types of three-point configurations in \mathbb{TP}^3 . These 35 configurations (with the tropical complexes they generate) are shown in Figure 5.6.

5.4 The Rank of a Matrix

The rank of a matrix M is one of the most important notions in linear algebra. This number can be defined in many different ways. In particular, the following three definitions are equivalent in classical linear algebra:

- The rank of M is the smallest integer r for which M can be written as the sum of r rank one matrices. A matrix has rank one if it is the product of a column vector and a row vector.
- The rank of M is the smallest dimension of any linear space containing the columns of M.
- The rank of M is the largest integer r such that M has a non-singular $r \times r$ minor.

Our objective is to examine these familiar definitions over the *tropical* semiring $(\mathbb{R}, \oplus, \odot)$. The set \mathbb{R}^d of real *d*-vectors and the set $\mathbb{R}^{d \times n}$ of real $d \times n$ -matrices are semimodules over the semiring $(\mathbb{R}, \oplus, \odot)$. The operations of matrix addition and matrix multiplication are well defined. All three definitions of rank make sense over the tropical semiring $(\mathbb{R}, \oplus, \odot)$:

Definition 5.4.1. The *Barvinok rank* of a matrix $M \in \mathbb{R}^{d \times n}$ is the smallest integer r for which M can be written as the tropical sum of r rank one matrices. Here, we say that a $d \times n$ -matrix has *rank one* if it is the tropical matrix product of a $d \times 1$ -matrix and a $1 \times n$ -matrix.

Definition 5.4.2. The Kapranov rank of a matrix $M \in \mathbb{R}^{d \times n}$ is the smallest dimension of any tropicalization of linear space in T^d (in the sense of 4.3) containing the columns of M. Here K is allowed to be any field of characteristic zero.

Definition 5.4.3. A square matrix $M = (m_{ij}) \in \mathbb{R}^{r \times r}$ is tropically singular if the minimum in the evaluation of the tropical determinant

 $\bigoplus_{\sigma \in S_r} m_{1\sigma_1} \odot m_{2\sigma_2} \odot \cdots \odot m_{r\sigma_r} = \min \{ m_{1\sigma_1} + m_{2\sigma_2} + \cdots + m_{r\sigma_r} : \sigma \in S_r \}$

is attained at least twice. Here S_r denotes the symmetric group on $\{1, 2, \ldots, r\}$. See also Proposition 5.1.7. The *tropical rank* of a matrix $M \in \mathbb{R}^{d \times n}$ is the largest integer r such that M has a non-singular $r \times r$ minor.

We will show that these three definitions are not equivalent:

Theorem 5.4.4. For every matrix M with entries in the tropical semiring,

 $tropical rank(M) \leq Kapranov rank(M) \leq Barvinok rank(M).$ (5.19)

Both of these inequalities can be strict.

The proof of Theorem 5.4.4 consists of Propositions 5.4.15, 5.4.17, 5.4.21 and Theorem 5.4.22. As we go along, several alternative characterizations of the Barvinok, Kapranov and tropical ranks will be offered. One of them arises from the fact that every $d \times n$ -matrix M defines a tropically linear map $\mathbb{R}^n \to \mathbb{R}^d$. The image of M is regarded as a tropical polytope in $\mathbb{TP}^{d-1} = \mathbb{R}^d/\mathbb{R}(1, 1, \ldots, 1)$. We shall see that the tropical rank of M equals the dimension of this tropical polytope plus one. The discrepancy between Definitions 5.4.1, 5.4.2 and 5.4.3 comes from the crucial distinction between tropical polytopes, tropicalized linear spaces, and tropical linear spaces.

We start out by studying the Barvinok rank (Definition 5.4.1). This notion of rank arises naturally in the context of combinatorial optimization. Barvinok, Johnson, Woeginger and Woodrooofe [BJWW98], building on earlier work of Barvinok, showed that for fixed r, the Traveling Salesman Problem can be solved in polynomial time if the distance matrix is the tropical sum of r matrices of tropical rank one (with \oplus as "max" instead of "min"). This motivates the definition and nomenclature of Barvinok rank as the smallest r for which $M \in \mathbb{R}^{d \times n}$ is expressible in this fashion. Since matrices of tropical rank one are of the form $X \odot Y^T$, for two column vectors $X \in \mathbb{R}^d$ and $Y \in \mathbb{R}^n$, this is equivalent to saying that M has a representation

$$M = X_1^T \odot Y_1 \oplus X_2^T \odot Y_2 \oplus \cdots \oplus X_r^T \odot Y_r.$$
 (5.20)

158

5.4. THE RANK OF A MATRIX

For example, here is a 3×3 -matrix which has Barvinok rank two:

$$M = \begin{pmatrix} 0 & 4 & 2 \\ 2 & 1 & 0 \\ 2 & 4 & 3 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \\ 2 \end{pmatrix} \odot (0, 4, 2) \oplus \begin{pmatrix} 3 \\ 0 \\ 3 \end{pmatrix} \odot (2, 1, 0).$$
(5.21)

This matrix also has tropical rank 2 and Kapranov rank 2. The column vectors lie on the tropical line in \mathbb{TP}^2 defined by $2 \odot x_1 \oplus 3 \odot x_2 \oplus 0 \odot x_3$.

We next present two reformulations of Barvinok rank: in terms of tropical convex hulls as in Section 5.3, and via a "tropical morphism" in matrix space.

Proposition 5.4.5. For a real $d \times n$ -matrix M, the following are equivalent:

- (a) M has Barvinok rank at most r.
- (b) The columns of M lie in the tropical convex hull of r points in \mathbb{TP}^{d-1} .
- (c) There are matrices $X \in \mathbb{R}^{d \times r}$ and $Y \in \mathbb{R}^{r \times n}$ such that $M = X \odot Y$. Equivalently, M lies in the image of the following tropical morphism, which is defined by matrix multiplication:

$$\phi_r : \mathbb{R}^{d \times r} \times \mathbb{R}^{r \times n} \to \mathbb{R}^{d \times n} , \ (X, Y) \mapsto X \odot Y.$$
 (5.22)

Proof. Let $M_1, \ldots, M_n \in \mathbb{R}^d$ be the column vectors of M. Let $X_1, \ldots, X_r \in \mathbb{R}^d$ and $Y_1, \ldots, Y_r \in \mathbb{R}^n$ be the columns of two unspecified matrices $X \in \mathbb{R}^{d \times r}$ and $Y \in \mathbb{R}^{n \times r}$. Let Y_{ij} denote the *j*th coordinate of Y_i . The following three algebraic identities are easily seen to be equivalent:

- (a) $M = X_1 \odot Y_1^T \oplus X_2 \odot Y_2^T \oplus \cdots \oplus X_r \odot Y_r^T$,
- (b) $M_j = Y_{1j} \odot X_1 \oplus Y_{2j} \odot X_2 \oplus \cdots \oplus Y_{rj} \odot X_r$ for all $j = 1, \ldots, n$, and

(c)
$$M = X \odot Y^T$$
.

Statement (b) says that each column vector of M lies in the tropical convex hull of X_1, \ldots, X_r . The entries of the matrix Y are the multipliers in that tropical convex combination. This shows that the three conditions (a), (b) and (c) in the statement of the proposition are equivalent.

We next take a closer look at the structure of the multiplication map ϕ_r .

Proposition 5.4.6. The map ϕ_r is piecewise-linear. The domains of linearity form a fan in $\mathbb{R}^{d \times r} \times \mathbb{R}^{r \times n}$. This fan is the common refinement of the normal fans of dn simplices of dimension r - 1.

Proof. Let $U = (u_{ij})$ and $V = (v_{jk})$ be matrices of indeterminates of format $d \times r$ and $r \times n$ respectively. The entries of the classical matrix product UV are the dn quadratic polynomials $u_{i1}v_{1k} + u_{i2}v_{2k} + \cdots + u_{ir}v_{rk}$. The Newton polytope of each such quadric is an (r-1)-dimensional simplex P_{ik} . Let $P = \sum_{i=1}^{d} \sum_{k=1}^{n} P_{ik}$ denote the Minkowski sum of these dn simplices. This is a polytope of dimension $(2 \cdot \min(d, n) - 1)(r-1)$ sitting inside $\mathbb{R}^{d \times r} \times \mathbb{R}^{r \times n}$.

The *ik*-coordinate of the tropical map ϕ_r takes a pair of matrices (X, Y) to the real number $\min(x_{i1} + y_{1k}, \ldots, x_{ir} + y_{rk})$. This function is the support function of the simplex P_{ik} . It is linear on each cone in the normal fan of P_{ik} . Hence ϕ_r is a linear map on the common refinement of the normal fans of the simplices P_{ik} . This common refinement is the normal fan of their Minkowski sum P. We conclude that ϕ_r is piecewise-linear on the normal fan of P. \Box

Corollary 5.4.7. If r = 2 then the map ϕ_2 is piecewise-linear with respect to the regions in an arrangement of dn hyperplanes in $\mathbb{R}^{d \times 2} \times \mathbb{R}^{2 \times n}$.

Proof. If r = 2 then each P_{ij} is a line segment, and their Minkowski sum P is a zonotope of dimension $2 \cdot \min(d, n) - 1$. The normal fan of the zonotope P is a hyperplane arrangement, and it follows from the previous proof that ϕ_r is piecewise linear on that hyperplane arrangement.

Example 5.4.8. Let d = n = 3 and r = 2. Then P is a four-dimensional zonotope with nine zones in $\mathbb{R}^{12} = \mathbb{R}^{3\times 2} \times \mathbb{R}^{2\times 3}$. This zonotope has 230 vertices, so the dual hyperplane arrangement has 230 maximal regions. The tropical morphism ϕ_2 maps each of these 230 regions linearly onto an 8-dimensional cone in $\mathbb{R}^{3\times 3}$. The image of ϕ_2 is the set of all tropically singular 3×3 -matrices. This is a polyhedral fan with 15 maximal cones. It is the codimension one skeleton of the normal fan of the four-dimensional Birkhoff polytope (the convex hull of all six 3×3 -permutation matrices).

By Proposition 5.4.5, the set of all Barvinok matrices of rank $\leq r$ is the image of the tropical morphism ϕ_r . In particular, this set is a polyhedral fan in $\mathbb{R}^{d \times n}$, as in the previous example. The distinction between the Barvinok rank and the Kapranov rank can be explained by the following general fact of tropical algebraic geometry: For most polynomial maps, the image of the tropicalization is strictly contained in the tropicalization of the image.

We next demonstrate that the Barvinok rank can be much larger than the other two notions of rank. The example to be considered is the $n \times n$ -matrix

$$C_n = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}.$$
 (5.23)

This looks like the identity matrix (in classical arithmetic) but it is not the identity matrix in tropical arithmetic. That honor belongs to the $n \times n$ -matrix whose diagonal entries are 0 and whose off-diagonal entries are ∞ .

Theorem 5.4.9. The Barvinok rank of the matrix C_n is the smallest positive integer r such that

$$n \leq \binom{r}{\left\lfloor \frac{r}{2} \right\rfloor}.$$

Proof. Let r be an integer and assume that $n \leq \binom{r}{\lfloor r/2 \rfloor}$. We first show that Barvinok rank $(C_n) \leq r$. Let S_1, \ldots, S_n be distinct subsets of $\{1, \ldots, r\}$ each having cardinality $\lfloor r/2 \rfloor$. For each $k \in 1, \ldots, r$, we define an $n \times n$ -matrix $X_k = (x_{ij}^k)$ with entries in $\{0, 1, 2\}$ as follows:

$$x_{ij}^k = 0$$
 if $k \in S_i \setminus S_j$, $x_{ij}^k = 2$ if $k \in S_j \setminus S_i$, and $x_{ij}^k = 1$ otherwise.

The matrix X_k has tropical rank one. To see this, let $V_k \in \{0,1\}^n$ denote the vector with *i*th coordinate equal to one or zero depending on whether kis an element of S_i or not. Then we have

$$X_k = V_k^T \odot (1 \odot (-V_k)).$$

To prove Barvinok rank $(C_n) \leq r$, it now suffices to establish the identity

$$C_n = X_1 \oplus X_2 \oplus \cdots \oplus X_r.$$

Indeed, all diagonal entries of the matrices on the right hand side are 1, and the off-diagonal entries of the right hand side are $\min(x_{ij}^1, x_{ij}^2, \ldots, x_{ij}^r) = 0$, because $S_i \setminus S_j$ is non-empty for $i \neq j$.

To prove the converse direction, we consider an arbitrary representation

$$C_n = Y_1 \oplus Y_2 \oplus \cdots \oplus Y_r$$

where the matrices $Y_k = (y_{ij}^k)$ have tropical rank one. For each k we set $T_k := \{(i, j) : y_{ij}^k = 0\}$. Since the matrices Y_k are non-negative and have tropical rank one, it follows that each T_k is a product $I_k \times J_k$, where I_k and J_k are subsets of $\{1, \ldots, n\}$. Moreover, we have $I_k \cap J_k = \emptyset$ because the diagonal entries of Y_k are not zero. For each $i = 1, \ldots, n$ we set

$$S_i := \{k : i \in I_k\} \subseteq \{1, \dots, r\}.$$

We claim that no two of the sets S_1, \ldots, S_n are contained in one another. Sperner's Theorem [AZ04, Chapter 23] then proves that $n \leq \binom{r}{\lfloor r/2 \rfloor}$. To prove the claim, observe that if $S_i \subset S_j$ then the entry $y_{i,j}^k$ cannot be zero for any k. Indeed, if $k \in S_i \subseteq S_j$ then $j \in I_k$ implies $j \notin J_k$. And if $k \notin S_i$ then $i \notin I_k$.

Example 5.4.10. The matrix C_6 has Barvinok rank 4. The upper bound is shown by the following tropical sum decomposition of C_6 :

$ \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} $	0 1 0 0	0 0 1 0	0 0 0 1	0 0 0 0	$ \begin{array}{c} 0\\0\\0\\0\\0\\0\end{array} \end{array} $	_		$\begin{pmatrix} 1\\ 1\\ 1\\ 0\\ 0 \end{pmatrix}$	1 1 1 0	1 1 1 0	2 2 2 1	2 2 2 1	2 2 2 1 1	\oplus	$\begin{pmatrix} 1 \\ 1 \\ 2 \\ 2 \\ 2 \\ 2 \end{pmatrix}$	1 1 2 2 2	0 0 1 1	0 0 1 1	0 0 1 1	$ \begin{array}{c} 1\\1\\2\\2\\2\end{array} \end{array} $	
$\begin{pmatrix} 0\\ 0 \end{pmatrix}$	0	0	0	1 0	$\begin{pmatrix} 0\\1 \end{pmatrix}$			$ \begin{bmatrix} 0\\0\\\\ 1\\\\ 2\\1 \end{bmatrix} $				1 1 1 2 1	$\begin{pmatrix} 1\\1\\\end{pmatrix}$ $\begin{pmatrix} 0\\1\\0\\\end{pmatrix}$		$\begin{pmatrix} 2\\1\\ \begin{pmatrix} 1\\0\\0 \end{pmatrix}$	2 1 2 1 1	1 0 2 1 1	1 0 2 1 1		$\begin{pmatrix} 2\\1 \end{pmatrix}$ 1 0 0	
							\oplus	$\begin{vmatrix} 1\\2\\1\\2 \end{vmatrix}$	1 0 1	1 2 1 2	1 0 1	1 2 1 2	1 0 1	\oplus	$\begin{bmatrix} 0\\ 0\\ 1\\ 1 \end{bmatrix}$	$1 \\ 2 \\ 2$	$1 \\ 2 \\ 2$	1 2 2	0 1 1	0 1 1	

Similarly, C_{36} has Barvinok rank 8, its 35×35 minors have Barvinok rank 7, and its 8×8 minors have Barvinok rank at most 5. Asymptotically,

Barvinok rank $(C_n) \sim \log_2 n$.

We will see in Examples 5.4.14 and 5.4.19 that the Kapranov rank and tropical rank of the matrix C_n are both two.

We now fix an algebraically closed field K of characteristic zero that has a surjective valuation val : $K^* \to \mathbb{R}$. If I is any ideal in $K[x_1, \ldots, x_d]$ then

5.4. THE RANK OF A MATRIX

we write V(I) for its variety in the *d*-dimensional algebraic torus $(K^*)^d$. The tropical variety $\mathcal{T}(I) \subset \mathbb{R}^d$ is the image of V(I) under the map deg, and, by the Fundamental Theorem, it coincides with the set of vectors $w \in \mathbb{R}^n$ such that the initial ideal $\operatorname{in}_w(I) = \langle \operatorname{in}_w(f) : f \in I \rangle$ contains no monomial. The initial ideal $\operatorname{in}_w(I)$ can be computed from any generating set of I by computing a Gröbner basis with respect to any term order that refines w. This shows that $\operatorname{trop}(I)$ is a polyhedral subcomplex of the Gröbner fan of I, and, as discussed earlier, it leads to an algorithm for computing $\operatorname{trop}(I)$.

Recall from Definition 5.4.2 that the Kapranov rank of a matrix $M \in \mathbb{R}^{d \times n}$ is the smallest dimension of any tropical linear space containing the columns of M. It is not completely apparent in this definition that the Kapranov rank of a matrix and its transpose are the same, but this follows from our next result. Let J_r denote the ideal generated by all the $(r + 1) \times (r+1)$ -subdeterminants of an $d \times n$ -matrix of indeterminates (x_{ij}) . This is a prime ideal of dimension $rd + rn - r^2$ in the polynomial ring $K[x_{ij}]$, and the generating determinants form a Gröbner basis. The variety $V(J_r)$ consists of all $d \times n$ -matrices with entries in K^* whose (classical) rank is at most r.

Theorem 5.4.11. For a matrix $M = (m_{ij}) \in \mathbb{R}^{d \times n}$ the following statements are equivalent:

- (a) The Kapranov rank of M is at most r.
- (b) The matrix M lies in the tropical determinantal variety $trop(J_r)$.
- (c) There exists a $d \times n$ -matrix $F = (f_{ij})$ with entries in K^* such that the rank of F is less than or equal to r and $\deg(f_{ij}) = m_{ij}$ for all i and j. We call F a lift of M, and we write $\operatorname{val}(F) = M$.

Proof. The equivalence of (b) and (c) is the Fundamental Theorem applied to the ideal J_r since, over the field K, lying in the variety of the determinantal ideal J_r is equivalent to having rank at most r. To see that (c) implies (a), consider the linear subspace of K^d spanned by the columns of F. This is an r-dimensional linear space over a field, so it is defined by an ideal Igenerated by d - r linearly independent linear forms in $K[x_1, \ldots, x_d]$. The tropical linear space trop(I) contains all the column vectors of $M = \deg(F)$.

Conversely, suppose that (a) holds, and let L be a tropical linear space of dimension r containing the columns of M. Pick a linear ideal I in $K[x_1, \ldots, x_d]$ such that $L = \operatorname{trop}(I)$. By applying the definition of tropical variety to the ideal I, we see that each column vector of M has a preimage in $V(I) \subset (K^*)^d$

under the valuation map. Let F be the $d \times n$ -matrix over K whose columns are these preimages. Then the column space of F is contained in the variety defined by I, so we have $\operatorname{rank}(F) \leq r$, and $\deg(F) = M$ as desired. \Box

Corollary 5.4.12. The Kapranov rank of a matrix $M \in \mathbb{R}^{d \times n}$ is the smallest rank of any lift of M.

Example 5.4.13. The following classical 3×3 -matrix has rank 2 over K:

$$F = \begin{pmatrix} 1 & t^4 & t^2 \\ t^2 & t & 1 \\ t^2 + t^5 & t^4 + t^6 & t^3 + t^4 \end{pmatrix}$$

We have val(F) = M, so F is a lift of the 3×3-matrix M in (5.21).

The ideal J_1 is generated by the 2 × 2-minors $x_{ij}x_{kl} - x_{il}x_{kj}$ of the $d \times n$ matrix (x_{ij}) . Therefore, a matrix of Kapranov rank one must certainly satisfy the linear equations $m_{ij} + m_{kl} = m_{il} + m_{kj}$. This happens if and only if there exist real vectors $X = (x_1, \ldots, x_d)$ and $Y = (y_1, \ldots, y_n)$ with

$$m_{ij} = x_i + y_j$$
 for all $i, j \iff m_{ij} = x_i \odot y_j$ for all $i, j \iff M = X^T \odot Y$.

Conversely, if such X and Y exist, we can lift M to a matrix of rank one by substituting $t^{m_{ij}}$ for m_{ij} . Therefore, a matrix M has Kapranov rank one if and only if it has Barvinok rank one. In general, the Kapranov rank can be much smaller than the Barvinok rank, as the following example shows.

Example 5.4.14. Let $n \ge 3$ and consider the matrix C_n in Theorem 5.4.9. The matrix C_n does not have Kapranov rank one, so its Kapranov rank is least two. Let a_3, a_4, \ldots, a_n be distinct scalars with $val(a_i) = 0$. The matrix

$$F_n = \begin{pmatrix} t & 1 & t + a_3 & t + a_4 & \cdots & t + a_n \\ 1 & t & 1 + a_3 t & 1 + a_4 t & \cdots & 1 + a_n t \\ t - a_3 & 1 & t & t - a_3 + a_4 & \cdots & t - a_3 + a_n \\ t - a_4 & 1 & t - a_4 + a_3 & t & \cdots & t - a_4 + a_n \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ t - a_n & 1 & t - a_n + a_3 & t - a_n + a_4 & \cdots & t \end{pmatrix}$$

has rank 2 because the *i*-th column (for $i \geq 3$) equals the first column plus a_i times the second column. Since $\deg(F_n) = C_n$, we conclude that C_n has Kapranov rank two. The tropicalized plane containing the columns of C_n is $\operatorname{trop}(U_{2,n})$, where $U_{2,n}$ is the uniform matroid as in Example 5.2.7.

164

5.4. THE RANK OF A MATRIX

The following proposition establishes half of Theorem 5.4.4.

Proposition 5.4.15. Every matrix $M \in \mathbb{R}^{d \times n}$ satisfies Kapranov rank $(M) \leq Barvinok rank(M)$, and this inequality can be strict.

Proof. Suppose that M has Barvinok rank r. Write $M = M_1 \oplus \cdots \oplus M_r$ where each M_i has Barvinok rank one. Then M_i has Kapranov rank one, so there exists a rank one matrix F_i over K such that $\deg(F_i) = M_i$. Moreover, by multiplying the matrices F_i by random complex numbers, we can choose F_i such that there is no cancellation of leading terms in t when we form the matrix $F = F_1 + \cdots + F_r$. This means $\deg(F) = M$. Clearly, the matrix F has rank $\leq r$. Theorem 5.4.11 implies that M has Kapranov rank $\leq r$. Example 5.4.14 shows that the inequality can be strict.

A general algorithm for computing the Kapranov rank of a matrix Minvolves computing a Gröbner basis of the determinantal ideal J_r . Suppose we wish to decide whether a given real $d \times n$ -matrix $M = (m_{ij})$ has Kapranov rank > r. To decide this question, we fix any term order \prec_M on the polynomial ring $\mathbb{Q}[x_{ij}]$ which refines the partial ordering on monomials given assigning weight m_{ij} to the variable x_{ij} , and we compute the reduced Gröbner basis \mathcal{G} of J_r in the term order \prec_M . For each polynomial g in \mathcal{G} , we consider its leading form $\operatorname{in}_M(g)$ with respect to the partial ordering coming from M. Note that $\operatorname{in}_{\prec_M}(\operatorname{in}_M(g)) = \operatorname{in}_{\prec_M}(g)$ for all $g \in \mathcal{G}$.

The ideal generated by the set of leading forms $\{in_M(g) : g \in \mathcal{G}\}$ is the initial ideal $in_M(J_r)$. Let x^{all} denote the product of all dn unknowns x_{ij} . The second step in our algorithm is to compute the saturation of the initial ideal:

$$\left(\operatorname{in}_{M}(J_{r}):\langle x^{all}\rangle^{\infty}\right) = \left\{f \in \mathbb{C}[x_{ij}]: f(x^{all})^{s} \in J_{r} \text{ for some } s \in \mathbb{N}\right\}.$$
 (5.24)

Computing such an ideal quotient, given the generators $in_M(g)$, is a standard operation in computational commutative algebra. It is a built-in command in software systems such as CoCoA, Macaulay 2 or Singular. We conclude:

Corollary 5.4.16. The matrix M has Kapranov rank > r if and only if (5.24) is the unit ideal $\langle 1 \rangle$.

Our next step is to prove the first inequality in Theorem 5.4.4.

Proposition 5.4.17. Every matrix $M \in \mathbb{R}^{d \times n}$ satisfies tropical rank $(M) \leq Kapranov rank(M)$.

Proof. If the matrix M has a tropically non-singular $r \times r$ minor, then any lift of M to the power series field K must have the corresponding $r \times r$ -minor non-singular over K, since the leading exponent of its determinant occurs only once in the sum. Consequently, no lift of M to K can have rank less than r. By Theorem 5.4.11, the Kapranov rank of M must be at least r. \Box

We next present a combinatorial formula for the tropical rank of a zeroone matrix, or any matrix which has only two distinct entries. We define the *support* of a vector in tropical space \mathbb{R}^d as the set of its zero coordinates. We define the *support poset* of a matrix M to be the set of all unions of supports of column vectors of M. This set is partially ordered by inclusion.

Proposition 5.4.18. The tropical rank of a zero-one matrix with no column of all ones equals the maximum length of a chain in its support poset.

The assumption that there is no column of all ones is needed because a column of zeroes and a column of ones represent the same point in \mathbb{TP}^{d-1} .

Proof. There is no loss of generality in assuming that every union of supports of columns of M is the support of a column. Indeed, the tropical sum of a set of columns gives a column whose support is the union of supports, and appending this column to M does not change the tropical rank since the tropical convex hull of the columns remains the same. Therefore, if there is a chain of length r in the support poset we may assume that there is a set of r columns with supports contained in one another. Since there is no column of ones, from this we can extract an $r \times r$ minor with zeroes on and below the diagonal and 1's above the diagonal, which is tropically non-singular.

Reciprocally, suppose there is a tropically non-singular $r \times r$ minor N. We claim that the support poset of N has a chain of length r, from which it follows that the support poset of M also has a chain of length r. Assume without loss of generality that the unique minimum permutation sum is obtained in the diagonal. This minimum sum cannot be more than one, because if n_{ii} and n_{jj} are both 1 then changing them for n_{ij} and n_{ji} does not increase the sum. If the minimum is zero, orienting an edge from i to j if entry ij of Nis zero yields an acyclic digraph, which admits an ordering. Rearranging the rows and columns according to this ordering yields a matrix with 1's above the diagonal and 0's on and below the diagonal. The tropical sum of the last i columns (which corresponds to union of the corresponding supports) then produces a vector with 0's exactly in the last i positions. Hence, there is a proper chain of supports of length r. If the minimum permutation sum in N is 1, then let n_{ii} be the unique diagonal entry equal to 1. The *i*-th row in N must consist of all 1's: if n_{ij} is zero, then changing n_{ij} and n_{ji} for n_{ii} and n_{jj} does not increase the sum. Changing this row of ones to a row of zeroes does not affect the support poset of N, and it yields a non-singular zero-one matrix with minimum sum zero to which we can apply the argument in the previous paragraph.

Example 5.4.19. The tropical rank of the matrix C_n in Theorem 5.4.9 equals two, since all its 3×3 minors are tropically singular, while the principal 2×2 minors are not. The supports of its columns are all the sets of cardinality n-1 and the support poset consists of them and the whole set $\{1, \ldots, n\}$. The maximal chains in the poset have indeed length two.

One of the important properties of rank in usual linear algebra is that it produces a matroid. Unfortunately, the definitions of tropical rank, Kapranov rank, and Barvinok rank all fail to do this.

Example 5.4.20. Consider the configuration of four points in the tropical plane \mathbb{TP}^2 given by the columns of

$$M = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 2 \\ 1 & 0 & 0 & -1 \end{pmatrix}.$$

By any of our three definitions of rank, the maximal independent sets of columns are $\{1, 2\}$, $\{1, 3, 4\}$, and $\{2, 3, 4\}$. These do not all have the same size, and so they cannot be the bases of a matroid. The obstruction here is that the sets $\{1, 2, 3\}$ and $\{1, 2, 4\}$ are collinear, but $\{1, 2, 3, 4\}$ is not. \Box

Despite this failure, there is a strong connection between tropical linear algebra and matroids. This allows us to construct matrices whose tropical and Kapranov ranks disagree. The smallest example we know is 7×7 . It is based on the *Fano matroid*. To explain our construction, we need the following definitions. Let M be a matroid. The *cocircuit matrix of* M, denoted $\mathbf{C}(M)$, has rows indexed by the elements of the ground set of M and columns indexed by the cocircuits of M, that is, the circuits of the dual matroid M^* . The matrix $\mathbf{C}(M)$ has a 0 in entry (i, j) if the *i*-th element is in the *j*-th cocircuit and a 1 otherwise.

In other words, $\mathbf{C}(M)$ is the zero-one matrix whose columns have the cocircuits of M as supports. (Here, the support of a column is its set of

zeroes.) As an example, the matrix C_n in Theorem 5.4.9 is the cocircuit matrix of the uniform matroid of rank 2 with n elements. Similarly, the cocircuit matrix of the uniform matroid $U_{n,r}$ has size $n \times \binom{n}{r-1}$ and its columns are all the zero-one vectors with exactly r-1 ones. The following results show that its tropical and Kapranov ranks equal r. The tropical polytopes associated with these matrices are known as *tropical hypersimplices*.

Proposition 5.4.21. The tropical rank of the cocircuit matrix C(M) is the rank of the matroid M.

Proof. This is a special case of Proposition 5.4.18 because the rank of M is the maximum length of a chain of non-zero covectors, and the supports of covectors are precisely the unions of supports of cocircuits. Note that $\mathbf{C}(M)$ cannot have a column of ones because every cocircuit is non-empty. \Box

Theorem 5.4.22. The Kapranov rank of C(M) is equal to the rank of M if and only if the matroid M is realizable over the field K.

Proof. Let M be a matroid of rank r on $\{1, \ldots, d\}$ which has n cocircuits. We first prove the only if direction. Suppose that $F \in K^{d \times n}$ is a rank r lift of the cocircuit matrix $\mathbf{C}(M)$. For each row f_i of F, let $v_i \in k^d$ be the vector of constant terms in $f_i \in K^d$. We claim that $V = \{v_1, \ldots, v_d\}$ is a representation of M. Here k is the residue field of K, which is also algebraically closed of characteristic zero. First note that V has rank at most r since every K-linear relation among the vectors f_i translates into a k-linear relation among the v_i . Our claim says that $\{i_1, \ldots, i_r\}$ is a basis of M if and only if $\{v_{i_1}, \ldots, v_{i_r}\}$ is a basis of V. Suppose $\{i_1, \ldots, i_r\}$ is a basis of M. Then, as in the proof of Proposition 5.4.18, we can find a square submatrix of $\mathbf{C}(M)$ using rows i_1, \ldots, i_r with 0's on and below the diagonal and 1's above it. This means that the lifted submatrix of constant terms is lower-triangular with nonzero entries along the diagonal. It implies that that v_{i_1}, \ldots, v_{i_r} are linearly independent, and, since rank $(V) \leq r$, they must be a basis. We also conclude rank(V) = r. If $\{i_1, \ldots, i_r\}$ is not a basis in M, there exists a cocircuit containing none of them; this means that some column of $\mathbf{C}(M)$ has all 1's in rows i_1, \ldots, i_r . Therefore, f_{i_1}, \ldots, f_{i_r} all have zero constant term in that coordinate, which means that v_{i_1}, \ldots, v_{i_r} are all 0 in that coordinate. Since the cocircuit is not empty, not all vectors v_j have an entry of 0 in that coordinate, and so $\{v_1, \ldots, v_r\}$ cannot be a basis. This shows that V represents M over k, which proves the only-direction.

5.4. THE RANK OF A MATRIX

For the if-direction, let us assume that M has no loops. This is no loss of generality because a loop corresponds to a row of 1's in $\mathbf{C}(M)$, which does not increase the Kapranov rank because every column has at least a zero. Assume M is representable over k and fix a $d \times n$ -matrix $A \in k^{d \times n}$ such that the rows of A represent M and the sets of non-zero coordinates along the columns of A are the cocircuits of M. Suppose $\{1, \ldots, r\}$ is a basis of M and let A' be the submatrix of A consisting of the first r rows. Write

$$A = \begin{pmatrix} \mathbf{I}_r \\ C \end{pmatrix} \cdot A'$$

where \mathbf{I}_r is the identity matrix and $C \in k^{(d-r) \times r}$. Observe that A, hence C, cannot have a row of zeroes (because M has no loops). Since k is an infinite field, there exists a matrix $B' \in k^{r \times n}$ such that all entries of the $d \times r$ -matrix $\begin{pmatrix} \mathbf{I}_r \\ C \end{pmatrix} \cdot B'$ are non-zero. We now define

$$F = \begin{pmatrix} \mathbf{I}_r \\ C \end{pmatrix} \cdot (A' + tB') \in K^{d \times n}.$$

This matrix has rank r and $\deg(F) = \mathbf{C}(M)$. This completes the proof. \Box

Corollary 5.4.23. Let M be a matroid which is not representable over a given field k. Then the Kapranov rank with respect to k of the tropical matrix C(M) exceeds its tropical rank.

This corollary furnishes many examples of matrices whose Kapranov rank exceeds their tropical rank. Consider, for example, the Fano and non-Fano matroids, depicted in Figure 5.7. They both have rank three and seven

Figure 5.7: The Fano (left) and non-Fano (right) matroids.

elements. The first is only representable over fields of characteristic two, the second only over fields of characteristic different from two. In particular, Corollary 5.4.23 applied to these two matroids implies that over every field there are matrices with tropical rank equal to three and Kapranov rank larger than that. Also, it shows that the Kapranov rank of a matrix may be different over two different algebraically closed fields k and k'.

More explicitly, the cocircuit matrix of the Fano matroid is

$$\mathbf{C}(M) = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

This matrix is the smallest known example of a matrix whose Kapranov rank over \mathbb{C} (four) is strictly larger than its tropical rank (three). Applied to nonrealizable matroids, such as the *Vamos* (rank 4, 8 elements, 41 cocircuits) or the *non-Pappus matroid* (rank 3, 9 elements, 20 cocircuits), Corollary 5.4.23 yields matrices with different Kapranov and tropical ranks over *every* field. One can also get examples in which the difference of the two ranks is arbitrarily large. Indeed, given matrices A and B, we can construct the matrix

$$M := \left(\begin{array}{cc} A & \infty \\ \infty' & B \end{array}\right)$$

where ∞ and ∞' denote matrices of the appropriate dimensions and whose entries are sufficiently large. Appropriate choices of these large values will ensure that the tropical and Kapranov ranks of M are the sums of those of A and of B. The difference between the Kapranov and tropical ranks of Mis equal to the sum of this difference for A and for B.

The construction in Theorem 5.4.22 is closely related to the tropical linear space of the matroid M. Theorem 5.2.6 shows that trop(M) is triangulated by the order complex of the lattice of flats of M. Since flats correspond to unions of cocircuits, the following result is easily derived:

Proposition 5.4.24. The tropical linear space trop(M) of the matroid M is equal to the tropical convex hull of the rows of the modified cocircuit matrix C'(M), where the 1's in C(M) are replaced by ∞ 's.

The set of all tropical linear combinations of a set $V = \{v_1, \ldots, v_n\} \subseteq \mathbb{R}^d$ determines a tropical polytope $\operatorname{tconv}(V)$ in $\mathbb{TP}^{d-1} = \mathbb{R}^d/\mathbb{R}(1,\ldots,1)$. For each sequence $S = (S_1, \ldots, S_d)$ of subsets $S_i \subseteq \{1, \ldots, n\}$, we denote by X_S the polytrope that is defined by the inequalities $x_k - x_j \leq v_{ik} - v_{ij}$ for all $k \in \{1, \ldots, d\}$ and i, j with $i \in S_j$. According to Theorem 5.3.16, the tropical

170

convex hull of V equals the union of the X_S which are bounded. The sequence S is the type of a given point $x \in X_S$ in the tropical polytope $\operatorname{tconv}(V)$. The dimension of a cell X_S of $\operatorname{tconv}(V)$ computed from the combinatorics of the set S: let G_S be the graph which has vertex set $1, \ldots, d$, with i and j connected by an edge if $S_i \cap S_j$ is nonempty. The dimension of X_S is one less than the number of connected components of the graph G_S .

Recall from Definition 5.4.3 that the tropical rank of a matrix is the size of the largest non-singular square minor, and that an $r \times r$ matrix M is non-singular if $\bigcirc_{i=1}^{r} M_{\sigma(i),i} = \sum_{i=1}^{r} M_{\sigma(i),i}$ achieves its minimum twice as σ ranges over the symmetric group S_r . Here is another characterization.

Theorem 5.4.25. Let $M \subset \mathbb{R}^{d \times n}$ be a matrix. Then the tropical rank of M is equal to one plus the dimension of the tropical convex hull of the columns of M, viewed as a tropical polytope in \mathbb{TP}^{d-1} .

Proof. Let $V = \{v_1, \ldots, v_n\}$ be the set of columns of M, and let $P = \operatorname{tconv}(V)$ be their tropical convex hull in \mathbb{TP}^{d-1} . Suppose that r is the tropical rank of M, that is, there exists a tropically non-singular $r \times r$ -submatrix of M, but all larger square submatrices are tropically singular.

We first show that $\dim(P) \ge r-1$. We fix a non-singular $r \times r$ -submatrix M' of M. Deleting the rows outside M' means projecting P into \mathbb{TP}^{r-1} , and deleting the columns outside M' means passing to a tropical subpolytope P' of the image. Both operations can only decrease the dimension, so it suffices to show $\dim(P') \ge r-1$. Hence, we can assume that M is itself a tropically non-singular $r \times r$ -matrix. Also, without loss of generality, we can assume that the minimum over $\sigma \in S_r$ of

$$f(\sigma) = \sum_{i=1}^{r} v_{\sigma(i),i}$$
(5.25)

is uniquely achieved when σ is the identity element $e \in S_r$. We now claim that the cell $X_{(\{1\},\ldots,\{r\})}$ in P = P' is of dimension r - 1. The inequalities defining this cell are $x_k - x_j \leq v_{jk} - v_{jj}$ for $j \neq k$. Suppose that this cell were not full-dimensional. By Farkas' Lemma, there would exist a non-negative linear combination of the inequalities $x_k - x_j \leq v_{jk} - v_{jj}$ which equals $0 \leq c$ for some non-positive real c. This linear combination would imply that some other $\sigma \in S_r$ has $f(\sigma) \leq f(e)$, a contradiction.

For the converse, suppose that $\dim(P) \ge r$. Pick a region X_S of dimension r. The graph G_S has r+1 connected components, so we can pick r+1

elements of $\{1, \ldots, n\}$ of which no two appear in a common S_j . Assume without loss of generality that this set is $\{1, \ldots, r+1\}$, so that $i \in S_j$ if and only if i = j, for $1 \le i, j \le r+1$. We claim that the square submatrix consisting of the first r+1 rows and columns of M is tropically non-singular. Indeed, we have:

$$f(\sigma) - f(e) = \sum_{i=1}^{r+1} v_{\sigma(i),i} - \sum_{i=1}^{r+1} v_{ii} = \sum_{i=1}^{r+1} (v_{\sigma(i),i} - v_{ii}),$$

but whenever $\sigma(i) \neq i$, $v_{\sigma(i),i} - v_{ii} > 0$ since $i \in S_i$ and $i \notin S_{\sigma(i)}$. Therefore, if σ is not the identity, we have $f(\sigma) - f(e) > 0$, and e is the unique permutation in S_{r+1} minimizing the expression (5.25). So M has tropical rank at least r+1. This is a contradiction, and we conclude $\dim(P) = r - 1$. \Box

A useful tool in tropical convexity is the computation of tropical polytopes via mixed subdivisions of the Minkowski sum of several copies of a simplex.

Definition 5.4.26. Let Δ^{d-1} be the standard (d-1)-simplex in \mathbb{R}^d , with vertex set $A = \{e_1, \ldots, e_d\}$. Let $n\Delta^{d-1}$ denote its dilation by a factor of n, which we regard as the convex hull of the Minkowski sum $A + A + \cdots + A$. Let $M = (v_{ij}) \subset \mathbb{R}^{d \times n}$ be a matrix. Consider the lifted simplices

 $P_i := \operatorname{conv}\{(e_1, v_{1i}), \dots, (e_d, v_{di})\} \subset \mathbb{R}^{d+1} \quad \text{for } i = 1, 2, \dots, n.$

The regular mixed subdivision of $n\Delta^{d-1}$ induced by M is the set of projections of the lower faces of the Minkowski sum $P_1 + \cdots + P_n$. Here, a face is called lower if its outer normal cone contains a vector with last coordinate negative.

There is natural a bijection between the cells X_S in the convex hull of the columns of M and the interior cells in the regular subdivision of a product of simplices induced by M. Via the Cayley trick, the latter biject to interior cells in the regular mixed subdivision defined above. Here we provide a short direct proof of the composition of these two bijections:

Lemma 5.4.27. Let $M \subset \mathbb{R}^{d \times n}$ and let $S = (S_1, \ldots, S_d)$, where each S_j is a subset of $\{1, \ldots, n\}$. Then, the following properties are equivalent:

1. The tropical convex hull in \mathbb{TP}^{d-1} of the columns of M contains a cell of type S.

5.4. THE RANK OF A MATRIX

- 2. There is a non-negative matrix M' such that M' is obtained from M by adding constants to rows or columns of M, and such that $M'_{ji} = 0$ precisely when $i \in S_j$.
- 3. The regular mixed subdivision of $n\Delta^{d-1}$ induced by M has as a cell the Minkowski sum $\tau_1 + \cdots + \tau_n$ where $\tau_i = \operatorname{conv}(\{e_j : i \in S_j\})$.

Moreover, the cells in (1) and (3) have complementary dimensions.

Proof. Adding a constant to a row of M amounts to translating the set of n points in \mathbb{TP}^{n-1} , while adding a constant to a column leaves the point set unchanged. Consider a cell X_S in the tropical convex hull, let x be any point in the relative interior of X_S and let M' be the (unique) matrix obtained by translating the point set by a vector -x and normalizing every column by adding a scalar so that its minimum coordinate equals 0. Conversely, for a matrix M' as in (2), consider the point x whose coordinates are the amounts added to the columns of M to obtain M'. The point x is in the tropical convex hull of the columns of M. Let S be its type. Then the modified matrix M' has zeroes precisely in entries (j, i) with $i \in S_j$, proving the equivalence of (1) and (2).

For the equivalence of (2) and (3), observe that adding a constant to a row or column of M does not change the mixed subdivision of $\sum P_i$. For a non-negative matrix M' with at least a zero in every column, the positions of the zero entries define the face of $\sum P_i$ in the negative vertical direction. Conversely, for every cell of the regular mixed subdivision, we can apply a linear transformation to give that cell height zero and all other vertices positive height (this is what it means to be in the lower envelope.) The resulting height function is precisely the matrix M' in (2), which proves the equivalence of (2) and (3). The assertion on dimensions is easy to prove. \Box

Corollary 5.4.28. Given a matrix M, the poset of types in the tropical convex hull of its columns and the poset of interior cells of the corresponding regular mixed subdivision are antiisomorphic.

Corollary 5.4.29. Let $M \subset \mathbb{R}^{n \times d}$. The tropical rank of M equals d minus the minimal dimension of an interior cell in the regular mixed subdivision of $n\Delta^{d-1}$ induced by M.

We can use these tools to prove that the tropical and Kapranov ranks of a matrix coincide if the latter is maximal. **Theorem 5.4.30.** If an $n \times d$ matrix M has Kapranov rank d then it has tropical rank d, too.

Proof. By Corollary 5.4.29, M has tropical rank d if and only if the corresponding regular mixed subdivision has an interior vertex. The theorem then follows from the next two lemmas.

Lemma 5.4.31. A $d \times n$ -matrix M has Kapranov rank less than d if and only if the corresponding regular mixed subdivision has a cell that intersects all facets of $n\Delta^{d-1}$.

Proof. If M has Kapranov rank less than d, then its column vectors lie in a tropical hyperplane. Since all tropical hyperplanes are translates of one another, there is no loss of generality in assuming that it is the hyperplane defined by $x_1 \oplus \cdots \oplus x_d$. That is, after normalization, all columns of M are non-negative and have at least two zeroes. Then, by Lemma 5.4.27, the zero entries of M define a cell B in the regular mixed subdivision none of whose Minkowski summands are single vertices. In particular, for every facet F of Δ^{d-1} and for every $i \in \{1, \ldots, n\}$, the *i*-th summand of B is at least an edge and hence it intersects F. Hence, B intersects all facets of $n\Delta^{d-1}$. For the converse suppose the regular mixed subdivision has a cell B which intersects all facets of $n\Delta^{d-1}$. We may assume that M gives height zero to the points in that cell and positive height to all the others. The intersection of B with the *j*-th facet is given by the zero entries in M after deletion of the *j*-th row. In particular, B intersects the *j*-th facet if and only if every column has a zero entry outside of the *j*-th row, and so B intersects all facets if and only if all columns of M have at least two zeroes, implying that these all lie in the hyperplane defined by $x_1 \oplus \cdots \oplus x_d$.

The cell in the preceding statement need not be unique. For example, if a tetrahedron is sliced by planes parallel to two opposite edges, then each maximal cell meets all the facets of the tetrahedron.

Lemma 5.4.32. In every polyhedral subdivision of a simplex which has no interior vertices, but arbitrarily many vertices on the boundary, there is a cell that intersects all of the facets.

Proof. Observe that there is no loss of generality in assuming that the polyhedral subdivision S is a triangulation. For a triangulation, we use Sperner's Lemma which states the following: If the vertices of a triangulation of Δ

are labeled so that (1) the vertices of Δ receive different labels and (2) the vertices in any face F of Δ receive labels among those of the vertices of F, then there is a fully labeled simplex. Our task is to give our triangulation a Sperner labeling with the property that every vertex labeled i lies in the i-th facet of the simplex. The way to obtain this is: the vertex opposite to facet i is labeled i + 1. More generally, the label i of a vertex v is taken so that v is contained in facet i but not on facet i - 1. All labels are modulo d.

By Theorem 5.4.25, if a matrix has tropical rank two, then the tropical convex hull of its columns is one-dimensional. Being contractible, this tropical polytope is a tree. Another way of showing this is via the corresponding regular mixed subdivision. Tropical rank 2 means that all the interior cells have codimension zero or one. Hence, the subdivision is constructed by slicing the simplex via a certain number of hyperplanes (which do not meet inside the simplex) and its dual graph is a tree. The special case when the matrix has Barvinok rank two is characterized by the following proposition.

Proposition 5.4.33. The following are equivalent for a matrix M:

- 1. It has Barvinok rank 2.
- 2. All its 3×3 minors have Barvinok rank 2.
- 3. The tropical convex hull of its columns is a path.

Proof. $(1) \Rightarrow (2)$ is trivial (the Barvinok rank of a minor cannot exceed that of the whole matrix) and $(3) \Rightarrow (1)$ is easy: if a tropical polytope is a path, then it is the tropical convex hull of its two endpoints. Proposition 5.4.5 then implies that the Barvinok rank is two.

For $(2) \Rightarrow (3)$ first observe that the case where M is 3×3 again follows from Proposition 5.4.5. We next prove the case where M is $d \times 3$ by contradiction. Since the tropical convex hulls of rows and of columns of a matrix are isomorphic as cell complexes, by Theorem 5.3.25, we can assume that the tropical convex hull of the *rows* of M is not a path. Then, there are three rows whose tropical convex hull is not a path, and their 3×3 minor has Barvinok rank 3. Finally, if M is of arbitrary size $d \times n$ and the tropical convex hull of its columns is not a path, consider three columns whose tropical convex hull is not a path and apply the previous case to them. \Box

Our goal is to show that if M has tropical rank 2 then it has Kapranov rank 2. We do this by constructing an explicit lift to a rank 2 matrix over K.

Lemma 5.4.34. Let M be a matrix of tropical rank two. Let x be a point in the tropical convex hull of the columns of M. Let M' be the matrix obtained by adding -x to every column and then normalizing columns to have zero as their minimal entry. After possibly reordering the rows and columns, M' has the following block structure:

$$M' := \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & A_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & A_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & A_k \end{pmatrix},$$

where the matrices A_i have all entries positive and every 2×2 minor has the property that the minimum of its four entries is achieved twice. Each **0** represents a matrix of zeroes of the appropriate size, and the first row and column blocks of M' may have size zero. Moreover, the tropical convex hull of the columns of M' is the union of the tropical convex hulls of the column vectors of the blocks augmented by the zero vector **0**, and two of these k trees meet only at the point **0**.

Proof. First, adjoin the column x to our matrix if it does not already exist; since x is in the convex hull of M, this will not change the tropical convex hull of the columns of M. We can then simply remove it at the end, when it is transformed into a column of all zeroes. Thus, we can assume that one of the columns of the matrix M' consists of all zeroes.

The asserted block decomposition means that any two given columns of M' have either equal or disjoint cosupports, where the cosupport of a column is the set of positions where it does not have a zero. To prove this, we observe that if it didn't then M' would have the following minor, where + denotes a strictly positive entry. (Recall that each column has a zero in it.)

$$\left(\begin{array}{ccc} 0 & + & + \\ 0 & 0 & + \\ 0 & ? & 0 \end{array}\right)$$

But this 3×3 -matrix is tropically non-singular. The assertion of the 2×2 minors follows from the fact that the non-negative matrix

$$\left(\begin{array}{rrrr} 0 & a & b \\ 0 & c & d \\ 0 & 0 & 0 \end{array}\right)$$

5.4. THE RANK OF A MATRIX

is tropically singular if and only if the minimum of a, b, c, d is achieved twice.

Finally, the assertion about the convex hulls is trivial, since any linear combination of column vectors from a given block will have all zero entries except in the coordinates corresponding to that block. Any path joining two such points from different blocks will pass through the origin. \Box

We next introduce a technical lemma for perturbing a power series lifting.

Lemma 5.4.35. Let A be a non-negative matrix with no zero column and suppose that the smallest entry in A occurs the most times in the first column. Let \tilde{A} be the matrix

$$\left(\begin{array}{cc} 0 & \mathbf{0} \\ \mathbf{0} & A \end{array}\right)$$

obtained by adjoining a row and a column of zeroes. If \tilde{A} has Kapranov rank two, then \tilde{A} has a rank-2 lift $F \in K^{d \times n}$ in which every 2×2 minor is non-singular and the *i*-th column can be written as a linear combination $\lambda_i u_1 + \mu_i u_2$ of the first two columns u_1 and u_2 , with $\deg(\lambda_i) \ge \deg(\mu_i) = 0$.

Proof. Starting with an arbitrary rank-2 lift, let F be obtained by adding to every row/column a K-linear combination of all other rows/columns with coefficients of sufficiently high degree but otherwise generic. This preserves the rank 2 of the lift and the degree of every entry, but makes every 2×2 minor of F non-singular. The *i*-th column of F is now a K-linear combination $\lambda_i u_1 + \mu_i u_2$ of the first two columns. If the degrees of λ_i and μ_i are different, then their minimum must be zero in order to get a degree zero element in the first entry of column *i*. But then deg(μ_i) > deg(λ_i) = 0 is impossible, because it would make the *i*-th column of A all zero. Hence deg(λ_i) > deg(μ_i) = 0.

If the degrees are equal, then they are non-positive in order to get degree zero for the first entry in $\lambda_i u_1 + \mu_i u_2$. But they cannot be equal and negative, or otherwise entries of positive degree in u_2 would produce entries of negative degree in u_i . Hence, $\deg(\lambda_i) = \deg(\mu_i) = 0$ in this case.

Corollary 5.4.36. Let A and B be non-negative matrices. Assume that

$$\tilde{A} := \begin{pmatrix} A & \mathbf{0} \\ \mathbf{0} & 0 \end{pmatrix} \quad and \quad \tilde{B} := \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{0} & B \end{pmatrix}$$

have Kapranov rank equal to 2. Then, the matrix

$$M := \left(\begin{array}{ccc} A & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & B \end{array} \right)$$

has Kapranov rank equal to 2 as well.

Proof. We may assume that neither A nor B has a zero column. Hence Lemma 5.4.35 applies to both of them. We number the rows of M from -kto k' and its columns from -l to l', where $k \times l$ and $k' \times l'$ are the dimensions of A and B respectively. In this way, A (respectively B) is the minor of negative (respectively, positive) indices. The row and column indexed zero consist of all zeroes. To further exhibit the symmetry between A and B the columns and rows in \tilde{A} will be referred to "in reverse". That is to say, the first and second columns of it are the ones indexed 0 and -1 in M.

We now construct a lifting $F = (a_{i,j}) \in k\{\{t\}\}^{d \times n}$ of M. We assume that we are given rank-2 lifts of \tilde{A} and \tilde{B} which satisfy the conditions of Lemma 5.4.35. Furthermore, we assume that the lift of the entry (0,0) is the same in both, which can be achieved by scaling the first row in one of them.

We use exactly those lifts of A and B for the upper-left and bottom-right corner minors of M. Our task is to complete that with an entry $a_{i,j}$ for every i, j with ij < 0, such that $\deg(a_{i,j}) = 0$ and the whole matrix still has rank 2. We claim that it suffices to choose the entry $a_{-1,1}$ of degree zero and sufficiently generic. That this choice fixes the rest of the matrix is easy to see: The entry $a_{1,-1}$ is fixed by the fact that the 3×3 minor

$$\left(\begin{array}{cccc} a_{-1,-1} & a_{-1,0} & a_{-1,1} \\ a_{0,-1} & a_{0,0} & a_{0,1} \\ a_{1,-1} & a_{1,0} & a_{1,1} \end{array}\right)$$

needs to have rank 2. All other entries $a_{i,-1}$ and $a_{i,1}$ are fixed by the fact that the entries $a_{i,-1}$, $a_{i,0}$ and $a_{i,1}$ (two of which come from either \tilde{A} or \tilde{B}) must satisfy the same dependence as the three columns of the minor above. For each $j = -l, \ldots, -2$ (respectively $j = 2, \ldots, l'$), let λ_j and μ_j be the coefficients in the expression of the *j*-th column of \tilde{A} (respectively, of \tilde{B}) as $\lambda_j u_0 + \mu_j u_{-1}$ (respectively, $\lambda_j u_0 + \mu_j u_1$). Then, $a_{i,j} = \lambda_j a_{i,0} + \mu_j a_{i,-1}$ (respectively, $a_{i,j} = \lambda_j a_{i,0} + \mu_j a_{i,1}$).

What remains to be shown is that if $a_{-1,1}$ is of degree zero and sufficiently generic then all the new entries are of degree zero too. For this, observe that if $j \in \{-l', \ldots, 2\}$ (resp. $j \in \{2, \ldots, l\}$ then $a_{i,j}$ is of degree zero as long as the coefficient of degree zero in $a_{i,-1}$ (resp. $a_{i,1}$) are different from the degree zero coefficients in the quotient $-\lambda_j a_{i,0}/\mu_j$ (here we are using the assumption that $\deg(\lambda_j) \ge \deg(\mu_j) \ge 0$). In terms of the choice of $a_{-1,1}$ this translates to the following determinant having non-zero coefficient in degree zero:

$$\begin{pmatrix} a_{i,-1} & a_{i,0} & -\lambda_j a_{i,0}/\mu_j \\ a_{-1,-1} & a_{-1,0} & a_{-1,1} \\ a_{0,-1} & a_{0,0} & a_{0,1} \end{pmatrix}, \quad (\text{respectively} \quad \begin{pmatrix} a_{0,-1} & a_{0,0} & a_{0,1} \\ a_{1,-1} & a_{1,0} & a_{1,1} \\ -\lambda_j a_{i,0}/\mu_j & a_{i,0} & a_{i,1} \end{pmatrix}).$$

That $a_{-1,1}$ and $a_{1,-1}$ sufficiently generic imply non-singularity of these matrices follows from the fact that the 2 × 2 minors

$$\left(\begin{array}{cc} a_{i,-1} & a_{i,0} \\ a_{0,-1} & a_{0,0} \end{array}\right), \qquad \left(\begin{array}{cc} a_{0,0} & a_{0,1} \\ a_{i,0} & a_{i,1} \end{array}\right).$$

come from the given lifts of \tilde{A} and \tilde{B} . Hence they are non-singular.

Theorem 5.4.37. Let M be a matrix of tropical rank 2. Then its Kapranov rank equals 2 as well.

Proof. The Kapranov rank of M is always at least the tropical rank, so we need only show that the Kapranov rank is less than or equal to 2. If the tropical convex hull P of the columns of M is a path, then M has Barvinok rank 2 (by Proposition 5.4.33) and thus Kapranov rank 2. Otherwise, let xbe a node of degree at least three in the tree P. We apply the method of Lemma 5.4.34. Since x has degree at least three, it follows that there are at least three blocks A_i . In particular, M has at least three columns. We induct on the number of columns of M. If M has exactly three columns, then each block A_i is a single column, and every row of M has at most one positive entry. It is easy to construct an explicit lift of rank 2: in each row, lift the positive entry α as $-t^{\alpha}$ and the zero entries as -1 and $1 + t^{\alpha}$. If there are rows of zeroes, lift them as (-1, -1, 2), for example.

Next, suppose that M has $m \ge 4$ columns. The two blocks with the fewest number of combined columns have at least 2 and at most m-2 rows all together. Possibly after adding a row and column of zeroes, this provides a decomposition of our matrix as

$$M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & B \end{pmatrix},$$

where both A and B have at least two columns. It follows that the minors

$$\left(\begin{array}{cc} 0 & \mathbf{0} \\ \mathbf{0} & A \end{array}\right) \qquad \text{and} \qquad \left(\begin{array}{cc} 0 & \mathbf{0} \\ \mathbf{0} & B \end{array}\right)$$

both have fewer columns than the original matrix. By the inductive hypothesis they have Kapranov rank 2. Applying Corollary 5.4.36 completes the inductive step of the theorem. $\hfill \Box$

5.5 Dressians

Our aim in this last section is to finally define tropical linear spaces. We introduce $\binom{n}{d}$ unknowns q_{σ} over the tropical semiring that are indexed by the *d*-sets $\sigma \in \binom{[n]}{d}$. These will serve as coordinates on the compactified tropical projective space $\overline{\mathbb{TP}}^{\binom{n}{d}-1}$. For any $\rho \in \binom{[n]}{d-1}$ and any $\tau \in \binom{[n]}{d+1}$ we consider the following quadratic tropical polynomial

$$\bigoplus_{i \in \tau \setminus \rho} q_{\tau \setminus \{i\}} \odot q_{\rho \cup \{i\}} \tag{5.26}$$

Each of these quadrics defines a tropical hypersurface in $\mathbb{TP}^{\binom{n}{d}-1}$. The intersection of these hypersurfaces is a tropical prevariety. This prevariety is denoted $\mathbb{Dr}(d, n)$ is called the *Dressian*. This name refers to Andreas Dress, who studied the elements of $\mathbb{Dr}(d, n)$ under the name of *valuated matroids*.

Fix a point $q \in \mathbb{D}r(d, n) \subset \overline{\mathbb{TP}}^{\binom{n}{d}-1}$. With any $\tau \in \binom{[n]}{d+1}$ we associate a tropical hyperplane $L_{\tau}(q)$ in \mathbb{R}^n . It is defined by the tropical linear form

$$\bigoplus_{i\in\tau} q_{\tau\setminus\{i\}} \odot w_i. \tag{5.27}$$

We finally consider the intersection of these tropical hyperplanes in \mathbb{R}^n :

$$L(q) := \bigcap_{\tau \in \binom{[n]}{d+1}} L_{\tau}(q).$$

This definition makes sense for any points $w \in \overline{\mathbb{TP}}^{\binom{n}{d}-1}$, and the set L(q) will always be a tropical prevariety in \mathbb{R}^n . However, this prevariety behaves like a *d*-dimensional linear space precisely when *w* comes from the Dressian.

Definition 5.5.1. A *d*-dimensional tropical linear space in \mathbb{R}^n to be any prevariety of the form L(q) where w is any point in the Dressian $\mathbb{D}r(d, n)$.

This definition is justified by the following three results.
Theorem 5.5.2. Every tropicalized linear space is a tropical linear space.

For tropicalized linear spaces, the desirable properties in our next theorem can be derived from the Fundamental Theorem and the Structure Theorem for tropical varieties. However, we shall see that these hold more generally.

Theorem 5.5.3. Every tropical linear space L(q) is a pure d-dimensional balanced polyhedral complex in \mathbb{R}^n . The image of L(q) in \mathbb{TP}^{n-1} is a contractible complex. Moreover, that image is a tropical cycle of degree one, which means that, for any generic point $p \in \mathbb{R}^n$, it intersects the complementary linear space $p + \operatorname{trop}(U_{n-d,n})$ transversally in precisely one point.

Finally, we need to show that Definition 5.5.1 is consistent with Definition 5.2.5, namely, that it includes the tropical linear spaces arising from matroids.

Theorem 5.5.4. If M is a matroid then trop(M) is a tropical linear space.

We begin by proving the first of the three theorems.

Proof of Theorem 5.5.2. Let X be a linear subspace of dimension d in K^n , and fix a matrix $A \in K^{d \times n}$ whose row space equals X. If σ is an ordered d-tuple from [n] then we write p_{σ} for the $d \times d$ -subdeterminant of the matrix A. This defines a vector $p \in K^{\binom{[n]}{d}}$ with coordinates p_{σ} . We write $q \in \mathbb{R}^{\binom{[n]}{d}}$ for the corresponding tropical vector with coordinates $q_{\sigma} = \operatorname{val}(p_{\sigma})$. The *circuits* of X are the non-zero linear forms in I_X that have minimal support. These have now coefficients in K, so, the valuations of these coefficients will usually be non-zero. Nevertheless, the argument of Proposition 5.2.2 remains valid in this case, and we can conclude that the circuits form a tropical basis of the linear ideal I_X . The tropicalization of the circuits are precisely the tropical linear forms (5.27). Therefore, $L(q) = \operatorname{trop}(X)$ as desired.

To complete the proof we need to show that q lies in the Dressian $\mathbb{D}r(d, n)$. But this follows from the fact that the vector p of maximal minors p_{σ} satisfies certain quadratic equations known as the *quadratic Plücker relations*:

$$\sum_{i \in \tau \setminus \rho} (\pm 1) \cdot p_{\tau \setminus \{i\}} \cdot p_{\rho \cup \{i\}} = 0$$

We shall return to these quadratic Plücker relations in more detail in our later section on Grassmannians. The only point we need here is that their tropicalizations are precisely relations (5.26), and this shows that $q \in \mathbb{D}r(d, n)$. Next, we establish the connections to matroids. To this end, we state yet one more axiom system for matroids, namely the *basis exchange axiom*.

Definition 5.5.5. A matroid is a pair $M = (E, \mathcal{B})$ where E is a finite set and \mathcal{B} is a collection of subsets of E, called the *bases* of M, that satisfies the following property: whenever σ and σ' are bases and $i \in \sigma \setminus \sigma'$ then there exists an element $j \in \sigma' \setminus \sigma$ such that $(\sigma \setminus \{i\}) \cup \{j\}$ is a basis as well.

Let $\mathcal{B} \subset {\binom{[n]}{d}}$ be any collection of *d*-subsets of [n]. We represent \mathcal{B} by its tropical incidence vector $q(\mathcal{B}) \in \{0,\infty\}^{\binom{[n]}{d}}$, which is defined by setting $q(\mathcal{B})_{\sigma} = 0$ if $\sigma \in \mathcal{B}$ and $q(\mathcal{B})_{\sigma} = \infty$ if $\sigma \notin \mathcal{B}$.

Lemma 5.5.6. The vector $q(\mathcal{B})$ lies in the Dressian $\mathbb{D}r(d, n)$ if and only if the collection \mathcal{B} of d-sets is the set of bases of a matroid $M = ([n], \mathcal{B})$.

Proof. The value of each tropical Plücker polynomial is either 0 or ∞ :

$$\bigoplus_{i \in \tau \setminus \rho} q_{\tau \setminus \{i\}}(\mathcal{B}) \odot q_{\rho \cup \{i\}}(\mathcal{B}) \in \{0, \infty\}.$$

This minimum fails to be attained twice precisely when there is only term with value 0. If we set $\sigma = \rho \cup \{i\}$ and $\sigma' = \tau \setminus \{i\}$ for that particular term, then σ and σ' are bases but the basis exchange promised in Definition 5.5.5 is not possible. This argument is reversible, and we see that (5.26) is precisely the basis exchange axiom when applied to vectors q in $\{0, \infty\}^{\binom{[n]}{d}}$.

We can paraphrase this succinctly as follows:

Corollary 5.5.7. The rank d matroids on [n] are the 0- ∞ -vectors in $\mathbb{D}r(d, n)$.

The derivation of the third theorem stated above has now become easy.

Proof of Theorem 5.5.4. Let $M = ([n], \mathcal{B})$ be a matroid of rank d. Any set $\tau \in {[n] \choose d+1}$ that has rank d in M contains a unique circuit C of M, namely, the elements of C are the indices such that $\tau \setminus \{i\}$ is a basis. This allows us to identify the tropical linear equations in (5.27) with the circuits of M:

$$\bigoplus_{i\in\tau} q_{\tau\setminus\{i\}}(\mathcal{B})\odot w_i \quad = \quad \bigoplus_{i\in C} w_i.$$

A vector $w \in \mathbb{R}^n$ lies in the tropical hyperplanes defined by these linear forms, as C runs over all circuits of M, if and only if w lies in the set trop(M) in Definition 5.2.5. This proves the desired identity trop $(M) = L(q(\mathcal{B}))$. \Box We now come to the geometric properties of tropical linear spaces L(q).

Proof of Theorem 5.5.3. Since L(q) is a prevariety, it is a polyhedral complex. Theorem 5.3.3 implies that it is contractible because L(q) is an intersection of tropical hyperplanes and hence is tropically convex. We next show that L(q) is pure of dimension d. This will be accomplished by showing that, locally at each point $w \in L(q)$, the linear space L(q) has the form $\operatorname{trop}(M_{q,w})$ for some rank d matroid $M_{q,w}$. Finally, we will prove that L(q) is a tropical cycle of degree one. These last two points still need to be written.

5.6 Exercises

Toric Connections

The theory of toric varieties is one of the main interfaces between combinatorics and algebraic geometry. In this chapter we will see how the tropical connection between these fields is intimately connected with the toric one.

A toric variety is a variety containing a dense copy of the algebraic torus \mathbb{T}^n with an action of \mathbb{T}^n on it. It decomposes into a union of \mathbb{T}^n -orbits. We first see that given a subvariety Z of a toric variety, tropical geometry answers the question "which torus orbits does Z intersect"?

A normal toric variety is determined by the combinatorial data of a rational polyhedral fan. For $Y \subset \mathbb{T}^n$, a choice of fan structure on $\operatorname{trop}(Y)$ then determines a toric variety with torus \mathbb{T}^n . The closure of Y in this toric variety is then a good choice of compactification of Y. This extends the story begun in Section 1.8. Conversely, a good choice of compactification of $Y \subset \mathbb{T}^n$ leads to a computation of $\operatorname{trop}(Y)$.

Degenerations of Y are also controlled by the tropical variety. We study these in Section 6.5, before turning to the tropical and toric approaches to intersection theory in the last section.

6.1 Toric Background

We assume familiarity with the basics of normal toric varieties as in [Ful93], [Oda88], or [CLS] and just briefly review notation here.

A toric variety is defined by a fan Σ in $N_{\mathbb{R}} = N \otimes \mathbb{R}$ for a lattice $N \cong \mathbb{Z}^n$. We denote by M the dual lattice $M = \text{Hom}(N, \mathbb{Z})$. We will work with toric varieties X_{Σ} defined over \Bbbk . The torus T of X_{Σ} is $N \otimes \Bbbk \cong (\Bbbk^*)^n$. We denote

Figure 6.1:

by $\Sigma(k)$ the set of k-dimensional cones of Σ .

Each cone $\sigma \in \Sigma$ determines a local chart $U_{\sigma} = \operatorname{Spec}(\Bbbk[\sigma^{\vee}])$, where $\sigma^{\vee} = \{\mathbf{u} \in M : \mathbf{u} \cdot \mathbf{v} \geq 0 \text{ for all } \mathbf{v} \in \sigma\}$ is the dual cone. Note that every affine normal toric variety has the form U_{σ} for some cone $\sigma \subset N_{\mathbb{R}}$. The cone σ also determines a *T*-orbit $\mathcal{O}_{\sigma} \cong (\Bbbk^*)^{n-\dim(\sigma)}$. The closure in X_{Σ} of the orbit \mathcal{O}_{σ} is denoted by $V(\sigma)$.

Example 6.1.1. 1. Let Σ have rays $\mathbf{e}_1, \ldots, \mathbf{e}_n, \mathbf{e}_0 = \sum_{i=1}^n \mathbf{e}_i$, and cones generated by any k of these rays for $k \leq n$. The case n = 2 is shown in Figure 6.1. Then $X_{\Sigma} \cong \mathbb{P}^n$.

The orbit corresponding to the cone generated by a subset of the \mathbf{e}_i indexed by a set $\sigma \subset \{0, \ldots, n\}$ is those points $[x_0 : \cdots : x_n] \in \mathbb{P}^n$ with $x_i = 0$ for $i \in \sigma$ and $x_i \neq 0$ for $i \notin \sigma$.

2. Let Σ be the fan in \mathbb{R}^2 with rays (1,0), (1,1), (0,1), and maximal cones pos((1,0), (1,1)) and pos((0,1), (1,1)). Then X_{Σ} is the blow-up of \mathbb{A}^2 at the origin.

The only smooth affine normal toric variety is $\mathbb{k}^d \times \mathbb{k}^{*n-d}$. This corresponds to a *d*-dimensional cone $\sigma \subset N_{\mathbb{R}}$ generated by part of a basis for N. In general a toric variety X_{Σ} is smooth if and only if every cone $\sigma \in \Sigma$ is generated by part of a basis for N. We call such Σ a smooth fan. Resolution of singularities of toric varieties is a combinatorial operation, and works in arbitrary characteristic. Specifically, given any fan Σ , there is a smooth fan $\tilde{\Sigma}$ that refines Σ , and the refinement of fans induces a proper birational map $\pi: X_{\tilde{\Sigma}} \to X_{\Sigma}$. See [Ful93, Section 2.2] for details.

6.2 Subvarieties of Toric Varieties

Let $\mathbb{T}^n = (\mathbb{K}^*)^n$, and let $Y \subseteq \mathbb{T}^n$ be a subvariety of \mathbb{T}^n . Fix a toric variety X_{Σ} , and let \overline{Y} be the closure of Y in X_{Σ} . We emphasize that we do not assume that X_{Σ} is a complete toric variety, so the support $|\Sigma|$ of Σ need not be all of \mathbb{R}^n . The following is a natural question in the context of toric geometry:

Question 6.2.1. Which \mathbb{T}^n -orbits of X_{Σ} does \overline{Y} intersect?



Figure 6.2: The torus orbits intersecting the \overline{Y} of Example 6.2.2

Example 6.2.2. Let $Y = V(x + y + 1) \subset (\mathbb{k}^*)^2$.

- 1. Let $X_{\Sigma} = \mathbb{P}^2$, with the torus $T = \{(x : y : 1) : x, y \in \mathbb{k}^*\}$ and homogeneous coordinates (x : y : z). Then \overline{Y} is the subvariety $\mathbb{V}(x + y + z) \subset \mathbb{P}^2$. Note that $\overline{Y} = Y \cup \{(1 : -1 : 0), (1 : 0 : -1), (0 : 1 : -1)\}$. The closure \overline{Y} thus intersects all *T*-orbits of \mathbb{P}^2 except the torus-fixed points $\{(1 : 0 : 0), (0 : 1 : 0), (0 : 0 : 1)\}$.
- 2. Let $\mathbb{P}_{\Sigma} = \mathbb{P}^1 \times \mathbb{P}^1$, with torus $T = \{(x : 1) \times (y : 1) : x, y \in \mathbb{k}^*\}$, and homogeneous coordinates $(x_1 : y_1) \times (x_2 : y_2)$. Then \overline{Y} is the subvariety of $\mathbb{P}^1 \times \mathbb{P}^1$ defined by the equation $x_1y_2 + x_2y_1 + x_2y_2 = 0$. Thus $\overline{Y} = Y \cup \{(-1 : 1) \times (0 : 1), (0 : 1) \times (-1 : 1), (1 : 0) \times (1 : 0)\}$. The closure \overline{Y} intersects four of the nine torus orbits of $\mathbb{P}^1 \times \mathbb{P}^1$: $T, \{(a : 1) \times (1 : 0) : a \in \mathbb{k}^*\}, \{(1 : 0) \times (a : 1) : a \in \mathbb{k}^*\}, (1 : 0) \times (1 : 0).$ These are illustrated in Figure 6.2.

Perhaps surprisingly, tropical geometry answers Question 6.2.1.

Theorem 6.2.3. Let \overline{Y} be a subvariety of a toric variety X_{Σ} , and let $Y = \overline{Y} \cap T$. Then for $\sigma \in \Sigma$ we have $\overline{Y} \cap O_{\sigma} \neq \emptyset$ if and only if trop(Y) intersects the relative interior of σ .

For the proof we first consider the case where the toric variety X_{Σ} is the product $\mathbb{A}^m_{\Bbbk} \times (\Bbbk^*)^{n-m}$ of an affine space and a torus.

Proposition 6.2.4. Let $\sigma \subset \mathbb{R}^n$ be the cone spanned by the first *m* basis vectors $\mathbf{e}_1, \ldots, \mathbf{e}_m$ of \mathbb{R}^m , so the affine toric variety $U_{\sigma} \cong \mathbb{A}^m \times (\mathbb{k}^*)^{n-m}$. Let $Y \subset T^m$ be a subvariety, and let \overline{Y} be the closure of Y inside U_{σ} . Then $\overline{Y} \cap \{(\mathbf{0}, x) : x \in (\mathbb{k}^*)^{n-m}\} \neq \emptyset$ if and only if $\operatorname{trop}(Y) \cap \operatorname{relint}(\sigma) \neq \emptyset$.

Proof.

6.3 Tropical Compactifications

Let Y be a subvariety of T^m . In the previous section we introduced the closure \overline{Y} of Y in a toric variety X_{Σ} with torus T^m . In this section we discuss the extra properties \overline{Y} has if Σ is chosen more carefully.

This section will have more of the content of [Tev07].

6.4 Geometric Tropicalization

Given a subvariety $Y \subset \mathbb{T}^n$, we saw in the last section how the tropical variety determines a good choice of compactifications of Y. In this section we explore the converse, and see how a sufficiently nice compactification of Y determines trop(Y). References include [HKT07] and [ST08].

6.5 Degenerations

The tropical variety of $Y \subset \mathbb{T}^n$ also determines degenerations of Y, and there is a beautiful interplay between the compactifications of the previous section and these degenerations, as we now explain.

6.6 Intersection theory

In this section we explain how some intersection-theoretic computations on the toric variety can be done tropically, and make the connection with [FS97].

6.7 Exercises

Elimination and Implicitization

Elimination theory is the art of computing the image of a morphism in algebraic geometry. The special case of a morphism from an algebraic torus into an affine space leads to the problem of implicitization. The operations of elimination and implicitization are fundamental also from applied and computational point of view. Tropical geometry provides excellent new tools for both theory and practice. This chapter will describe the successes of tropical geometry in elimination and implicitization, expanding far beyond the example of parametric plane curves in Section 1.5 of the Introduction. This will include material from papers including [DFS07, ST08, STY07, SY08].

7.1 Tropical Maps

A tropical map is a function $\phi : \mathbb{R}^d \to \mathbb{R}^n$ whose coordinates ϕ_i are tropical polynomials. Such a map arises as the tropicalization $\phi = \operatorname{trop}(f)$ of a map $f := (K^*)^d \to K^n$ whose coordinates f_1, \ldots, f_n are Laurent polynomials in $K[x_1^{\pm 1}, \ldots, x_d^{\pm d}]$. Here, $\phi_i = \operatorname{trop}(f_i)$ for $i = 1, 2, \ldots, n$. The following inclusion always holds, but it is usually strict:

 $\operatorname{image}(\operatorname{trop}(f)) \subset \operatorname{trop}(\operatorname{image}(f)).$

The right hand side is a tropical variety in \mathbb{R}^n , specified by the Zariski closure X = image(f), and the right hand side is a polyhedral complex inside it. In this section we study this inclusion, and we characterize the domains of linearity of the tropical map ϕ . Particular emphasis placed on the case when f is the natural parametrization of a secant variety of a toric variety.

7.2 Projections and Tropical Bases

In this section we examine the pushforward-formula for tropical varieties and their multiplicities, and we explain how this can be used to compute the tropicalization of the image of a variety under a monomial map. In practise, this amounts to a linear projection of balanced polyhedral complexes. When the projection is generic, this leads to a theorem of Hept and Theobald [HT09] which states that every Laurent polynomial ideal has a small tropical basis.

7.3 Discriminants and Resultants

Tropical geometry is used to develop a new approach to the theory of discriminants and resultants in the sense of Gelfand, Kapranov and Zelevinsky [GKZ08]. The tropical A-discriminant is the tropicalization of the dual variety of the projective toric variety given by an integer matrix A. This tropical algebraic variety is shown to coincide with the Minkowski sum of the row space of A and the tropicalization of the kernel of A. This leads to an explicit positive formula for all the extreme monomials of any A-discriminant. This section will be developed from the material in [DFS07].

7.4 Mixed Fiber Polytopes

We present the polyhedral approach to elimination developed by Esterov, Khavanski, Sturmfels, Tevelev and Yu. For a Laurent polynomial map with generic coefficients, we give an explicit formula for the tropical variety of the image, along with its multiplicities. When the image is a hypersurface, the output is the Newton polytope of the defining polynomial. This method be used to compute mixed fiber polytopes, including secondary polytopes. This section will be developed from the material in [SY08].

7.5 Parametrized Surfaces

In this section we apply Geometric Tropicalization (Section 6.4) to the problem of finding the equations of a parametrized surface. The primary application is in the case when the surface is embedded in 3-space, and we seek to find its defining polynomial. This section builds on the work in [ST08], it makes a connection to resolution of singularities, and it offers a sneak preview of material to appear in the dissertation of Angelica Cueto.

7.6 Hadamard Products

The Hadamard product $X \star Y$ of two subvarieties of an algebraic torus \mathbb{T}^n is the Zariski closure of the set of coordinatewise products of points in Xwith points in Y. We argue that this operation is important for applications, notably in algebraic statistics, and we show how tropical geometry can be used to compute and study Hadamard products. This is based on the identity

$$\operatorname{trop}(X \star Y) = \operatorname{trop}(X) + \operatorname{trop}(Y).$$

This generalizes Kapranov's Horn uniformization which underlies Section 4.3.

7.7 Exercises

192

Realizability

In Chapter 3 we introduced a tropical variety as the tropicalization of a subvariety of an algebraic torus. We then showed in Theorem 3.3.4 that the tropical variety of an irreducible d-dimensional variety is the support of a pure d-dimensional weighted balanced polyhedral complex. Many authors define a tropical variety to be the support of such a polyhedral complex. We now switch to that definition, and we examine the question which tropical varieties are realizable by classical varieties over a field K with a valuation.

8.1 Hypersurfaces

A tropical hypersurface is a weighted balanced polyhedral complex of pure codimension one in \mathbb{R}^n . In this section we show that every tropical hypersurface is geometrically dual to a regular polyhedral subdivision of a lattice polytope. This implies that every tropical hypersurface is realizable by a classical hypersurface in $\mathbb{T}^n = (K^*)^n$, and the field K does not matter here.

8.2 Matroids

In Chapter 4 we associated a balanced fan with every matroid. In this section we extend this construction to matroid subdivisions of matroid polytopes, and we prove that these tropical linear spaces are precisely the irreducible tropical varieties of degree one. A tropical linear space is realizable if and only if the underlying matroid subdivision is realizable over the algebraic closure of the residue field k. The smallest example of a non-realizable linear space lives in \mathbb{R}^7 and has dimension three. It is based on the Fano matroid.

8.3 Curves

A tropical curve is a balanced weighted graph in \mathbb{R}^n . We show that these curves may fail to be realizable even if n = 3. The question which tropical curves are realizable is a rich area of current research. Complete answers are beginning to emerge but they require deep methods from complex geometry (due to Brugallé and Mikhalkin) or arithmetic geometry (due to Speyer). The purpose of this section is to offer a first introduction to this topic.

8.4 Prevarieties

A tropical prevariety is a finite intersection of tropical hypersurfaces in \mathbb{R}^n . It follows from the existence of finite tropical bases that every tropical variety is a prevariety, but the converse is far from true. Many classical varieties, such as Grassmannians and determinantal varieties, come with nice generating sets for their prime ideal, and it is natural to ask whether these generators form a tropical basis. The answer is negative in general: not all naturally defined prevarieties are tropical varieties. The discrepency between the Grassmannian (Section 4.4) and the Dressian (Section 5.5) is the typical scenario. This issue relates to the realizability problem because non-realizable linear spaces correspond to those points in the Dressian that are not in the Grassmannian.

8.5 Exercises

Further Topics

This chapter will contain some brief vignettes of material that has not received sufficient coverage in the rest of the book. The section titles below give some ideas of topics that might be covered here.

- 9.1 Berkovich Spaces
- 9.2 Abstract Tropical Intersection Theory
- 9.3 Tropical Curves and Riemann-Roch
- 9.4 Tropical Moduli Spaces

196

Bibliography

- [AK06] Federico Ardila and Caroline J. Klivans. The Bergman complex of a matroid and phylogenetic trees. J. Combin. Theory Ser. B, 96(1):38–49, 2006.
- [AZ04] Martin Aigner and Günter M. Ziegler. *Proofs from The Book.* Springer-Verlag, Berlin, third edition, 2004. Including illustrations by Karl H. Hofmann.
- [BCOQ92] François Louis Baccelli, Guy Cohen, Geert Jan Olsder, and Jean-Pierre Quadrat. Synchronization and linearity. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, 1992. An algebra for discrete event systems.
- [Ber71] George M. Bergman. The logarithmic limit-set of an algebraic variety. *Trans. Amer. Math. Soc.*, 157:459–469, 1971.
- [BG84] Robert Bieri and J. R. J. Groves. The geometry of the set of characters induced by valuations. J. Reine Angew. Math., 347:168– 195, 1984.
- [BJWW98] Alexander Barvinok, David S. Johnson, Gerhard J. Woeginger, and Russell Woodroofe. The maximum traveling salesman problem under polyhedral norms. In *Integer programming and combinatorial optimization (Houston, TX, 1998)*, volume 1412 of *Lecture Notes in Comput. Sci.*, pages 195–201. Springer, Berlin, 1998.
- [BNS87] Robert Bieri, Walter D. Neumann, and Ralph Strebel. A geometric invariant of discrete groups. *Invent. Math.*, 90(3):451–477, 1987.

- [BS80] Robert Bieri and Ralph Strebel. Valuations and finitely presented metabelian groups. *Proc. London Math. Soc. (3)*, 41(3):439–464, 1980.
- [CGQ04] Guy Cohen, Stéphane Gaubert, and Jean-Pierre Quadrat. Duality and separation theorems in idempotent semimodules. *Linear Algebra Appl.*, 379:395–422, 2004. Tenth Conference of the International Linear Algebra Society.
- [CLO07] David Cox, John Little, and Donal O'Shea. Ideals, varieties, and algorithms. Undergraduate Texts in Mathematics. Springer, New York, third edition, 2007. An introduction to computational algebraic geometry and commutative algebra.
- [CLS] David Cox, John Little, and Hal Schenck. Toric varieties. Draft of book in progress. Samples available from Cox's webpage.
- [DFS07] Alicia Dickenstein, Eva Maria Feichtner, and Bernd Sturmfels. Tropical discriminants. J. Amer. Math. Soc., 20(4):1111–1133 (electronic), 2007.
- [DSS05] Mike Develin, Francisco Santos, and Bernd Sturmfels. On the rank of a tropical matrix. In *Combinatorial and computational* geometry, volume 52 of Math. Sci. Res. Inst. Publ., pages 213– 242. Cambridge Univ. Press, Cambridge, 2005.
- [Eis95] David Eisenbud. Commutative algebra, volume 150 of Graduate Texts in Mathematics. Springer-Verlag, New York, 1995. With a view toward algebraic geometry.
- [FS97] William Fulton and Bernd Sturmfels. Intersection theory on toric varieties. *Topology*, 36(2):335–353, 1997.
- [FS05] Eva Maria Feichtner and Bernd Sturmfels. Matroid polytopes, nested sets and Bergman fans. *Port. Math.* (N.S.), 62(4):437– 468, 2005.
- [Ful93] William Fulton. Introduction to toric varieties, volume 131 of Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1993. , The William H. Roever Lectures in Geometry.

[GGP97]	Israel M. Gelfand, Mark I. Graev, and Alexander Postnikov. Combinatorics of hypergeometric functions associated with posi- tive roots. In <i>The Arnold-Gelfand mathematical seminars</i> , pages 205–221. Birkhäuser Boston, Boston, MA, 1997.
[GKZ08]	I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. <i>Discriminants, resultants and multidimensional determinants.</i> Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, 2008. Reprint of the 1994 edition.
[GM07a]	Andreas Gathmann and Hannah Markwig. The Caporaso-Harris formula and plane relative Gromov-Witten invariants in tropical geometry. <i>Math. Ann.</i> , 338(4):845–868, 2007.
[GM07b]	Andreas Gathmann and Hannah Markwig. The numbers of trop- ical plane curves through points in general position. J. Reine Angew. Math., 602:155–177, 2007.
[GM08]	Andreas Gathmann and Hannah Markwig. Kontsevich's formula and the WDVV equations in tropical geometry. <i>Adv. Math.</i> , 217(2):537–560, 2008.
[Har77]	Robin Hartshorne. <i>Algebraic geometry</i> . Springer-Verlag, New York, 1977. Graduate Texts in Mathematics, No. 52.
[HKT07]	Paul Hacking, Sean Keel, and Jenia Tevelev. Stable pair, trop- ical, and log canonical compact moduli of del pezzo surfaces. arXiv.0702505, 2007.
[HT09]	Kerstin Hept and Thorsten Theobald. Tropical bases by regular projections. <i>Proc. Amer. Math. Soc.</i> , 137(7):2233–2241, 2009.
[Hun80]	Thomas W. Hungerford. <i>Algebra</i> , volume 73 of <i>Graduate Texts</i> in <i>Mathematics</i> . Springer-Verlag, New York, 1980. Reprint of the 1974 original.

[IMS07] Ilia Itenberg, Grigory Mikhalkin, and Eugenii Shustin. Tropical algebraic geometry, volume 35 of Oberwolfach Seminars. Birkhäuser Verlag, Basel, 2007.

- [JK08] Michael Joswig and Katja Kulas. Tropical and ordinary convexity combined. arXiv:0801.4835, 2008.
- [Kar78] Richard M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Math.*, 23(3):309–311, 1978.
- [Ked01] Kiran S. Kedlaya. The algebraic closure of the power series field in positive characteristic. *Proc. Amer. Math. Soc.*, 129(12):3461– 3470 (electronic), 2001.
- [Kuh55] H. W. Kuhn. The Hungarian method for the assignment problem. Naval Res. Logist. Quart., 2:83–97, 1955.
- [LMS01] G. L. Litvinov, V. P. Maslov, and G. B. Shpiz. Idempotent functional analysis. An algebraic approach. *Mat. Zametki*, 69(5):758– 797, 2001.
- [Mac01] Diane Maclagan. Antichains of monomial ideals are finite. *Proc. Amer. Math. Soc.*, 129(6):1609–1615 (electronic), 2001.
- [Mar07] Thomas Markwig. A field of generalised Puiseux series for tropical geometry. arXiv:0709.3784, 2007.
- [Mik] Grigory Mikhalkin. Tropical geometry. Book in preparation.
- [Mik04] Grigory Mikhalkin. Decomposition into pairs-of-pants for complex algebraic hypersurfaces. *Topology*, 43(5):1035–1065, 2004.
- [Mik05] Grigory Mikhalkin. Enumerative tropical algebraic geometry in \mathbb{R}^2 . J. Amer. Math. Soc., 18(2):313–377 (electronic), 2005.
- [Oda88] Tadao Oda. Convex bodies and algebraic geometry, volume 15 of Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1988. An introduction to the theory of toric varieties, Translated from the Japanese.
- [Pas85] Donald S. Passman. The algebraic structure of group rings.
 Robert E. Krieger Publishing Co. Inc., Melbourne, FL, 1985.
 Reprint of the 1977 original.

- [Pay09] Sam Payne. Fibers of tropicalization. Math. Z., 262(2):301–311, 2009.
- [Pin98] Jean-Eric Pin. Tropical semirings. In Idempotency (Bristol, 1994), volume 11 of Publ. Newton Inst., pages 50–69. Cambridge Univ. Press, Cambridge, 1998.
- [Poo93] Bjorn Poonen. Maximally complete fields. *Enseign. Math. (2)*, 39(1-2):87–106, 1993.
- [PR04] Mikael Passare and Hans Rullgård. Amoebas, Monge-Ampère measures, and triangulations of the Newton polytope. *Duke Math. J.*, 121(3):481–507, 2004.
- [PS04] Lior Pachter and Bernd Sturmfels. Tropical geometry of statistical models. Proc. Natl. Acad. Sci. USA, 101(46):16132–16137 (electronic), 2004.
- [PS05] Lior Pachter and Bernd Sturmfels, editors. Algebraic statistics for computational biology. Cambridge University Press, New York, 2005.
- [PT05] Mikael Passare and August Tsikh. Amoebas: their spines and their contours. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 275–288. Amer. Math. Soc., Providence, RI, 2005.
- [Pur08] Kevin Purbhoo. A Nullstellensatz for amoebas. *Duke Math. J.*, 141(3):407–445, 2008.
- [Ram02] Jörg Rambau. TOPCOM: triangulations of point configurations and oriented matroids. In *Mathematical software (Beijing, 2002)*, pages 330–340. World Sci. Publ., River Edge, NJ, 2002.
- [Rib99] Paulo Ribenboim. The theory of classical valuations. Springer Monographs in Mathematics. Springer-Verlag, New York, 1999.
- [Sim88] Imre Simon. Recognizable sets with multiplicities in the tropical semiring. In Mathematical foundations of computer science, 1988 (Carlsbad, 1988), volume 324 of Lecture Notes in Comput. Sci., pages 107–120. Springer, Berlin, 1988.

- [SS04] David Speyer and Bernd Sturmfels. The tropical Grassmannian. Adv. Geom., 4(3):389–411, 2004.
- [ST08] Bernd Sturmfels and Jenia Tevelev. Elimination theory for tropical varieties. *Math. Res. Lett.*, 15(3):543–562, 2008.
- [Sta97] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997. With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original.
- [Str84] Ralph Strebel. Finitely presented soluble groups. In *Group the*ory, pages 257–314. Academic Press, London, 1984.
- [Stu96] Bernd Sturmfels. Gröbner bases and convex polytopes, volume 8 of University Lecture Series. American Mathematical Society, Providence, RI, 1996.
- [STY07] Bernd Sturmfels, Jenia Tevelev, and Josephine Yu. The Newton polytope of the implicit equation. *Mosc. Math. J.*, 7(2):327–346, 351, 2007.
- [SY08] Bernd Sturmfels and Josephine Yu. Tropical implicitization and mixed fiber polytopes. In Software for algebraic geometry, volume 148 of IMA Vol. Math. Appl., pages 111–131. Springer, New York, 2008.
- [Tev07] Jenia Tevelev. Compactifications of subvarieties of tori. Amer. J. Math., 129(4):1087–1104, 2007.
- [The02] Thorsten Theobald. Computing amoebas. *Experiment. Math.*, 11(4):513–526 (2003), 2002.
- [Zie95] Günter M. Ziegler. Lectures on polytopes, volume 152 of Graduate Texts in Mathematics. Springer-Verlag, New York, 1995.