


Algorithmic Aspects of Gröbner Fans and Tropical Varieties



Anders Nedergaard Jensen

Ph.D. Dissertation
July 2007

Department of Mathematical Sciences
Faculty of Science, University of Aarhus

Abstract

The Gröbner fan of a polynomial ideal $I \subseteq k[x_1, \dots, x_n]$ is a polyhedral complex in \mathbb{R}^n whose maximal cones are in bijection with the reduced Gröbner bases of I . In tropical algebraic geometry the tropical variety of an ideal is defined. It is the image of an algebraic variety over the Puiseux series field under the negative valuation map. Another description of it is as a certain subcomplex of the Gröbner fan. In this dissertation we study the structure of both polyhedral fans and suggest algorithms for computing them.

Contents

1	Introduction	5
1.1	Summary	12
2	Preliminaries	15
2.1	Convex geometry	15
2.2	Gröbner bases	17
2.3	Primary decomposition	20
2.4	Krull dimension	21
2.5	Laurent polynomials and saturation	23
3	The Gröbner fan of a polynomial ideal	27
3.1	Definitions	27
3.1.1	Proof: The Gröbner fan is a fan	31
3.2	Reverse search property	36
3.2.1	Proof: The Gröbner fan has the reverse search property	39
4	Algorithms for Gröbner fans	43
4.1	Finding facets	44
4.2	Local change	44
4.3	The generic Gröbner walk	47
4.4	Computing the search edge	49
4.5	Exploiting symmetry	49
5	A non-regular Gröbner fan	51
5.1	The proof	52
5.1.1	A necessary condition	52
5.1.2	The certificate	53
5.1.3	Correctness of the subgraph	55
5.2	Further remarks	56
5.2.1	Homogenizing the ideal	56
5.2.2	A program for finding the example	57
6	Tropical varieties	59
6.1	Tropical varieties	62
6.2	Examples and the basic structure	67
6.3	Reduction to $\mathbb{Q}[x_1, \dots, x_n]$	70
6.4	Decomposing tropical varieties	72

6.4.1	Saturated initial ideals and decomposition	73
7	Tropical algorithms	77
7.1	Tropical bases	77
7.2	Computing tropical prevarieties	79
7.2.1	Tropical curves	82
7.3	Connectedness of tropical varieties of prime ideals	84
7.4	Traversing tropical varieties	85
8	A proof of the Bieri Groves theorem	89
8.1	An equivalent theorem	89
8.2	A proof of the equivalent theorem	92
9	Software and examples	95
9.1	Gfan	95
9.2	Computational results and examples	97

Chapter 1

Introduction

Finding the solutions to a system of polynomial equations is a classical mathematical problem. For linear equations the problem is solved by Gaussian elimination which turns the coefficient matrix of the linear system into a matrix of reduced row-echelon form. The reduced row-echelon form is unique once the ordering of the variables (columns) has been fixed. In the case of general polynomial equations the analogue of Gaussian elimination is Buchberger's algorithm which computes a reduced Gröbner basis; see [12] – or the textbooks [14] and [40] for an introduction.

Example 1.0.1 Consider the polynomial system

$$y = x^2 \quad \wedge \quad (x - 5)^2 + y^2 = 5^2$$

describing the intersection of a parabola and a circle in the plane \mathbb{R}^2 . The reduced lexicographic Gröbner basis of $I = \langle y - x^2, (x - 5)^2 + y^2 - 5^2 \rangle = \langle y - x^2, x^2 - 10x + y^2 \rangle$ with $x < y$ is the equivalent system $\{x^4 + x^2 - 10x, y - x^2\}$. Solving the equation $0 = x^4 + x^2 - 10x = x(x - 2)(x + 1 + 2i)(x + 1 - 2i)$ and using backward substitution we get the complete list of real solutions: $(x, y) \in \{(0, 0), (2, 4)\}$.

Another important application of Gröbner bases is for the ideal membership problem which is the problem of deciding if a given polynomial is in a polynomial ideal or not. The answer is yes if and only if the remainder of a division of the polynomial modulo a Gröbner basis of the ideal is zero. Most problems in computer algebra are reduced to Gröbner basis computations. For example the intersection of two polynomial ideals and the dimension and degree of a variety can also be computed with Gröbner bases. For many of these problems the ordering of the terms play an important role.

As for linear systems the reduced Gröbner basis is unique once the ordering of the monomials has been fixed. Such an ordering is called a term order. If n is the number of variables in the system, then in the linear case we have $n!$ choices of orderings of the variables while in the general polynomial case there are uncountably many choices of term orders as soon as $n \geq 2$. It is a surprising fact that the set of reduced Gröbner bases is finite. Which term orderings give which reduced Gröbner bases is the study of the first part of this thesis. This




Figure 1.1: The Gröbner fan of the ideal I_A in Example 1.0.2 for $s = 5$. The Gröbner fan is a three-dimensional object in this example. To get a two-dimensional picture we have drawn the intersection of the Gröbner fan with the standard two-simplex $\text{conv}(e_1, e_2, e_3) \subseteq \mathbb{R}^3$.

information is encoded by the Gröbner fan of the polynomial ideal generated by the system. The Gröbner fan was defined by Mora and Robbiano in the 1988 paper [44].

The Gröbner fan is a polyhedral complex in \mathbb{R}^n . A polyhedral complex is a collection of polyhedra in \mathbb{R}^n with nice intersection properties. If all polyhedra are cones then the complex is called a fan. The full-dimensional polyhedral cones in a Gröbner fan of an ideal I are in bijection with the reduced Gröbner bases of I .

In the following example we will see a family of ideals in $\mathbb{Q}[x_1, x_2, x_3]$ parametrized by an integer s . Each ideal is generated by three polynomials. The example shows how the reduced Gröbner basis with respect to a certain term order changes when we change the generators. Also the number of reduced Gröbner bases changes when s is increased. This can be seen in the Gröbner fan picture in Figure 1.1 where the number of cones increases.

Example 1.0.2 Let $s \in \mathbb{N} \setminus \{0\}$ and $A = [1 + 2s, 3 + 2s, 5 + 2s] \in \mathbb{N}^{1 \times 3}$. Define the toric ideal $I_A := \langle x^u - x^v : A(u - v) = 0 \text{ and } u, v \in \mathbb{N}^3 \rangle \subseteq \mathbb{Q}[x_1, x_2, x_3]$, where $x^w = x_1^{w_1} + \dots + x_3^{w_3}$. The ideal is generated by $\{x_2^2 - x_1x_3, x_1^{s+2}x_2 - x_3^{s+1}, x_1^{s+3} - x_2x_3^s\}$ which is a reverse lexicographic Gröbner basis. The reduced lexicographic Gröbner basis with $x_1 \succ x_2 \succ x_3$ is the $5 + s$ element set

$$\{x_1x_3 - x_2^2, x_2^{2s+5} - x_3^{2s+3}, x_1x_2^{2s+3} - x_3^{2s+2}, \dots, x_1^{s+2}x_2 - x_3^{s+1}, x_1^{s+3} - x_2x_3^s\}.$$

Varying the term order we get $2s + 9$ different reduced Gröbner bases. These are indexed by the maximal cones in the Gröbner fan of I_A ; see Figure 1.1.

The family was found by computer experiments. To actually show that the listed set is a reduced Gröbner basis requires some work but is doable by performing Buchberger's algorithm by hand. To prove that the number of reduced Gröbner bases is correct requires a lot more work.

With the many applications of Gröbner bases there is an enormous interest in making efficient algorithms for their computation. The example above suggests that the complexity of a Gröbner basis depends on the term order. Many experiments support this observation. A general rule is that degree reverse lexicographic Gröbner bases are much easier to compute than lexicographic Gröbner bases. For many problems, unfortunately, we are interested in a lexicographic Gröbner basis or in other difficult bases. One approach for computing a lexicographic Gröbner basis more efficiently is to start by computing an reverse lexicographic basis and then change it into the desired lexicographic basis by some algorithm. This last step is called Gröbner basis conversion. For zero-dimensional ideals the step can be performed with the FGLM algorithm [18]. In that case the conversion problem basically reduces to linear algebra. For general ideals the Gröbner walk [13] applies. The idea is to walk along a straight line in the Gröbner fan towards the desired term order and change the Gröbner basis when passing from one full-dimensional cone to the next. The step of moving from one full-dimensional cone to the next is relatively easy. See [38] for results on the complexity of this local step.

It is important to realize that there is no Gröbner basis algorithm that will perform well on all examples. Sometimes the Gröbner walk performs well and sometimes the complexity i.e. the number of cones in the Gröbner fan makes it impractical while other algorithms succeed. It is our hope that a better understanding of the Gröbner fan will lead to a better understanding of the complexity of Gröbner bases. This could be either a deep understanding or simply an understanding of how the Gröbner cones group together and make Gröbner walks impractical. In that respect Figure 3.3 in Section 3.1 says more than a thousand words.

The goal of the first half of the thesis is to properly define Gröbner fans and prove properties that will be useful for computing them. In order to develop algorithms for computing Gröbner fans a good understanding of their structure is required. For this reason a proof that the Gröbner fan is actually a polyhedral complex was carefully worked out, see Section 3.1.1. This proof was missing in the original paper [44] and in [53] the inhomogeneous case was not covered.

One of the important properties of Gröbner fans is that if the ideal is homogeneous then there exists a polytope called the state polytope whose normal fan equals the Gröbner fan. If for example the ideal is a principal ideal then the Gröbner fan will be the normal fan of the Newton polytope of the generator. The idea of the state polytope appeared in [5]. In the book [53] which has become a standard reference for Gröbner fans and state polytopes a construction is given.

Many of the algorithmic problems concerning Gröbner fans were already solved by the introduction of the Gröbner walk in [13]. Also the original papers [5] and [44] contained some algorithms. In [53] two fundamentally different algorithms for computing the Gröbner fan of an ideal are described. The first one [53, Algorithm 3.2] works for homogeneous ideals and builds up an inequality description of the state polytope. This algorithm was implemented by Alyson Reeves; see [53, page 29]. Since the algorithm works with a vertex-facet description of the state polytope it is difficult to implement in practise and we

choose to study the second proposed algorithm. The second algorithm enumerates the maximal cones in the fan by applying the local Gröbner basis change procedure from the Gröbner walk. For homogeneous toric ideals the algorithm was already studied in [32]. The toric case is very nice since all polynomials ever appearing are monomials and binomials. It turns out that all essential properties translates to the general homogeneous case. In fact, mathematically the general situation is not more difficult, but for implementations the toric case is easier.

Since the local basis change procedure [13] is already known, implementing the algorithm is a matter of applying a global enumeration method. In [32] the reverse search strategy was applied — a strategy originally applied for enumerating vertices of a polytope; see [2]. Thus the application of this method to Gröbner fans relies on the existence of the state polytope and a certain orientation of its edges. Having implemented the reverse search and having a seemingly working algorithm for non-homogeneous ideals, it was the author's and one coauthor's (Rekha Thomas) belief for a long time that a state polytope would always exist, but a construction was missing. It was not until our software Gfan [33] found the example in Theorem 5.0.1 that we realized that the situation was more subtle: a state polytope does not always exist, but the reverse search always works.

Computing the Gröbner fan of an ideal and taking the union of all Gröbner bases we get a *universal* Gröbner bases — a set that is a Gröbner basis with respect to any term order. Sometimes people are interested in finding a Gröbner basis with a special property. For example if an ideal has a square-free initial ideal then it automatically follows that the ideal is a radical ideal ([54, Proposition 5.3]). In these cases Gröbner fan algorithms are relevant.

Two questions arise in the Gröbner walk. How do we get an interior point of the target Gröbner cone and how can we be sure that we only leave a Gröbner cone through a facet? In [58] these problems were solved by using a degree bound on the Gröbner bases to construct a vector that is guaranteed to be in the lexicographic Gröbner cone. Since the degree bound is big arbitrary length arithmetic must be handled in practical implementations. The generic Gröbner walk which we present solves both problems in a simple way by symbolically perturbing the target weight vector. This perturbation also turns out to be useful for the reverse search traversal of the Gröbner fan.

The Gröbner fan is a complicated object which may seem random with little structure except from the property of indexing Gröbner bases. In our introduction so far our attention has been on the full-dimensional cones in the Gröbner fan since they correspond to Gröbner bases. Any cone in the Gröbner fan has an algebraic object associated to it, namely an initial ideal of I . For a full-dimensional cone the initial ideal is the monomial ideal generated by the initial terms of the Gröbner basis. In general the initial ideal is not a monomial ideal. The cones whose initial ideals are monomial-free turn out to form a very interesting piecewise linear object called a tropical variety.

The second half of this dissertation is concerned with the tropical variety of a polynomial ideal I . Originally tropical varieties were studied by Bergman as logarithmic limit sets of algebraic varieties; see [7]. We will use two different




Figure 1.2: A line in $(\mathbb{C}\{\{t\}\}^*)^2$, its tropicalization in \mathbb{R}^2 and the Gröbner fan of the ideal defining the line. See Example 1.0.3. The Gröbner fan has two full-dimensional cones corresponding to the two differently marked Gröbner bases $\{\underline{x} + y - 1\}$ and $\{\underline{y} + x - 1\}$. Since, for instance, $(3, 3)$ is in the tropical variety, there is a zero of I of the form $(c_1 t^{-3} + \dots, c_2 t^{-3} + \dots) \in (\mathbb{C}\{\{t\}\}^*)^2$ where $c_1, c_2 \in \mathbb{C}^*$ and the dots are some higher order terms.

definitions of the tropical variety of an ideal I . One definition will be used for proofs and algorithms while the other is used for intuition and motivation. Having already investigated Gröbner fans, it is natural for us to define the tropical variety as a certain subcomplex of the Gröbner fan of I . It is the subcomplex of all cones whose initial ideals are monomial-free. This definition was introduced in [51]. Another definition comes from considering a variety V in the algebraic torus of the Puiseux series field $(\mathbb{C}\{\{t\}\}^*)^n$ defined by I and taking its image under the negative coordinatewise valuation map. We explain this in the following. The function $\text{val} : \mathbb{C}\{\{t\}\}^* \rightarrow \mathbb{Q}$ takes a Puiseux series to the exponent of the first term in the series. The function extends to $\text{val} : (\mathbb{C}\{\{t\}\}^*)^n \rightarrow \mathbb{Q}^n$. Let $V \subseteq (\mathbb{C}\{\{t\}\}^*)^n$ be a variety defined by a polynomial ideal I . The tropicalization of V which equals the tropical variety of I is defined as $-\text{val}(V) \subseteq \mathbb{R}^n$ where the closure is taken in the usual topology. We should think of the tropical variety as a tropical shadow of V .

Example 1.0.3 Let $I = \langle x + y - 1 \rangle \subseteq \mathbb{C}[x, y]$ be an ideal. The ideal defines a line in $(\mathbb{C}\{\{t\}\}^*)^2$; see Figure 1.2. We are faced with the usual difficulties with drawings in algebraic geometry, namely that the complex plane is drawn as an axis. In fact, it is not just the complex plane that is drawn as an axis but the whole Puiseux series field. A lot of information is missing in the picture to the left. Applying the $-\text{val}$ function we get the tropical variety in the second drawing. It is not clear at all that this is how the tropical variety looks. Finally we show the Gröbner of I and observe that the tropical variety is a subcomplex.

In general the tropicalization of an algebraic variety is a piecewise linear object. Here we list three important theorems concerning tropical varieties.

- Theorem 6.1.17 which states that the above definitions are equivalent.
- Bieri Groves Theorem 6.4.3 which states that the tropical variety of a d -dimensional prime ideal is a pure d -dimensional polyhedral complex.

- Theorem 7.3.2 which states that the tropical variety of a prime ideal is connected in codimension one.

The first theorem has been proved several places; see Section 6.1. The second theorem appeared in [8]. The third theorem appeared in [10] and is crucial for one of our algorithms for computing tropical varieties.

An introduction to tropical geometry is given in Chapter 6 where also a list of applications is given. Tropical geometry is a relatively new field with a lot of activity. Only few attempts have been made for computing tropical varieties. These have mainly been concerned with special cases such as linear ideals or generators with generic coefficients; see [56]. Of course Speyer's and Sturmfels' observation in [51] that tropical varieties can be defined as a subcomplex of the Gröbner fan immediately gives algorithms for computing them. Our work in [10] is the first detailed study of such algorithms while a study of the complexity of the problem of computing tropical varieties was carried out in [57].

The study of tropical varieties can be seen as a way of turning an algebraic problem into combinatorial one. This of course is a well-know strategy in mathematics, we may study toric varieties by looking at their defining polytopes and polynomial ideals by looking at monomial initial ideals. Properties of the combinatorial object carry over to the original problem. The best example of this for tropical varieties is that the dimension is the same for a variety in the algebraic torus and its tropicalization.

Being less philosophical, we should consider the tropical variety as what it really is. The tropical variety of an ideal I shows how to find the zeros of I in the Puiseux series field, namely it consists of all possible first exponents of a solution to I . Thus the first step for solving a polynomial system over the Puiseux series field is to compute a tropical variety. With this in mind we should reconsider the first two drawings in Figure 1.2. Which drawing tells us most about the Puiseux series solutions of the ideal? An algorithm which takes as input a point in a tropical variety of I and lifts it to a Puiseux series solution is presented in [42]. It uses the software Gfan described in Section 9.1. In this setting we may think of the thesis as a thesis about solving polynomial systems.

While Puiseux series solutions may seem of little practical importance, tropical varieties do have applications in solving of polynomial systems. The state of the art method for solving polynomial systems for zero-dimensional ideals numerically is the homotopy method. The idea here is to trace solution paths between solutions of an easy start system to solutions of the original system. In the polyhedral homotopy method the start systems are found by computing a tropical variety; see [31]. The variety is a tropical prevariety and its computation appears as a substep of our algorithms.

This thesis is mainly concerned with the combinatorial and computational aspects of the Gröbner fan and its subfan. On one hand the algorithms are developed as tools for studying the structure of the Gröbner fan and allows us to produce examples illustrating certain features. One such example is the non-regular Gröbner fan in Theorem 5.0.1. On the other hand in the process of developing algorithms we discover new results such as Theorem 3.2.6. Thus there is a lot of interaction between the process of developing algorithms and the

process of discovering theorems. The word “algorithmic” in the title might as well have been “combinatorial” or “algebraic”. We chose the word “algorithmic” to emphasize our goal.

Two topics which are highly relevant for algorithms for Gröbner fans and tropical varieties are not covered in detail in this dissertation which has its main focus on mathematics. The first topic is computational complexity. While Gaussian reduction for a square matrix can be done using $O(n^3)$ field operations, the complexity of Buchberger’s algorithm is doubly exponential in the worst case. See [6, page 511-514] for an overview of Gröbner basis complexity. Example 1.0.2 above shows that the size of a reduced Gröbner basis for a toric ideal can be single exponential in the size of the generators. The high complexity of Gröbner basis computations is not unexpected. In computational commutative algebra NP-hard problems appear all the time. For example Gröbner bases of toric ideals can be used for solving integer programs; see [53]. An other example of an NP-hard problem appears in Section 4.5 about symmetry where one needs to check if two polynomials belong to the same orbit under the action of a symmetry group. The graph-isomorphism problem, to decide if two graphs are isomorphic, is a special instance of that problem. Even finding the dimension of a monomial ideal is difficult complexity-wise; see Section 2.4. As already mentioned complexity questions for tropical varieties were studied in [57].

The second topic which is not covered in detail is computational geometry. Many of the algorithms presented in this thesis rely on methods for doing polyhedral computations. For the first half of the thesis these computations amount to finding redundant inequalities of full dimensional cones and computing interior points of polyhedra. Dantzig’s simplex algorithm will solve these problems. In the second half, lower-dimensional cones are considered and techniques for finding extreme rays, dimension and writing cones in a unique form become relevant. The problem of finding extreme rays is solved by the double description method in [45]. The half open cones in Section 7.2 can be implemented by considering closed cones in a space of dimension one higher. In [60] a useful extension of the reverse search appeared which allows the search to compute all cones in a fan, not just the maximal ones, without producing duplicates. This method was used to compute the f-vectors of the Gröbner fans in Section 9.2.

1.1 Summary

This dissertation summarizes work done in the following papers

- Fukuda, J., Thomas: “Computing Gröbner fans” [21]
- Fukuda, J., Lauritzen, Thomas: “The generic Gröbner walk” [19]
- J.: “A non-regular Gröbner fan” [37]
- Bogart, J., Speyer, Sturmfels, Thomas: “Computing tropical varieties” [10]

and contains a few results from

- J., Markwig, Markwig: “An algorithm for lifting points in a tropical variety” [42].

In Chapter 2 we recall definitions and results concerning convex geometry, Gröbner bases, primary decomposition, Krull dimension, the Laurent polynomial ring and saturation. Section 2.1 and Section 2.2 were partly taken from [21].

Chapter 3 is basically taken from [21] with some changes and some additions. In this chapter we present our uniform definition of Gröbner fans which works for both homogeneous and non-homogeneous ideals. We also present a proof that the Gröbner fan is actually a polyhedral complex. In the process of proving this result many useful propositions are developed. Most of these are slight generalizations of theorems already known. This is the content of the first half of Chapter 3 which is thus mainly about the local structure of the Gröbner fan. The second half is concerned with the global structure. We define what it means for a fan to have the reverse search property and explain why this property is important for computations. We finish the chapter by proving that Gröbner fans have the reverse search property. The definition of the reverse property and the result that Gröbner fans have this property is new and appeared [21].

Chapter 4 is partly taken from [21] and is partly new. The algorithms for Gröbner fans are presented. Since the local change procedure was already known, the main contribution in this chapter is the introduction of the generic Gröbner walk. We presented the generic Gröbner walk in [19]. The presentation in Section 4.3 is simplified in the sense that only the target point is perturbed. This suffices for solving the two mentioned problems with Gröbner walks, see Section 4.3, and to efficiently define the search edge for reverse search, see Section 4.4.

In Chapter 5 we present an example that shows that the (restricted) Gröbner fan of a non-homogeneous ideal is not always the normal fan of a polyhedron. This settles an open problem and emphasizes the importance of the reverse search property in the previous chapter. The example also shows that care must be taken when homogenizing an ideal. The example first appeared in [37] and most of the chapter was taken from there. We discuss the homogenization process for the example and show how to compute initial ideals outside the Gröbner region.

Chapter 6 is an introduction to tropical varieties. After a general introduction to tropical geometry we define tropical varieties in terms of t -initial ideals. The notion of t -initial ideals appeared in the writing process of [42]. At this point we were heavily influenced by Bernd Sturmfels and it is not clear who came up with the name and the definition. The idea of using initial ideals in the definition of tropical varieties goes back to [51] and [54]. The rest of the chapter is mainly concerned with the polyhedral structure of these varieties, the reduction to the constant coefficient case and how tropical varieties decompose into varieties defined by prime ideals. We also show that it suffices to consider tropical varieties of homogeneous ideals. Essentially, most of these results were already known but some were maybe not clearly stated. We provide proofs for the theorems we need. This chapter can be seen as a collection of observations in the papers [10] and [42] and other before them ([51], [54]...). One useful result that is new is Lemma 6.3.5 which states that the computation of a tropical variety of an ideal defined with coefficients in an algebraic field extensions of \mathbb{Q} can be reduced to a computation of a tropical variety defined by an ideal in $\mathbb{Q}[x_1, \dots, x_{n+1}]$.

Chapter 7 is about tropical algorithms. The algorithms presented here appeared in [10]. Many new results which appeared in [10] are presented in this chapter. We prove that every ideal has a tropical basis. This had already been claimed in [51] without a proof. Our proof is constructive. We discuss algorithms for checking if an ideal contains a monomial and for computing intersections of tropical hypersurfaces. In the case of tropical curves we present a new algorithm for computing a tropical basis. We present algorithms for computing connected components of tropical varieties which are also new. Speyer's Theorem 7.3.2 which was proved in [10] states that the tropical variety of a prime ideal is connected. We do not prove Theorem 7.3.2 in this thesis.

In Chapter 8 we give an almost self contained proof of Bieri Groves Theorem 6.4.3 which states that the tropical variety of a prime ideal is pure. While the theorem was originally stated in terms of valuations we prove it for the initial ideal definition of tropical varieties. The theorem has also been proved in this setting in [54].

The algorithms presented in this dissertation have all been implemented in the software Gfan [33]. In Chapter 9 we give an example of how this software is used and list some Gröbner fans and tropical varieties that have been computed with the software.

Acknowledgments

I would like to thank my advisor Niels Lauritzen for inspiring project oriented teaching. In particular for teaching a course in the fall of 1998 based on the book [53] which later led to my master thesis project [34]. Also thanks for the open-minded guidance.

This dissertation would not have been written if it was not for Rekha Thomas who kindly invited me to Seattle for an extended period of time during the academic year 2003-2004. Starting up many new projects this led to my return to academia.

The many polyhedral discussions with Komei Fukuda have been extremely useful in theory and practise. I would like to thank Komei Fukuda for co-advicing and hospitality through the academic year 2004-2005.

I would also like to thank coauthors and fellow students for many interesting conversations, Bernd Sturmfels for his many project proposals and this year's visitors at the Institute for Mathematics and its Applications for introducing me to many new mathematical topics and for making this a wonderful year in general.

Thanks to the following institutions for their hospitality and support: Institut für Operations Research, ETH Zürich, Department of Mathematics, University of Washington, Institute for Mathematics and its Applications, University of Minnesota and Institut for Matematiske Fag, University of Aarhus.

Anders Nedergaard Jensen
Århus, July 2007

Chapter 2

Preliminaries

In this chapter we recall some basic definitions and results. Five topics are covered: convex geometry, Gröbner bases, primary decomposition, Krull dimension and saturation. Only the sections about convex geometry and Gröbner bases are needed for the first half of this thesis. The reader is advised to skip Section 2.3, 2.4 and 2.5 for now.

2.1 Convex geometry

A standard reference for polyhedral and convex geometry is Ziegler's book [61]. Here we will just give some definitions and state a few theorems. Although many of the stated and implicitly stated theorems are intuitively clear they are not all easy to prove. An example is Definition 2.1.5 which states that the Minkowski sum of two polyhedra is a polyhedron. This follows from the Minkowski-Weyl theorem for polyhedra; see [61].

We recall the definition of a fan in \mathbb{R}^n . A *polyhedron* in \mathbb{R}^n is a set of the form $\{x \in \mathbb{R}^n : Ax \leq b\}$ where A is a matrix and b is a vector. Bounded polyhedra are called *polytopes*. If $b = 0$ the set is a *polyhedral cone*. The *dimension* $\dim(P)$ of a polyhedron P is the dimension of the smallest affine subspace containing it. A *face* of a polyhedron P is either the empty set or a non-empty subset of P which is the set of maximizers of a linear form over P . We use the following notation for the face maximizing a form $\omega \in \mathbb{R}^n$:

$$\text{face}_\omega(P) := \{p \in P : \langle \omega, p \rangle = \max_{q \in P} \langle \omega, q \rangle\}.$$

A face of P is called a *facet* if its dimension is one smaller than the dimension of P .

Definition 2.1.1 A collection \mathcal{C} of polyhedra in \mathbb{R}^n is a *polyhedral complex* if:

1. all non-empty faces of a polyhedron $P \in \mathcal{C}$ are in \mathcal{C} , and
2. the intersection of any two polyhedra $A, B \in \mathcal{C}$ is a face of A and a face of B .




Figure 2.1: The four collections of cones in Example 2.1.2.

The *support* $\text{supp}(\mathcal{C})$ of \mathcal{C} is the union of all the members of \mathcal{C} . A polyhedral complex is a *fan* if it only consists of cones. If the support of a fan is \mathbb{R}^n , the fan is said to be *complete*. A fan is *pure* if all its maximal cones have the same dimension.

An example of a polyhedral complex is the set of non-empty faces of a polyhedron. If the smallest dimension of a non-empty polyhedron in a complex \mathcal{C} is m and the largest dimension is M then the *f-vector* of \mathcal{C} is the $(M - m + 1)$ -dimensional vector whose i th entry is the number of polyhedra in \mathcal{C} of dimension $m + i - 1$. The f-vector of the complex consisting of the faces of the cube $[0, 1] \times [0, 1] \times [0, 1] \subseteq \mathbb{R}^3$ is $(8, 12, 6, 1)$.

Example 2.1.2 In Figure 2.1 four collections of cones in the plane are drawn:

- The first collection consists of 1 zero-dimensional cone, 5 rays and 3 two-dimensional cones. This is a fan which is not pure and not complete.
- The second collection consists of 1 zero-dimensional cone, 5 rays and 4 two-dimensional cones. This is a pure fan which is not complete.
- The third collection consists of 1 line and 2 regions. This is a pure complete fan.
- The fourth collection consists of 1 zero-dimensional cone, three rays, one line and three regions of which one is a half space. This is *not* a polyhedral fan since the intersection of the half space with one of the other regions is not a face of the half space.

Notice that for a finite non-empty fan the intersection of all cones is a subspace. This subspace is the smallest non-empty face of every cone in the fan.

A simple way to construct a fan is by taking the *normal fan* of a polyhedron.

Definition 2.1.3 Let $P \subseteq \mathbb{R}^n$ be a polyhedron. For a face F of P we define its *normal cone*

$$N_P(F) := \overline{\{\omega \in \mathbb{R}^n : \text{face}_\omega(P) = F\}}$$

with the closure being taken in the usual topology. The *normal fan* of P is the fan consisting of the normal cones $N_P(F)$ as F runs through all non-empty faces of P .

The following equation is satisfied for a non-empty face F :

$$\dim(N_P(F)) + \dim(F) = n.$$

It is clear that the normal fan of a polytope is complete. In Figure 2.1 the third fan is the normal fan of a horizontal linesegment. The other two are not normal fans. Not all complete fans arise as the normal fan of a polyhedron [22, page 25].

Definition 2.1.4 The *common refinement* of two fans \mathcal{F}_1 and \mathcal{F}_2 in \mathbb{R}^n is defined as

$$\mathcal{F}_1 \wedge \mathcal{F}_2 = \{C_1 \cap C_2\}_{(C_1, C_2) \in \mathcal{F}_1 \times \mathcal{F}_2}.$$

The common refinement of two fans is a fan.

Definition 2.1.5 The *Minkowski sum* of two polyhedra $P, Q \subseteq \mathbb{R}^n$ is the polyhedron $P + Q := \{p + q : (p, q) \in P \times Q\}$.

Proposition 2.1.6 Let $P, Q \subseteq \mathbb{R}^n$ be two polyhedra. The normal fan of $P + Q$ is the common refinement of the normal fan of P and the normal fan of Q .

We write $\text{int}(V)$ for the interior of a subset V of a topological space.

Definition 2.1.7 The *relative interior* of a polyhedron $P \subseteq \mathbb{R}^n$ is the interior of $P \cap L \subseteq L$ where L is the smallest affine subspace of \mathbb{R}^n containing P . Here L has its topology induced from \mathbb{R}^n . We denote the relative interior by $\text{rel int}(P)$.

2.2 Gröbner bases

We assume that the reader knows the basics of Gröbner basis theory like the division algorithm for multivariate polynomials, S-polynomials and Buchberger's algorithm. If not, the books [14] and [40] are recommended. Besides making our notation clear we make a few important points about marked Gröbner bases and homogeneous ideals in this section.

Let $R = k[x_1, \dots, x_n]$ be the polynomial ring in n variables over a field k and let $I \subseteq R$ be an ideal. For $\alpha \in \mathbb{N}^n$ we use the notation $x^\alpha := x_1^{\alpha_1} \dots x_n^{\alpha_n}$ for a monomial in R . By a *term order* on R we mean a total ordering \preceq on all monomials in R such that:

1. For all $\alpha \in \mathbb{N}^n \setminus \{0\} : 1 \prec x^\alpha$ and
2. for $\alpha, \beta, \gamma \in \mathbb{N}^n : x^\alpha \prec x^\beta \Rightarrow x^\alpha x^\gamma \prec x^\beta x^\gamma$.

A total ordering is an antisymmetric, transitive and total relation. By “ \prec ” we mean $x^\alpha \prec x^\beta \Leftrightarrow x^\alpha \preceq x^\beta \wedge x^\alpha \neq x^\beta$. We prefer the symbol \prec over the symbol \preceq when denoting term orders. By a *term* we mean a monomial together with its coefficient. Term orders are used for ordering terms, ignoring the coefficients. For a vector $\omega \in \mathbb{R}_{\geq 0}^n$ and a term order \prec we define the new term order \prec_ω as follows:

$$x^\alpha \prec_\omega x^\beta \iff \langle \omega, \alpha \rangle < \langle \omega, \beta \rangle \vee (\langle \omega, \alpha \rangle = \langle \omega, \beta \rangle \wedge x^\alpha \prec x^\beta).$$

Let \prec be a term order. For a non-zero polynomial $f \in R$ we define its *initial term*, $\text{in}_{\prec}(f)$, to be the unique maximal term of f with respect to \prec . In the same way for $\omega \in \mathbb{R}^n$ we define the *initial form*, $\text{in}_{\omega}(f)$, to be the sum of all terms of $f \in R$ whose exponents maximize $\langle \omega, \cdot \rangle$. Notice that for a non-zero polynomial f we have $\text{in}_{\prec_{\omega}}(f) = \text{in}_{\prec}(\text{in}_{\omega}(f))$. The *Newton polytope*, $\text{New}(f)$, of a polynomial f is the convex hull of its exponent vectors.

Remark 2.2.1 Let $f \in k[x_1, \dots, x_n]$ be a polynomial and $P \subseteq \mathbb{R}^n$ its Newton polytope. Notice that $\text{in}_u(f) = \text{in}_v(f) \Leftrightarrow \text{face}_u(P) = \text{face}_v(P)$ and that $\text{in}_u(f)$ is a monomial if and only if $\text{face}_u(P)$ has dimension 0 or, equivalently, the normal cone $N_P(\text{face}_u(P))$ is full-dimensional.

The ω -degree of a term cx^{α} is $\langle \omega, \alpha \rangle$ and the ω -degree of a non-zero polynomial f is the maximal ω -degree of the terms of $\text{in}_{\omega}(f)$. The *initial ideals* of an ideal I with respect to \prec and ω are defined as

$$\text{in}_{\prec}(I) = \langle \text{in}_{\prec}(f) : f \in I \setminus \{0\} \rangle \text{ and } \text{in}_{\omega}(I) = \langle \text{in}_{\omega}(f) : f \in I \rangle.$$

Note that $\text{in}_{\prec}(I)$ is a monomial ideal while $\text{in}_{\omega}(I)$ might not be. A monomial in $R \setminus \text{in}_{\prec}(I)$ (with coefficient 1) is called a *standard monomial* of $\text{in}_{\prec}(I)$.

Although initial ideals are defined with respect to not necessarily positive vectors, Gröbner bases are only defined with respect to true term orders:

Definition 2.2.2 Let $I \subseteq R$ be an ideal and \prec a term order on R . A generating set $\mathcal{G} = \{g_1, \dots, g_m\}$ for I is called a *Gröbner basis* for I with respect to \prec if

$$\text{in}_{\prec}(I) = \langle \text{in}_{\prec}(g_1), \dots, \text{in}_{\prec}(g_m) \rangle.$$

The Gröbner basis \mathcal{G} is *minimal* if $\{\text{in}_{\prec}(g_1), \dots, \text{in}_{\prec}(g_m)\}$ generates $\text{in}_{\prec}(I)$ minimally. A minimal Gröbner basis is *reduced* if the initial term of every $g \in \mathcal{G}$ has coefficient 1 and all other monomials in g are standard monomials of $\text{in}_{\prec}(I)$.

We use the term *marked Gröbner basis* for a Gröbner basis where the initial terms have been distinguished from the non-initial ones (they have been marked). For example, $\{\underline{x}^2 + xy + y^2\}$ and $\{x^2 + xy + \underline{y}^2\}$ are marked Gröbner bases for the ideal $\langle x^2 + xy + y^2 \rangle$ while $\{x^2 + \underline{xy} + y^2\}$ is not since xy is not the initial term of $x^2 + xy + y^2$ with respect to any term order.

Given a marking of one term of each polynomial in a set $\mathcal{G} \subseteq R$, the division algorithm produces the remainder of a polynomial $f \in R$ modulo \mathcal{G} . Here the elements of \mathcal{G} can be considered as exchange rules and the algorithm terminates when no more rules apply. The remainder depends on the order in which these rules or reduction steps are applied. If \mathcal{G} is a marked Gröbner basis then the remainder is unique and we call it the *normal form* of f modulo \mathcal{G} . The normal form does not depend on a term order but only on the markings of the Gröbner basis.

For a term order \prec and an ideal I , Buchberger's algorithm guarantees the existence of a unique marked reduced Gröbner basis. We denote it by $\mathcal{G}_{\prec}(I)$. Buchberger's algorithm is a completion procedure that keeps adding remainders of S-polynomials to a generating set of I until Buchberger's S-criterion is satisfied:

Theorem 2.2.3 *Let $I \subseteq R$ be an ideal and \prec a term order. A polynomial set $\mathcal{G} \subseteq R$ marked according to \prec is a Gröbner basis of I with respect to \prec if for all $g_1, g_2 \in \mathcal{G}$ some remainder of the division algorithm run on $S(g_1, g_2)$ modulo \mathcal{G} is zero.*

Here $S(g_1, g_2)$ is the S -polynomial $c_2x^{(v_1 \vee v_2) - v_1}g_1 - c_1x^{(v_1 \vee v_2) - v_2}g_2$, assuming that $c_i x^{v_i}$ is the marked term of g_i with $c_i \in k$ and $v_i \in \mathbb{N}^n$ for $i = 1, 2$ and with $v_1 \vee v_2$ being the coordinate-wise maximum of v_1 and v_2 .

Remark 2.2.4 For two term orders \prec and \prec' , if $\text{in}_\prec(I) = \text{in}_{\prec'}(I)$ then $\mathcal{G}_\prec(I) = \mathcal{G}_{\prec'}(I)$. To see this, consider a polynomial $g \in \mathcal{G}_\prec(I)$. Since $\mathcal{G}_\prec(I)$ is reduced only the marked term of g is in $\text{in}_\prec(I)$. Hence $\text{in}_{\prec'}(g)$ which is in $\text{in}_{\prec'}(I) = \text{in}_\prec(I)$ must be the same marked term. This shows that the S -polynomials of elements in $\mathcal{G}_\prec(I)$ are the same no matter which of the two term orders we consider. Furthermore, the division algorithm only depends on the markings in $\mathcal{G}_\prec(I)$. Thus all S -polynomials have remainder zero. This proves that $\mathcal{G}_\prec(I)$ is a Gröbner basis with respect to \prec' . It is also reduced and by uniqueness of reduced Gröbner bases we get $\mathcal{G}_\prec(I) = \mathcal{G}_{\prec'}(I)$.

Conversely, given a marked Gröbner basis $\mathcal{G}_\prec(I)$, the initial ideal $\text{in}_\prec(I)$ can easily be read off.

Let $\omega \in \mathbb{R}^n$. A polynomial $f \in R$ is ω -homogeneous if $\text{in}_\omega(f) = f$. An ideal $I \subseteq R$ is ω -homogeneous if it is generated by ω -homogeneous elements. Hilbert's basis theorem states that I has a finite generating set. Each generator in the set can be expressed in terms of finitely many ω -homogeneous generators. This proves that I has a finite generating set consisting of ω -homogeneous elements. Using the ω -homogeneous generating set, a polynomial $f \in I$ can be uniquely written as a finite sum $f = \sum_i f_i$ where the f_i 's are ω -homogeneous, have different ω -degrees and belong to I . It is now easy to deduce the following lemmas.

Lemma 2.2.5 *Let $I \subseteq R$ be an ideal and $\omega \in \mathbb{R}^n$. Then I is ω -homogeneous if and only if $\text{in}_\omega(I) = I$.*

Lemma 2.2.6 *Let $I \subseteq R$ be an ω -homogeneous ideal with $\omega \in \mathbb{R}^n$ and $v \in \mathbb{R}^n$. Then $\text{in}_{v+s\omega}(I) = \text{in}_v(I)$ for any $s \in \mathbb{R}$.*

We say that an ideal is homogeneous if it is ω -homogeneous for some positive vector $\omega \in \mathbb{R}_{>0}^n$.

Notice that if Buchberger's algorithm gets ω -homogeneous generators as input, then the output is also ω -homogeneous. In particular, all reduced Gröbner bases for an ω -homogeneous ideal consists of ω -homogeneous elements. As a consequence equations defining the subspace of vectors for which I is homogeneous can be read off from any reduced Gröbner basis. We conclude that $\mathcal{G}_\prec(I) = \mathcal{G}_{\prec_\omega}(I)$ for any term order if I is ω -homogeneous with $\omega \in \mathbb{R}_{>0}^n$.

Remark 2.2.7 In fact, we may be less strict with our orderings if we are given homogeneous generators for an ideal. Let \prec be an ordering which is a term order except that it does not satisfy $1 \prec x^\alpha$ for $\alpha \in \mathbb{N}^n \setminus \{0\}$. Then \prec has the

property that it will give the same reduced Gröbner basis as \prec_ω when run on an ω -homogeneous generating set. One example of this is the *reverse lexicographic* “term order” which only makes sense for homogeneous generating sets (with respect to a positive grading ω). For $x_n \prec x_{n-1} \prec \cdots \prec x_1$ it is defined as follows

$$x^u \prec x^v \Leftrightarrow \exists j : u_j > v_j \wedge \forall i > j : u_i = v_i.$$

Definition 2.2.8 Let $A \in GL_n(\mathbb{R})$ be an invertible matrix. If the first non-zero entry in each column is positive then the matrix defines a *matrix term order* \prec_A in the following way. We define $x^a \prec_A x^b$ whenever the first non-zero entry $A(a - b)$ is negative, where $a, b \in \mathbb{N}^n$.

Lemma 2.2.9 Let $A \in \mathbb{R}^{n \times n}$ be a matrix defining a term order and let $a, b \in \mathbb{N}^n$. We define $A_\varepsilon = \varepsilon^0 A_1 + \varepsilon^1 A_2 + \cdots + \varepsilon^{n-1} A_n$ where A_i is the i th row of A . For all $\varepsilon > 0$ sufficiently small we have $x^a \prec_A x^b$ if and only if $A_\varepsilon \cdot a < A_\varepsilon \cdot b$.

Proof. Assume $a \neq b$. Let $M \in \mathbb{R}_{>0}$ be a number larger than the numerical value of any entry in $A(a - b)$. Let L be the first non-zero entry of $A(a - b)$. For every positive $\varepsilon < \frac{|L|}{nM}$ the sign of $A_\varepsilon(a - b)$ equals the sign of L . \square

An important result concerning term orders is the following theorem which states that any term order has a matrix representation.

Theorem 2.2.10 ([49]) Let \prec be a term order on $k[x_1, \dots, x_n]$. There exists a matrix $T \in \mathbb{R}^{n \times n}$ with rows τ_1, \dots, τ_n such that for $u, v \in \mathbb{N}^n$

$$x^u \prec x^v \Leftrightarrow \exists j : \langle \tau_j, u \rangle < \langle \tau_j, v \rangle \wedge \forall i < j : \langle \tau_i, u \rangle = \langle \tau_i, v \rangle.$$

2.3 Primary decomposition

By an algebraic variety we mean the zero-set of an ideal $I \subseteq k[x_1, \dots, x_n]$ where k is a field. We denote it by $V(I) \subseteq k^n$. If we want to consider the zero-set over a larger field k' or a zero-set in the algebraic torus $(k^*)^n = (k \setminus \{0\})^n$ we write $V_{k'}(I)$ or $V_{k^*}(I)$ respectively. In the algebraic torus the variety $V_{k^*}(I)$ can also be defined by an ideal in the Laurent polynomial ring $k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. A variety is *irreducible* if it cannot be written as a union of two other varieties in a non-trivial way. A variety can be written as a finite union of irreducible varieties. The corresponding algebraic notion is that of primary decomposition of ideals. Primary decompositions will be important in the second half of the thesis where we see that tropical varieties can be decomposed according to primary decompositions of their defining ideals. Primary decomposition will also be import for proofs in Chapter 8.

Let R be a commutative ring and $I \subseteq R$ an ideal. The *associated primes* of I are the prime ideals of the form $(I : f) := \{g \in R : gf \in I\}$ where $f \in R$. The set of associated primes of I is denoted by $\text{Ass}(I)$. The ideals in $\text{Ass}(I)$ which are minimal with respect to inclusion are called the *minimal* associated primes of I . The set of minimal associated primes is denoted by $\text{minAss}(I)$.

Associated primes that are not minimal are called *embedded*. It turns out that the minimal associated primes are exactly the minimal prime ideals among all prime ideals containing I ; see [1, Proposition 4.6].

An ideal $Q \subsetneq R$ is primary if for $ab \in Q$ either $a \in Q$ or there exists $m \in \mathbb{N}$ such that $b^m \in Q$. The radical $P = \sqrt{Q} := \{a \in R : a^m \in Q \text{ for some } m \in \mathbb{N}\}$ of a primary ideal is a prime ideal. We say that Q is *P-primary* if Q is primary with $\sqrt{Q} = P$. The intersection $Q_1 \cap Q_2$ of two P -primary ideals is again P -primary; see [1, Lemma 4.3]. For an ideal $I \subseteq R$ a primary decomposition is an expression of I as a finite intersection of primary ideals:

$$I = \bigcap_i Q_i.$$

A primary decomposition is *minimal* if the radicals $\sqrt{Q_i}$ are all distinct and no Q_i can be left out in the intersection. By intersecting primary ideals with the same radical and removing primary ideals that are not needed we can make any primary decomposition minimal.

The radicals of the Q_i 's in a minimal primary decomposition are exactly the associated primes of I ; see [1, Theorem 4.5]. For a minimal primary decomposition and $P \in \min\text{Ass}(I)$ the primary component Q_i with $\sqrt{Q_i} = P$ is unique; see [1, Corollary 4.11].

In a Noetherian ring every ideal has a primary decomposition ([1, Theorem 7.13]). In particular, any ideal in the polynomial ring or Laurent polynomial ring over a field has a primary decomposition.

Algorithms for computing primary decompositions in polynomial rings exist (see [26]) and rely on methods for factoring polynomials. In the case where $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ is generated by polynomials in $\mathbb{Q}[x_1, \dots, x_n]$ the field \mathbb{Q} may need to be extended. An implementation that does extensions as needed can be found in the algebra system Singular [28] under the name “absolute primary decomposition”.

2.4 Krull dimension

The best reference for the notion of dimension of an ideal is the book [29] since it starts with the rather algebraic definition of the Krull dimension of a ring and ends up by showing how to compute the dimension of a polynomial ideal.

Definition 2.4.1 The *Krull dimension* of a ring R is the supremum of the lengths of chains of prime ideals in R .

$$P_0 \subsetneq P_1 \subsetneq P_2 \subsetneq \dots P_d \subsetneq R$$

Here the length is the number of prime ideals excluding P_0 (and R which we do not consider to be a prime ideal).

All rings that we shall consider have finite Krull dimension.

By the *dimension of an ideal* $I \subset R$ we mean the Krull dimension of the quotient ring R/I and we denote it by $\dim(I)$. There is an inclusion preserving

bijection between the prime ideals in R containing I and the prime ideals in R/I by taking P' to $P = P' + I$. For this reason an equivalent definition of the dimension of an ideal would be the length d of the longest chain of inclusions

$$I \subseteq P'_0 \subsetneq P'_1 \subsetneq \dots \subsetneq P'_d \subsetneq R$$

of prime ideals P'_i containing I . Here we do not include P'_0 in the length of the chain. If $I = R$ we say that the dimension of I is -1 .

In a longest chain, P'_0 is always a minimal associated prime of I . It follows that $\dim(I)$ is the maximum of the dimensions of the minimal primes $\min\text{Ass}(I)$.

Given generators of a monomial ideal in $k[x_1, \dots, x_n]$ it is a simple, although time consuming, process to compute the dimension.

Lemma 2.4.2 *Let k be a field and $I \subseteq k[x_1, \dots, x_n]$ an ideal generated by a set of monomials M . The dimension $\dim(I)$ is the number of elements in the largest subset $\sigma \subseteq \{x_1, \dots, x_n\}$ such that $M \cap k[x]_{x \in \sigma} = \emptyset$. Here $k[x]_{x \in \sigma}$ is the polynomial ring in the variables of σ .*

Example 2.4.3 Let $I = \langle x^2, xy \rangle \subseteq k[x, y, z]$ where k is a field. A longest chain of inclusions of prime ideals in $k[x, y, z]/I$ is

$$\langle x + I \rangle \subsetneq \langle x + I, y + I \rangle \subsetneq \langle x + I, y + I, z + I \rangle \subsetneq k[x, y, z]/I.$$

Hence $\dim(I) = \dim(k[x, y, z]/I) = 2$. Notice that the chain cannot be extended on the left by $\langle 0 + I \rangle$ since this is not a prime ideal in $k[x, y, z]/I$. The corresponding chain of inclusions in $k[x, y, z]$ of prime ideals containing I is

$$I \subseteq \langle x \rangle \subsetneq \langle x, y \rangle \subsetneq \langle x, y, z \rangle \subsetneq k[x, y, z].$$

The ideal $\langle x \rangle$ is a minimal associated prime of I . The intersection $\{x^2, xy\} \cap k[y, z] = \emptyset$ while $\{x^2, xy\} \cap k[x, y, z] \neq \emptyset$.

Theorem 2.4.4 [29, Corollary 5.3.9 and Corollary 7.5.6] *Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal and \prec a term order. Then*

$$\dim(I) = \dim(\text{in}_{\prec}(I)).$$

Remark 2.4.5 For ideals homogeneous with respect to the grading $(1, 1, \dots, 1)$ the above theorem is proved using Hilbert functions; see [29, Chapter 5]. The book [29] also has a proof for the case of non-homogeneous ideals. That proof relies on the notion of flatness; see [29, Chapter 7 and Remark 5.3.18]. In light of this thesis another approach for the non-homogeneous case is the following. It is relatively easy to prove the theorem for degree term orders using homogenization; see [29, Corollary 5.3.14]. Initial ideals with respect to term orders correspond to maximal cones in the Gröbner fan which we shall define. We can go between any two cones in the Gröbner fan by a series of “flips”. For every flip the new and old initial ideal is also the initial ideal of a “facet-initial ideal”. This ideal is homogeneous with respect to some positive grading. If we can show that the dimension is preserved under a flip we would have another proof for the theorem. Preservation of dimension under a flip can be shown using Definition 5.4 and Lemma 5.5 in the *arxiv.org* version of [11].

Using Theorem 2.4.4, Lemma 2.4.2 and Buchberger's algorithm we can compute the dimension of any ideal $I \subseteq k[x_1, \dots, x_n]$.

Definition 2.4.6 Let $A \in \mathbb{Z}^{d \times n}$ be a matrix and k a field. The *toric ideal* of A is

$$I_A := \langle x^u - x^v : Au = Av \text{ for } u, v \in \mathbb{N}^n \rangle \subseteq k[x_1, \dots, x_n].$$

Lemma 2.4.7 [53, Lemma 4.2] *The toric ideal I_A has dimension $\text{rank}(A)$.*

2.5 Laurent polynomials and saturation

Gröbner bases work for polynomial rings while tropical varieties are naturally defined for ideals in the Laurent polynomial ring. Let k be a field. Given an ideal $I \subseteq k[x_1, \dots, x_n]$ we may consider the ideal it generates in the Laurent polynomial ring $k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. In this section we prove that the dimension of the two ideals is the same under certain conditions and we show how to computationally go from the Laurent polynomial ring to the polynomial ring.

Definition 2.5.1 Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal and $f \in k[x_1, \dots, x_n]$ a polynomial. The *saturation* of I with f is the ideal defined by

$$(I : f^\infty) := \{g \in k[x_1, \dots, x_n] : gf^m \in I \text{ for some } m \in \mathbb{N}\}.$$

The above definition can be made more general by considering other rings than polynomial rings and saturating with ideals instead of polynomials. However, we will only need the definition for polynomial rings with f being a monomial. If $(I : x_1 \cdots x_n^\infty) = I$ we say that the ideal I is *saturated*.

In [53, Chapter 12] the identity $(I : x_1 \cdots x_n^\infty) = ((\cdots (I : x_1^\infty) \cdots) : x_n^\infty)$ and a method for computing saturations with respect to a single variable using reverse lexicographic Gröbner bases was given. The single variable saturation goes as follows.

Algorithm 2.5.2 [53, Algorithm 12.1]

Input: A set of generators for a homogeneous ideal $I \subseteq k[x_1, \dots, x_n]$.

Output: A Gröbner basis for the ideal $(I : x_n^\infty)$.

```
{
  Return {sat(f, x_n) : f ∈ G_<(I)};
}
```

Here $\text{sat}(f, x_n)$ denotes the polynomial $x^{-m}f$ where m is the highest power such that x^m divides f and $<$ is the (degree) reverse lexicographic term order with $x_1 > \cdots > x_n$.

Let $\sigma \subseteq \{x_1, \dots, x_n\}$ and $I \subseteq \mathbb{C}[x_1, \dots, x_n]$. Let us consider the geometric meaning of having $I = (I : x_\sigma^\infty)$ where x_σ is the product of the variables in σ . Either $1 \in I$ or $x_\sigma^m \notin I$ for all m . In the first case $V(I) = \emptyset$. In the second case $x_\sigma \notin \sqrt{I} = I(V(I))$ by Hilbert's Nullstellensatz. That x_σ does not vanish on $V(I)$ means that there is a point $p \in V(I)$ such that the coordinates of p indexed by σ are non-zero. If $x_\sigma = x_1 \cdots x_n$ then this means that $V(I) \cap (\mathbb{C}^*)^n \neq \emptyset$. It

is easy to see that $V(I : f^\infty) \subseteq V(I)$ and that if $p \in V(I)$ and $f(p) \neq 0$ then $p \in V(I : f^\infty)$. In fact, we remove all components contained in the coordinate hyperplanes when saturating I with respect to $x_1 \cdots x_n$. This follows from Lemma 2.5.8 below which states that any associated prime of $(I : x_1 \cdots x_n^\infty)$ is saturated and thus must have a zero in the algebraic torus $(\mathbb{C}^*)^n$.

Lemma 2.5.3 *Let $I \subseteq k[x_1, \dots, x_n]$ be a saturated ideal. Then*

$$\langle\langle I \rangle\rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]} \cap k[x_1, \dots, x_n] = I.$$

Proof. The inclusion “ \supseteq ” is clear. For the other inclusion suppose f is in the left hand side. The f has the form $\sum_i m_i g_i$ where $g_i \in I$ and m_i are Laurent monomials. Multiplying the sum by a sufficiently large power $(x_1 \cdots x_n)^m$ the sum will involve only monomials in $k[x_1, \dots, x_n]$. This shows that $(x_1 \cdots x_n)^m f \in I$. However, since I is saturated we also get $f \in I$. \square

Lemma 2.5.4 *Let $I \subseteq k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be an ideal then*

$$\langle\langle I \cap k[x_1, \dots, x_n] \rangle\rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]} = I.$$

Proof. The inclusion “ \subseteq ” is clear. For the other inclusion suppose $f \in I$ then for a sufficiently high power $m \in \mathbb{N}$ we have $(x_1 \cdots x_n)^m f \in I \cap k[x_1, \dots, x_n]$. This proves that $f \in \langle\langle I \cap k[x_1, \dots, x_n] \rangle\rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]}$. \square

Remark 2.5.5 The above lemmas show that there is a bijection between saturated ideals in $k[x_1, \dots, x_n]$ and ideals in the Laurent polynomial ring.

Lemma 2.5.6 *Let $P \subseteq k[x_1, \dots, x_n]$ be a saturated prime ideal. Then the ideal $\langle P \rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]}$ is prime.*

Proof. Suppose $a, b \in k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ with $ab \in \langle P \rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]}$. Then the polynomials $(x_1 \cdots x_n)^m a$ and $(x_1 \cdots x_n)^m b$ belong to $k[x_1, \dots, x_n]$ for sufficiently large $m \in \mathbb{N}$. By Lemma 2.5.3 their product $(x_1 \cdots x_n)^{2m} ab$ is in P . Since P is prime either $(x_1 \cdots x_n)^m a$ or $(x_1 \cdots x_n)^m b$ is in P . This proves that a or b is in $\langle P \rangle_{k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]}$. \square

Lemma 2.5.7 *Let $P' \subseteq k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be a prime ideal. Then the intersection $P' \cap k[x_1, \dots, x_n]$ is prime.*

Proof. Clearly, if $a, b \in k[x_1, \dots, x_n]$ satisfy $ab \in P'$ then either a or b is in P' and thus in $(P' \cap k[x_1, \dots, x_n])$. \square

Lemma 2.5.8 *Let $I \subseteq k[x_1, \dots, x_n]$ be a saturated ideal. Then any $P \in \text{Ass}(I)$ satisfies $(P : x_1 \cdots x_n^\infty) = P$.*

Proof. Suppose this was not the case then we would be able to find $g \in (P : x_1 \cdots x_n^\infty) \setminus P$. The polynomial g must satisfy $g(x_1 \cdots x_n)^m \in P$ for some power m . The ideal P is a prime ideal which shows that $x_1 \cdots x_n \in P$. However, P has the form $(I : f)$ for some $f \in k[x_1, \dots, x_n]$ which shows that $(x_1 \cdots x_n)f \in I$. Using that $(I : x_1 \cdots x_n^\infty) = I$ we see that $f \in I$. This proves that $P = (I : f) = \langle 1 \rangle$ contradicting that P is a prime ideal. \square

The following proposition shows the ideals that correspond to each other in the sense of Remark 2.5.5 have the same dimension.

Proposition 2.5.9 *Let $J \subseteq k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be an ideal of dimension d . Then the dimension of $I := J \cap k[x_1, \dots, x_n]$ is d .*

Proof. Recall that the codimension $\text{codim}(P)$ of a prime ideal P is the maximal number of strict inclusions in a chain of prime ideal contained in P . We start by proving the theorem in the case where J is prime. The codimension of J is $n - d$. We have a chain of inclusions of prime ideals P'_i :

$$P'_0 \subsetneq \cdots \subsetneq J = P'_{n-d} \subsetneq \cdots \subsetneq P'_n \subsetneq k[x_1^{\pm 1}, \dots, x_n^{\pm 1}].$$

Applying the correspondence of ideals in the two rings together with the inclusion preservation we get an inclusion of prime ideals P_i

$$P_0 \subsetneq \cdots \subsetneq I = P_{n-d} \subsetneq \cdots \subsetneq P_n \subsetneq k[x_1, \dots, x_n]$$

proving that $\text{codim}(I) \geq n - d$ and $\dim(I) \geq d$. Since $\text{codim}(I) + \dim(I) \leq \dim(k[x_1, \dots, x_n]) = n$ we get $\dim(I) = d$.

If J is not prime then Lemma 2.5.8 shows that the minimal primes of I are saturated and thus in bijection with the minimal primes of J . The dimension of an ideal is the maximal dimension of one of its minimal associated primes. Preservation of dimension for prime ideals now proves the claim. \square

The proposition and Algorithm 2.5.2 together give a method for computing the dimension of an ideal in the Laurent polynomial ring.

Chapter 3

The Gröbner fan of a polynomial ideal

In this chapter we define the Gröbner fan of a polynomial ideal and show that it is a fan in the sense of a polyhedral complex. We explain the reverse search technique for traversing graphs and show that the Gröbner fan has what we shall call the reverse search property.

In this chapter, Chapter 4 and Chapter 5 we let $R = k[x_1, \dots, x_n]$ be the polynomial ring in n variables over a field k and let $I \subseteq R$ be an ideal.

3.1 Definitions

The *Gröbner fan* and the *restricted Gröbner fan* of I are n -dimensional polyhedral fans defined in [44]. We will present our definition below.

Given an ideal $I \subseteq R$, a natural equivalence relation on \mathbb{R}^n is induced by taking initial ideals:

$$u \sim v \iff \text{in}_u(I) = \text{in}_v(I). \quad (3.1)$$

We introduce the following notation for the closures of the equivalence classes:

$$C_{\prec}(I) = \overline{\{u \in \mathbb{R}^n : \text{in}_u(I) = \text{in}_{\prec}(I)\}} \quad \text{and}$$

$$C_v(I) = \overline{\{u \in \mathbb{R}^n : \text{in}_u(I) = \text{in}_v(I)\}}.$$

The closed set $C_0(I)$ is also known as the *homogeneity space* of I . We define $\text{homog}(I) := \dim(C_0(I))$.

Remark 3.1.1 It follows from Lemma 2.2.6 that any cone $C_v(I)$ is invariant under translation by any vector $\omega \in C_0(I)$.

Remark 3.1.2 It is well known that for a fixed ideal I there are only finitely many sets $C_{\prec}(I)$ and they cover $\mathbb{R}_{\geq 0}^n$, see [44] and [53, Theorem 1.2]. Secondly, every initial ideal $\text{in}_{\prec}(I)$ is of the form $\text{in}_{\omega}(I)$ for some $\omega \in \mathbb{R}_{> 0}^n$, see [53, Proposition 1.11]. This will follow from Lemma 2.2.9, Theorem 2.2.10 and Lemma 3.1.9 below. Consequently, every $C_{\prec}(I)$ is of the form $C_{\omega}(I)$.

A third observation is that the equivalence classes are not convex in general since we allow the vectors to be anywhere in \mathbb{R}^n :

Example 3.1.3 Let $I = \langle x - 1, y - 1 \rangle$. The ideal I has five initial ideals: $\langle x - 1, y - 1 \rangle$, $\langle x, y \rangle$, $\langle x, y - 1 \rangle$, $\langle x - 1, y \rangle$ and $\langle 1 \rangle$. In particular, for $u = (-1, 3)$ and $v = (3, -1)$ we have $\text{in}_u(I) = \text{in}_v(I) = \langle 1 \rangle$ but $\text{in}_{\frac{1}{2}(u+v)}(I) = \langle x, y \rangle$.

The following proposition which we shall prove in the next section is very important since it describes equivalence classes in terms of equations and inequalities.

Proposition 3.1.4 *Let \prec be a term order and $v \in C_{\prec}(I)$. For $u \in \mathbb{R}^n$*

$$\text{in}_u(I) = \text{in}_v(I) \iff \forall g \in \mathcal{G}_{\prec}(I), \text{in}_u(g) = \text{in}_v(g).$$

The proposition is a little more general than Proposition 2.3 in [53] as it allows the vectors u and v to have negative components. For fixed \prec and v as in Proposition 3.1.4, we get that $C_v(I)$, the closure of the equivalence class of v , is a polyhedral cone since each $g \in \mathcal{G}_{\prec}(I)$ introduces the equation $\text{in}_u(g) = \text{in}_v(g)$ which is equivalent to having u satisfy a set of linear equations and strict linear inequalities, see Example 3.1.5. The closure is obtained by making the strict inequalities non-strict. Under the assumptions of Proposition 3.1.4 we may write this in the following way:

$$u \in C_v(I) \iff \forall g \in \mathcal{G}_{\prec}(I), \text{in}_v(\text{in}_u(g)) = \text{in}_v(g). \quad (3.2)$$

As we saw in Example 3.1.3, not all equivalence classes are convex. However, for an arbitrary v , $C_v(I)$ is a convex polyhedral cone if it contains a strictly positive vector. To see this, notice that there must exist a vector $p \in \mathbb{R}_{>0}^n$ with $\text{in}_p(I) = \text{in}_v(I)$ and, by Lemma 3.1.14 below, $p \in C_{\prec_p}(I)$ for any \prec . Hence the equivalence class of v is of the form required in Proposition 3.1.4.

Example 3.1.5 Let $I = \langle x + y + z, x^3z + x + y^2 \rangle \subseteq \mathbb{Q}[x, y, z]$ and let \prec be the lexicographic term order with $x \prec y \prec z$. Then $\mathcal{G}_{\prec}(I) = \{y^2 + x - x^3y - x^4, z + y + x\}$. If $v = (1, 4, 5)$ then $\text{in}_v(I) = \text{in}_{\prec}(I) = \langle y^2, z \rangle$ and $C_v(I) = C_{\prec}(I)$. By Proposition 3.1.4, $\text{in}_u(I) = \text{in}_v(I)$ if and only if the following two equations are satisfied:

$$\begin{aligned} \text{in}_u(z + y + x) &= z \iff u_z > \max\{u_x, u_y\}, \text{ and} \\ \text{in}_u(y^2 + x - x^3y - x^4) &= y^2 \iff 2u_y > \max\{u_x, 3u_x + u_y, 4u_x\}. \end{aligned}$$

Introducing non-strict inequalities we obtain a description of $C_{\prec}(I)$. This cone is simplicial and has the cones $C_{(0,0,1)}(I)$, $C_{(1,3,3)}(I)$ and $C_{(-2,-1,-1)}(I)$ as extreme rays and $C_{(1,3,4)}(I)$, $C_{(-2,-1,0)}(I)$ and $C_{(-1,2,2)}(I)$ as facets; see Figure 3.1. Since $(-2, -1, 0)$ is in $C_{\prec}(I)$ a description of vectors u in $C_{(-2,-1,0)}(I)$ is given by:

$$\begin{aligned} \text{in}_{(-2,-1,0)}(\text{in}_u(z + y + x)) &= z \iff u_z \geq \max\{u_x, u_y\}, \text{ and} \\ \text{in}_{(-2,-1,0)}(\text{in}_u(y^2 + x - x^3y - x^4)) &= y^2 + x \iff 2u_y = u_x \geq \max\{3u_x + u_y, 4u_x\}. \end{aligned}$$




Figure 3.1: The Gröbner fan of the ideal in Example 3.1.5 has 7 three-dimensional, 14 two-dimensional and 8 one-dimensional cones. The intersections of the two-dimensional cones with the hyperplane $x + y + z = 1$ are drawn as lines. The dotted part of the figure shows the combinatorial structure outside the hyperplane. The gray triangle indicates the positive orthant.

Definition 3.1.6 The *Gröbner fan* of an ideal $I \subseteq R$ is the set of the closures of all equivalence classes intersecting the positive orthant together with their proper faces. The cones in a Gröbner fan are called *Gröbner cones*.

This is a variation of the definitions appearing in the literature. The advantage of this variant is that it gives well-defined and *nice* fans in the homogeneous and non-homogeneous case simultaneously. By *nice* we mean that all cones in this fan are closures of equivalence classes. It is not clear a priori that the Gröbner fan is a polyhedral complex. A proof is given in the next section (Theorem 3.1.19). The support of the Gröbner fan of I is called the *Gröbner region* of I . We define the *restricted Gröbner fan* of an ideal to be the common refinement (Definition 2.1.4) of the Gröbner fan and the faces of the non-negative orthant. The support of the restricted Gröbner fan is $\mathbb{R}_{\geq 0}^n$.

Example 3.1.7 The Gröbner fan of the principal ideal $\langle x^4 + x^4y - x^3y + x^2y^2 + y \rangle$ consists of one 0-dimensional cone, three 1-dimensional cones and two 2-dimensional cones, see Figure 3.1. The same is true for the restricted Gröbner fan. Notice, however, that in the restricted Gröbner fan one of the 1-dimensional cones and one of the 2-dimensional cones are not closures of equivalence classes of the equivalence relation (3.1).

Example 3.1.8 [53, Example 3.9] Consider the ideal $I = \langle a^5 - 1 + c^2 + b^3, b^2 - 1 + c + a^2, c^3 - 1 + b^5 + a^6 \rangle \subseteq \mathbb{Q}[a, b, c]$. The Gröbner fan of I has 360 full-dimensional cones and the Gröbner region is $\mathbb{R}_{\geq 0}^3$. This means that the restricted Gröbner fan equals the Gröbner fan. The intersection of the fan with the standard simplex in \mathbb{R}^3 is shown in Figure 3.3.




Figure 3.2: The Gröbner fan of the ideal in Example 3.1.7 is shown on the left. The restricted Gröbner fan is on the right. In the middle the Newton polytope of the generator is drawn with the shape of its normal fan indicated.




Figure 3.3: The Gröbner fan of the ideal in Example 3.1.8 intersected with the standard 2-simplex. The a -axis is on the right, the b -axis on the left and the c -axis at the top.

3.1.1 Proof: The Gröbner fan is a fan

In this section we prove that the Gröbner fan is a fan i.e., that it is a polyhedral complex consisting of cones. Recall, in general the Gröbner fan is not complete and its support is larger than $\mathbb{R}_{\geq 0}^n$. In [44] there is no proof that the Gröbner fan is a fan in the sense of a polyhedral complex. A proof that the Gröbner fan is a polyhedral complex under the assumption that the ideal is homogeneous is given in [53, Chapter 2]. We present a complete proof for the general case. Many of the results we need in the proof are generalizations of known results needed in the proof in [53]. However, we do not rely on these references for the sake of being self-contained.

We fix the ideal $I \subseteq R$ in the following theorems. The most important step is the proof of Proposition 3.1.4 which tells us that the closure of an equivalence class is a polyhedral cone. Then we prove that the relative interior of any face in the Gröbner fan is an equivalence class (Proposition 3.1.16) and, finally, that the intersection of two cones in the fan is a face of both (Proposition 3.1.18).

To prove Proposition 3.1.4 we start by proving a similar statement for the equivalence classes arising from initial ideals with respect to term orders.

Lemma 3.1.9 *Let \prec be a term order. For $v \in \mathbb{R}^n$,*

$$\text{in}_v(I) = \text{in}_{\prec}(I) \iff \forall g \in \mathcal{G}_{\prec}(I), \text{in}_v(g) = \text{in}_{\prec}(g).$$

Proof. \Rightarrow : Let $g \in \mathcal{G}_{\prec}(I)$. Since $\mathcal{G}_{\prec}(I)$ is reduced, only one term from g , $\text{in}_{\prec}(g)$, can be in $\text{in}_{\prec}(I) = \text{in}_v(I)$. The initial ideal $\text{in}_v(I)$ is a monomial ideal, implying that all terms of an element in the ideal must be in the ideal too. Hence, the initial form $\text{in}_v(g) \in \text{in}_v(I)$ has to be equal to $\text{in}_{\prec}(g)$.

\Leftarrow : We must show that $\text{in}_v(I) = \text{in}_{\prec}(I)$ where $\text{in}_{\prec}(I) = \langle \text{in}_{\prec}(g) \rangle_{g \in \mathcal{G}_{\prec}(I)}$. The “ \supseteq ” inclusion is clear since $\text{in}_{\prec}(g) = \text{in}_v(g) \in \text{in}_v(I)$ for all $g \in \mathcal{G}_{\prec}(I)$.

To prove the “ \subseteq ” inclusion, since $\text{in}_v(I) = \langle \text{in}_v(f), f \in I \rangle$, it suffices to show that $\text{in}_v(f) \in \text{in}_{\prec}(I)$ for all $f \in I$. Pick $f \in I$ and reduce it to zero using the division algorithm (e.g. [14, Chapter 2]) with $\mathcal{G}_{\prec}(I)$ and \prec . We may write

$$f = m_1 g_{i_1} + \cdots + m_r g_{i_r} \tag{3.3}$$

where m_j is a monomial and g_{i_j} is an element from $\mathcal{G}_{\prec}(I)$. The division algorithm guarantees that $\text{in}_{\prec}(f) \geq m_j \text{in}_{\prec}(g_{i_j})$ with respect to \prec since monomials are substituted with monomials less than the original ones with respect to \prec in the division process. Exactly the same thing is true for v -degrees since v and \prec agree on $\mathcal{G}_{\prec}(I)$. Thereby, any monomial on the right hand side in (3.3) has v -degree less than or equal to the v -degree of the left hand side. Consequently,

$$\text{in}_v(f) = \sum_{j \in J} m_j \text{in}_v(g_{i_j})$$

with j running through a subset such that $m_j \text{in}_v(g_{i_j})$ has the same v -degree as $\text{in}_v(f)$. Since $\text{in}_v(g) \in \text{in}_{\prec}(I)$, the initial form $\text{in}_v(f) \in \text{in}_{\prec}(I)$. \square

By Lemma 3.1.9 the equivalence class of $\text{in}_{\prec}(I)$ is open. Since $\text{in}_{\prec}(I)$ is of the form $\text{in}_v(I)$ for some v (see Remark 3.1.2), the equivalence class of $\text{in}_{\prec}(I)$ is also non-empty and hence full-dimensional. Thus we have proved that the equivalence class of a term order is a full dimensional open polyhedral cone.

Corollary 3.1.10 *Let \prec be a term order and $v \in \mathbb{R}^n$. Then*

$$v \in C_{\prec}(I) \Leftrightarrow \forall g \in \mathcal{G}_{\prec}(I) : \text{in}_{\prec}(\text{in}_v(g)) = \text{in}_{\prec}(g).$$

Proof. Lemma 3.1.9 tells us that v lies in the interior of $C_{\prec}(I)$ if and only if $\text{in}_v(g) = \text{in}_{\prec}(g)$ for all $g \in \mathcal{G}_{\prec}(I)$. Relaxing the resulting strict inequalities to non-strict inequalities we get a description of $C_{\prec}(I)$. This relaxation is exactly the one given by $\text{in}_{\prec}(\text{in}_v(g)) = \text{in}_{\prec}(g)$ for all g in $\mathcal{G}_{\prec}(I)$. \square

Lemma 3.1.11 *A polynomial $f \in \text{in}_v(I)$ can be written in the form $f = \sum_i \text{in}_v(c_i)$ where $c_i \in I$ and all summands in the sum have different v -degrees.*

Proof. The initial ideal $\text{in}_v(I)$ is generated by v -homogeneous polynomials, implying that all v -homogeneous components of f are in $\text{in}_v(I)$. Let h be a maximal v -homogeneous component of f . We need to show that h is the initial form of an element in I with respect to v . We may write h as $\text{in}_v(a_1) + \cdots + \text{in}_v(a_s)$ for some polynomials a_1, \dots, a_s in I . Since h is v -homogeneous we can rewrite h as the sum $\sum_{j \in J} \text{in}_v(a_j)$ of forms having the same v -degree as h . We pull out the initial form and get $h = \text{in}_v(\sum_{j \in J} a_j)$. \square

Lemma 3.1.12 *Let \prec be a term order. If $v \in C_{\prec}(I)$ then $\text{in}_{\prec}(\text{in}_v(I)) = \text{in}_{\prec}(I)$.*

Proof. Let $g \in \mathcal{G}_{\prec}(I)$. Since $v \in C_{\prec}(I)$, by Corollary 3.1.10, we have the equality $\text{in}_{\prec}(g) = \text{in}_{\prec}(\text{in}_v(g))$ and hence $\text{in}_{\prec}(I) = \langle \text{in}_{\prec}(g) \rangle_{g \in \mathcal{G}_{\prec}(I)} \subseteq \text{in}_{\prec}(\text{in}_v(I))$.

We now prove that $\text{in}_{\prec}(\text{in}_v(I)) \subseteq \text{in}_{\prec}(I)$. Notice that $\text{in}_{\prec}(\text{in}_v(I))$ is generated by initial terms of elements $f \in \text{in}_v(I) \setminus \{0\}$ with respect to \prec . Suppose $f \in \text{in}_v(I) \setminus \{0\}$. It suffices to show that $\text{in}_{\prec}(f) \in \text{in}_{\prec}(I)$. Using Lemma 3.1.11 we may write $f = \sum_{i=1}^s \text{in}_v(c_i)$ where $c_1, \dots, c_s \in I$ and $\text{in}_v(c_1), \dots, \text{in}_v(c_s)$ are v -homogeneous each with distinct degree, so that no cancellations occur. Consequently $\text{in}_{\prec}(f)$ equals $\text{in}_{\prec}(\text{in}_v(c_j))$ for some j . We wish to prove that $\text{in}_{\prec}(\text{in}_v(c_j)) \in \text{in}_{\prec}(I)$. We use the division algorithm with $\mathcal{G}_{\prec}(I)$ and \prec to rewrite c_j

$$c_j = m_1 g_{i_1} + \cdots + m_r g_{i_r}$$

where m_1, \dots, m_r are monomials and g_{i_1}, \dots, g_{i_r} belong to $\mathcal{G}_{\prec}(I)$. Let M be the v -degree of c_j . In the division algorithm we sequentially reduce c_j to zero. In each step, the v -degree of c_j will decrease or stay the same since we subtract the product of a monomial and an element from $\mathcal{G}_{\prec}(I)$ where the v -degree of the product already appeared in c_j by Corollary 3.1.10. Equivalently, the product of the monomial and the element from $\mathcal{G}_{\prec}(I)$ are “added” to the right hand side of the equation. We are done when $c_j = 0$ and or equivalently, the original

c_j is written as the above sum with every term having v -degree less or equal to M . Consequently, we have

$$\text{in}_v(c_j) = \sum_{j' \in J'} \text{in}_v(m_{j'} g_{i_{j'}})$$

for a suitable J' . The division algorithm guarantees that the exponent vectors of $\text{in}_\prec(m_1 g_{i_1}), \dots, \text{in}_\prec(m_r g_{i_r})$ are distinct. Since $v \in C_\prec(I)$, they equal $\text{in}_\prec(\text{in}_v(m_1 g_{i_1})), \dots, \text{in}_\prec(\text{in}_v(m_r g_{i_r}))$. The maximal one of these with respect to \prec cannot cancel in the sum. Hence $\text{in}_\prec(\text{in}_v(c_j)) = \text{in}_\prec(m_{j'} g_{i_{j'}})$ for some j' which implies that $\text{in}_\prec(\text{in}_v(c_j)) \in \text{in}_\prec(I)$ as needed. \square

An easy corollary is a method for computing Gröbner bases for initial ideals.

Corollary 3.1.13 *Let \prec be a term order. If $v \in C_\prec(I)$ then*

$$\mathcal{G}_\prec(\text{in}_v(I)) = \{\text{in}_v(g)\}_{g \in \mathcal{G}_\prec(I)}.$$

Proof. By Corollary 3.1.10, $\langle \text{in}_\prec(\text{in}_v(g)) \rangle_{g \in \mathcal{G}_\prec(I)} = \langle \text{in}_\prec(g) \rangle_{g \in \mathcal{G}_\prec(I)} = \text{in}_\prec(I)$. By Lemma 3.1.12, $\text{in}_\prec(I)$ equals $\text{in}_\prec(\text{in}_v(I))$. Thus we have the equality of ideals $\text{in}_\prec(\text{in}_v(I)) = \langle \text{in}_\prec(\text{in}_v(g)) \rangle_{g \in \mathcal{G}_\prec(I)}$. This proves that $\{\text{in}_v(g)\}_{g \in \mathcal{G}_\prec(I)}$ is a Gröbner basis of $\text{in}_v(I)$ with respect to \prec . It is reduced since $\mathcal{G}_\prec(I)$ is minimal and reduced. \square

We are now able to give a proof for Proposition 3.1.4 which claimed that given $v \in C_\prec(I)$ and $u \in \mathbb{R}^n$, $\text{in}_u(I) = \text{in}_v(I) \iff \forall g \in \mathcal{G}_\prec(I), \text{in}_u(g) = \text{in}_v(g)$.

Proof of Proposition 3.1.4. \Leftarrow : Since $\text{in}_u(g) = \text{in}_v(g)$ for all $g \in \mathcal{G}_\prec(I)$, we get that $\text{in}_\prec(\text{in}_u(g)) = \text{in}_\prec(\text{in}_v(g))$ for all $g \in \mathcal{G}_\prec(I)$. Since $v \in C_\prec(I)$, by Corollary 3.1.10, $\text{in}_\prec(g) = \text{in}_\prec(\text{in}_v(g))$ for all $g \in \mathcal{G}_\prec(I)$ and hence $\text{in}_\prec(g) = \text{in}_\prec(\text{in}_u(g))$ for all $g \in \mathcal{G}_\prec(I)$ and $u \in C_\prec(I)$ by Corollary 3.1.10. The Gröbner basis $\mathcal{G}_\prec(\text{in}_u(I))$ is then $\{\text{in}_u(g)\}_{g \in \mathcal{G}_\prec(I)}$ by Corollary 3.1.13. We get the same Gröbner basis for $\text{in}_v(I)$. Hence, $\text{in}_u(I) = \text{in}_v(I)$.

\Rightarrow : Let $g \in \mathcal{G}_\prec(I)$. We need to show that $\text{in}_u(g) = \text{in}_v(g)$. Since the basis is reduced, only one term of g , namely $\text{in}_\prec(g)$, is in $\text{in}_\prec(I)$. We start by proving that the term $\text{in}_\prec(g)$ is a term in $\text{in}_v(g)$ and a term in $\text{in}_u(g)$. For $\text{in}_v(g)$ we apply Corollary 3.1.10 which says $\text{in}_\prec(g) = \text{in}_\prec(\text{in}_v(g))$. For $\text{in}_u(g)$ we apply Lemma 3.1.12 and get $\text{in}_\prec(\text{in}_u(g)) \in \text{in}_\prec(\text{in}_u(I)) = \text{in}_\prec(\text{in}_v(I)) = \text{in}_\prec(I)$. Only one term of g is in $\text{in}_\prec(I)$, so $\text{in}_\prec(\text{in}_u(g)) = \text{in}_\prec(g)$. If the difference $\text{in}_u(g) - \text{in}_v(g)$, belonging to $\text{in}_u(I) = \text{in}_v(I)$, is non-zero we immediately reach a contradiction since the difference contains no terms from $\text{in}_\prec(I) = \text{in}_\prec(\text{in}_v(I))$. \square

We have now proved that every equivalence class of a vector v in a $C_\prec(I)$ is a relatively open convex polyhedral cone. By Definition 3.1.6 and the argument following Proposition 3.1.4 in the previous section all sets in the Gröbner fan are in fact cones. We now argue that the relative interior of every cone in the Gröbner fan is an equivalence class.

Lemma 3.1.14 *Let \prec be a term order. If $v \in \mathbb{R}_{\geq 0}^n$ then $v \in C_{\prec_v}(I)$.*

Proof. This follows from Corollary 3.1.10 since $\text{in}_{\prec_v}(\text{in}_v(g)) = \text{in}_{\prec_v}(g)$ for all $g \in \mathcal{G}_{\prec_v}(I)$. \square

Corollary 3.1.15 *If \prec is a term order and $v \in \mathbb{R}_{\geq 0}^n$ then*

$$\text{in}_{\prec_v}(I) = \text{in}_{\prec}(\text{in}_v(I)).$$

Proof. By Lemma 3.1.14 $v \in C_{\prec_v}(I)$. By Lemma 3.1.12 $\text{in}_{\prec_v}(I) = \text{in}_{\prec_v}(\text{in}_v(I))$. The Gröbner basis $\mathcal{G}_{\prec_v}(\text{in}_v(I))$ is v -homogeneous and thereby also a Gröbner basis with respect to \prec with the same initial terms which generate the initial ideal $\text{in}_{\prec_v}(\text{in}_v(I)) = \text{in}_{\prec}(\text{in}_v(I))$. \square

In the sense of Remark 2.2.7 the requirement $v \in \mathbb{R}_{> 0}^n$ is not needed in Corollary 3.1.15 and Lemma 3.1.14 if I is homogeneous. Recall also Remark 3.1.1.

Proposition 3.1.16 *The relative interior of a cone in the Gröbner fan is an equivalence class (with respect to $u \sim u' \Leftrightarrow \text{in}_u(I) = \text{in}_{u'}(I)$).*

Proof. By definition every cone in the fan is the face of the closure of an equivalence class for a positive vector $v \in \mathbb{R}_{> 0}^n$. Let \prec' be an arbitrary term order and define \prec as \prec'_v . According to Lemma 3.1.14 the vector v belongs to $C_{\prec}(I)$. Notice that by (3.2), $C_v(I) \subseteq C_{\prec}(I)$ since for all $u \in C_v(I)$ and $g \in \mathcal{G}_{\prec}(I)$, the condition $\text{in}_{\prec}(\text{in}_u(g)) = \text{in}_{\prec}(\text{in}_v(\text{in}_u(g))) = \text{in}_{\prec}(\text{in}_v(g)) = \text{in}_{\prec}(g)$ of Corollary 3.1.10 is satisfied. By (3.2) the closed set $C_v(I)$ is cut out by some equations and non-strict inequalities. The relative interior of any face of $C_v(I)$ can be formed from this inequality system by changing a subset of the inequalities to strict inequalities and the remaining ones to equations. So let u be a vector in the relative interior of some face of $C_v(I)$. The vector u is in $C_v(I) \subseteq C_{\prec}(I)$. We may use Proposition 3.1.4 to conclude that a vector $u' \in \mathbb{R}^n$ is equivalent to u if and only if it satisfies the inequality system mentioned above — that is, if and only if it is in the relative interior of the face. \square

It remains to be shown that the intersection of two cones in the Gröbner fan is a face of both cones (Proposition 3.1.18). We need a few observations.

Corollary 3.1.17 *Let C be a cone in the Gröbner fan. If $v \in C$ then for $u \in \mathbb{R}^n$,*

$$\text{in}_u(I) = \text{in}_v(I) \Rightarrow u \in C.$$

Proof. The vector v is in the relative interior of some face of C . This face is also in the Gröbner fan. By Proposition 3.1.16 u is in the relative interior of the same face and, consequently, also in C . \square

By Remark 3.1.2 there are only finitely many initial ideals given by term orders and, consequently, only finitely many reduced Gröbner bases of I . It follows that there can only be finitely many equivalence classes of the type described in Proposition 3.1.4 and Proposition 3.1.16.

Proposition 3.1.18 *Let C_1 and C_2 be two cones in the Gröbner fan of I . Then the intersection $C_1 \cap C_2$ is a face of C_1 .*

Proof. The intersection $C_1 \cap C_2$ is a cone. By Corollary 3.1.17, C_1 and C_2 are unions of equivalence classes. Further, if $v \in C_1 \cap C_2$, then again by Corollary 3.1.17, the entire equivalence class of v is both in C_1 and in C_2 and hence in $C_1 \cap C_2$. Hence $C_1 \cap C_2$ is a union of equivalence classes.

Let u be a vector in such an equivalence class E contained in $C_1 \cap C_2$. Then u is in the relative interior of one of the faces of C_1 which is a cone in the Gröbner fan. By Proposition 3.1.16 the set of vectors in the relative interior of this face is exactly E . Hence every such equivalence class is the relative interior of a face of C_1 and its closure is the face.

Look at the \mathbb{R} -span of each equivalence class contained in $C_1 \cap C_2$. These spans must be different for every face of C_1 . We claim that there can be only one maximal dimensional cone/span. If there were two cones then their convex hull would be in $C_1 \cap C_2$ and have dimension at least one higher and thus cannot be covered by the finitely many lower dimensional equivalence classes — a contradiction.

Let E be the maximal dimensional equivalence class contained in $C_1 \cap C_2$. We will argue that $\overline{E} = C_1 \cap C_2$. The inclusion $\overline{E} \subseteq C_1 \cap C_2$ is already clear since $C_1 \cap C_2$ is closed. To prove the other inclusion suppose $\omega \in C_1 \cap C_2 \setminus \overline{E}$. Then $\text{conv}(\overline{E}, \omega) \setminus \overline{E}$ is contained in $C_1 \cap C_2$ and has dimension at least the dimension of E . This is a contradiction since $\text{conv}(\overline{E}, \omega) \setminus \overline{E}$ cannot be covered by finitely many lower dimensional equivalence classes. This completes the proof. \square

Theorem 3.1.19 *The Gröbner fan is a polyhedral complex of cones and hence a fan.*

Proof. We already argued using Proposition 3.1.4 and Lemma 3.1.14 that the Gröbner fan consists of polyhedral cones. The first condition for being a polyhedral complex is satisfied by definition. The second condition is Proposition 3.1.18. \square

We finish this section by a proof of the following very important proposition which follows from Corollary 3.1.10 and Corollary 3.1.13.

Proposition 3.1.20 [53, Proposition 1.13] *Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal and $u, v \in \mathbb{R}^n$. Furthermore, suppose that I is homogeneous or $u \in \mathbb{R}_{>0}^n$. Then for $\varepsilon > 0$ sufficiently small*

$$\text{in}_{u+\varepsilon v}(I) = \text{in}_v(\text{in}_u(I)).$$

Proof. Fix a term order \prec . The key observation is that for $\varepsilon > 0$ sufficiently small we have $u + \varepsilon v \in C_{(\prec_v)_u}(I)$. This follows from Corollary 3.1.10 and an argument similar to the proof of Lemma 2.2.9. The proposition now follows from Corollary 3.1.13:

$$\text{in}_{u+\varepsilon v}(I) = \langle \text{in}_{u+\varepsilon v}(g) : g \in \mathcal{G}_{(\prec_v)_u}(I) \rangle = \langle \text{in}_v(\text{in}_u(g)) : g \in \mathcal{G}_{(\prec_v)_u}(I) \rangle$$




Figure 3.4: A graph. The same graph with a cycle-free orientation and a unique sink. The same graph cut down to a tree by removing edges until no vertex has more than one outgoing edge.

$$= \langle \text{in}_v(g) : g \in \mathcal{G}_{\prec_v}(\text{in}_u(I)) \rangle = \text{in}_v(\text{in}_u(I)).$$

The second equality holds for $\varepsilon > 0$ sufficiently small since $\mathcal{G}_{\prec_v}(\text{in}_u(I))$ is finite. \square

Remark 3.1.21 It follows that $\dim(C_u(I)) = \text{homog}(\text{in}_u(I))$ under the same assumptions. To see this, observe that the set of v 's keeping the right hand side of the equation in Proposition 3.1.20 equal to $\text{in}_u(I)$ is $C_0(\text{in}_u(I))$. On the other hand, the vectors keeping the left hand side equal to $\text{in}_v(I)$ generate the span of $C_\omega(I)$.

3.2 Reverse search property

By the *graph* of a pure full-dimensional fan we mean the set of maximal cones with two cones being connected if they share a common facet. In this section we will prove that the *reverse search technique* [3] can be used for traversing the graph of a Gröbner fan. This follows from the main theorem, Theorem 3.2.6, which says that the graph of a Gröbner fan can be oriented easily without cycles and with a unique sink. In Definition 3.2.4 we define what we mean by this.

We start by explaining how a graph with this special kind of orientation can be traversed by reverse search. The idea is to define a spanning tree of the graph which can be easily traversed. The following is a simple but very useful proposition which we shall not prove.

Proposition 3.2.1 *Let $G = (V, E)$ be an oriented graph without cycles and with a unique sink s . If for every vertex $v \in V \setminus \{s\}$ some outgoing search edge $e_v = (v, \cdot)$ is chosen then the set of chosen search edges is a spanning tree for G .*

The spanning tree in Proposition 3.2.1 is referred to as the *search tree*. The proposition implies that the graph is connected. An example is given in Figure 3.2

Notice that we can find the sink by starting at any vertex and walking along a unique path of search edges until we get stuck, in which case we are at

the sink. Consequently, the sink is the root of the oriented spanning tree. A corollary to the proposition is the *reverse search algorithm* for traversing G :

Algorithm 3.2.2 Let $G = (V, E)$ be the oriented graph of Proposition 3.2.1 and suppose the choice of a search edge e_v for each vertex $v \neq s$ has been made. Calling the following recursive procedure with $v = s$ will output all vertices in G .

Output_subtree(v)

Input: A vertex v in the graph G .

Output: The set of vertices in the subtree with root v .

```
{
  Output  $v$ ;
  Compute the edges of form  $(\cdot, v) \in E$ ;
  For every oriented edge  $(u, v) \in E$ 
    If  $(e_u = (u, v))$  Output_subtree( $u$ );
}
```

This algorithm does not have to store a set of “active” vertices as is usually needed in depth- and breadth-first traversals. It is even possible to formulate the algorithm completely without recursion avoiding the need for a recursion stack. In that sense the algorithm is *memory-less*. A real world analogue of this memory-less method is the strategy for walking through a labyrinth by keeping your right hand on the wall.

We give an example of how the edge graph of a polytope or, equivalently, the graph of its normal fan can be oriented.

Example 3.2.3 Let $P \subset \mathbb{R}^n$ be a polytope whose vertices have positive integer coordinates and let \prec be a term order on R . The following is an orientation of the edge graph of P without cycles and with a unique sink: An edge (p, q) is oriented from p to q if and only if $x^p \prec x^q$.

This defines an orientation of the graph of the normal fan of a polytope for any term order. We would like to mimic this orientation for any pure full-dimensional fan in \mathbb{R}^n . For simplicity we shall restrict ourselves to fans whose $(n - 1)$ -dimensional cones allow rational normals. In view of Propositions 3.1.4 and 3.1.16 this is no restriction for Gröbner fans. Have a look at Figure 3.2 while reading the following definition.

Definition 3.2.4 A pure full-dimensional fan in \mathbb{R}^n is said to have the *reverse search property* if for any term order \prec the following is an acyclic orientation of its graph with a unique sink: If (C_1, C_2) is an edge then C_1 and C_2 are n -dimensional cones with a common facet F . Let $p, q \in \mathbb{N}^n$ such that $q - p \neq 0$ is a normal for F with all points in $C_1 \setminus F$ having negative inner product with $q - p$ and all points in $C_2 \setminus F$ having positive inner product with $q - p$. We orient the edge in direction from C_1 to C_2 if and only if $x^p \prec x^q$.

Note that the orientation of an edge in Definition 3.2.4 does not depend on the particular choice of p and q . Note also that for normal fans of polytopes




Figure 3.5: A schematic drawing of the situation in Definition 3.2.4. Keep Example 3.2.3 in mind. The direction of the edge between C_1 and C_2 depends on the relative orientation of \prec (considered as a generic vector) and $p - q$.




Figure 3.6: A fan not having the reverse search property, see Example 3.2.5.

this orientation agrees with the orientation of the edge graphs of the polytopes in Example 3.2.3. Not every fan has the reverse search property:

Example 3.2.5 Figure 3.2 shows a fan with support $\mathbb{R}_{\geq 0}^3$ intersected with the standard simplex. The intersection is the non-dotted part of the figure. For every shared 2-dimensional facet the orientation of its edge with respect to a term order of form $\prec_{(1,1,1)}$ is indicated by an arrow. The graph has a cycle. The reason is that the vector $(1, 1, 1)$ is in the interior of the cone over the dotted triangle and therefore induces the shown orientation with any tie-breaking.

Example 3.2.3 on the other hand shows that any normal fan of a polytope has the reverse search property. If I is a homogeneous ideal the Gröbner fan of I is known to be the normal fan of the *state polytope* of I , see [53] for a proof. As a consequence the Gröbner fan will have the reverse search property. The reverse search orientation of a fan with respect to any term order can be carried out on any fan covering $\mathbb{R}_{\geq 0}^n$ and being the normal fan of a polyhedron. Since the restricted Gröbner fan of any 0-dimensional or principal ideal satisfies these conditions (see Chapter 5) it is clear that these fans have the reverse search property.

We have shown in [37] (see Chapter 5) that this line of reasoning cannot be

applied to Gröbner fans in general. In particular, in Theorem 5.0.1 an ideal is presented whose restricted Gröbner fan is not the normal fan of a polyhedron. For this reason we need a non-trivial argument to prove the following theorem:

Theorem 3.2.6 *The Gröbner fan of any ideal $I \subseteq R$ has the reverse search property.*

The proof is given in Section 3.2.1. In Chapter 4 we argue that all parts of Algorithm 3.2.2 (finding adjacent edges, finding adjacent vertices and finding search edges) can be implemented efficiently for Gröbner fans.

3.2.1 Proof: The Gröbner fan has the reverse search property

In this section we prove Theorem 3.2.6. We start by recalling how the polynomial ring can be graded by semigroups. This leads to a more general notion of homogeneous ideals.

Definition 3.2.7 By a *grading* on $R = k[x_1, \dots, x_n]$ we mean a pair (A, \mathcal{A}) consisting of an abelian semigroup A and a semigroup homomorphism:

$$\mathcal{A} : \mathbb{N}^n \rightarrow A$$

such that $\mathcal{A}^{-1}(a)$ is finite for all $a \in A$. The \mathcal{A} -degree of a term cx^b is $\mathcal{A}(b)$. A polynomial is \mathcal{A} -homogeneous if all its terms have the same \mathcal{A} -degree. An ideal is \mathcal{A} -homogeneous if it is generated by a set of \mathcal{A} -homogeneous polynomials.

For a grading (A, \mathcal{A}) on R we get the direct sum of k -vector spaces

$$R = \bigoplus_{a \in A} R_a$$

where R_a denotes the k -subspace of R consisting of \mathcal{A} -homogeneous polynomials of degree a . Any reduced Gröbner basis of an \mathcal{A} -homogeneous ideal I consists of \mathcal{A} -homogeneous polynomials. In particular, by generalizing the argument of Lemma 3.1.11 we get the direct sum

$$I = \bigoplus_{a \in A} I_a$$

where I_a denotes the k -subspace of I consisting of \mathcal{A} -homogeneous polynomials of degree a . The \mathcal{A} -homogeneous part I_a is a k -subspace of R_a . We define the \mathcal{A} -graded Hilbert function:

$$H_{I, \mathcal{A}} : A \rightarrow \mathbb{N} \tag{3.4}$$

$$a \mapsto \dim_k(R_a/I_a) \tag{3.5}$$

Remark 3.2.8 For a monomial ideal I the standard monomials of degree a form a basis for R_a/I_a . Hence $H_{I, \mathcal{A}}(a)$ counts the number of standard monomials of degree a .

In general, as the following well-known proposition shows, the Hilbert function can be found by looking at a monomial initial ideal:

Proposition 3.2.9 *Let I be an \mathcal{A} -homogeneous ideal and \prec a term order then*

$$H_{I,\mathcal{A}} = H_{\text{in}_{\prec}(I),\mathcal{A}}.$$

Proof. The linear map taking a polynomial to its unique normal form by the division algorithm on $\mathcal{G}_{\prec}(I)$ induces an isomorphism of k -vector spaces

$$R_a/I_a \rightarrow R_a/\text{in}_{\prec}(I)_a.$$

□

Consider a shared facet of the cones C_1 and C_2 in the Gröbner fan with a relative interior point v . The “edge ideal” $\text{in}_v(I)$ is homogeneous with respect to any vector in the relative interior of the facet and consequently also homogeneous with respect to any vector in the span of the facet. Since C_1 and C_2 both contain positive vectors, so does $\text{span}_{\mathbb{R}}(C_v(I))$. Recall that $C_v(I)$ is the closure of the equivalence class of v . Pick a basis $u_1, \dots, u_{n-1} \in \mathbb{N}^n$ for $\text{span}_{\mathbb{R}}(C_v(I))$ with u_1 being a positive vector. The vectors induce a grading $\mathcal{A}_v : \mathbb{N}^n \rightarrow \mathbb{N}^{n-1}$ on R by

$$\mathcal{A}_v(b) = (\langle u_1, b \rangle, \dots, \langle u_{n-1}, b \rangle)$$

for $b \in \mathbb{N}^n$. The initial ideal $\text{in}_v(I)$ is \mathcal{A}_v -homogeneous.

Lemma 3.2.10 *Let \prec be a term order, I an ideal, (C_1, C_2) a directed edge with respect to the orientation in Definition 3.2.4 and M_1 and M_2 the initial ideals of C_1 and C_2 respectively. Let v be a relative interior point in the shared facet. Then $\text{in}_{\prec}(\text{in}_v(I)) = M_2$.*

Proof. Choose a positive interior point ω_2 of C_2 . We claim that the following identities hold:

$$M_2 = \text{in}_{\omega_2}(I) = \text{in}_{\prec_{\omega_2}}(\text{in}_{\omega_2}(I)) = \text{in}_{\prec_{\omega_2}}(I) = \text{in}_{\prec_{\omega_2}}(\text{in}_v(I)) = \text{in}_{\prec}(\text{in}_v(I)).$$

The first one holds by the choice of ω_2 . The second one is clear since $\text{in}_{\omega_2}(I)$ is a monomial ideal. The third one holds by Lemma 3.1.12 and Lemma 3.1.14. By Lemma 3.1.12 the fourth equality holds since $v \in C_{\prec_{\omega_2}}(I) = C_{\omega_2}(I)$. To prove the last equality we look at the reduced Gröbner basis $\mathcal{G}_{\prec}(\text{in}_v(I))$. If we can show that $\text{in}_{\prec_{\omega_2}}(g) = \text{in}_{\prec}(g)$ for all elements $g \in \mathcal{G}_{\prec}(\text{in}_v(I))$ then we know that $\mathcal{G}_{\prec}(\text{in}_v(I))$ is also a Gröbner basis with respect to \prec_{ω_2} and the generators for the initial ideal $\text{in}_{\prec_{\omega_2}}(\text{in}_v(I))$ are exactly the same as those for $\text{in}_{\prec}(\text{in}_v(I))$. This would complete the proof.

The reduced Gröbner basis $\mathcal{G}_{\prec}(\text{in}_v(I))$ is \mathcal{A}_v -homogeneous. For an element g this implies that the difference between two of its exponent vectors must be perpendicular to the shared facet. By Definition 3.2.4 there exists a normal $q - p$ of the facet with $x^p \prec x^q$ and $\langle \omega_2, q - p \rangle > 0$. Since \prec and \prec_{ω_2} agree on one normal vector they must agree on all exponent differences of elements in $\mathcal{G}_{\prec}(\text{in}_v(I))$. □

Notice that by Proposition 3.2.9 any initial ideal $\text{in}_{\prec}(\text{in}_v(I))$ of $\text{in}_v(I)$ has the same \mathcal{A}_v -graded Hilbert function as $\text{in}_v(I)$.

By a flip we mean a move from one vertex in the graph to a neighbor. For a degree $a \in \mathbb{N}^{n-1}$ we call $\mathcal{A}_v^{-1}(a)$ the *fiber* over a . The \mathcal{A}_v -graded Hilbert function of an initial ideal $\text{in}_{\prec}(\text{in}_v(I))$ counts the number of standard monomials inside each fiber. A flip preserves the Hilbert function. We may think of this as monomials in the monomial initial ideal moving around in the fiber. We wish to keep track of how the monomials move when we walk in the oriented graph. We define exactly what we mean by “moving around”:

Definition 3.2.11 Let $\prec, M_1, M_2, u_1, \dots, u_{n-1}$ and v be as above with $\text{in}_{\prec}(\text{in}_v(I)) = M_2$. Let N_1 and N_2 be the monomials in M_1 and M_2 respectively. We define the bijection $\phi_{\prec M_1 M_2} : N_1 \rightarrow N_2$ in the following way: For a monomial $x^b \in N_1$ look at the monomials $B_1 \subseteq N_1$ and $B_2 \subseteq N_2$ with the same \mathcal{A} -degree as x^b . Since taking initial ideals preserves the \mathcal{A} -graded Hilbert function, $|B_1| = |B_2|$. Sort B_1 and B_2 with respect to \prec . The bijection $\phi_{\prec M_1 M_2}$ is now defined by taking the first element of B_1 to the first element of B_2 , the second element of B_1 to the second element of B_2 and so on.

The following lemma is from [41, Lemma 4.1]:

Lemma 3.2.12 Let \leq_1 and \leq_2 be two term orders. If f_1^1, \dots, f_s^1 is a vector space basis for I_a such that $\text{in}_{\leq_1}(f_1^1), \dots, \text{in}_{\leq_1}(f_s^1)$ is a basis for $\text{in}_{\leq_1}(I)_a$, then there exists a basis f_1^2, \dots, f_s^2 for I_a such that $\text{in}_{\leq_2}(f_1^2), \dots, \text{in}_{\leq_2}(f_s^2)$ is a basis for $\text{in}_{\leq_2}(I)_a$ and

$$\begin{aligned} \text{in}_{\leq_2}(f_1^2) &\leq_1 \text{in}_{\leq_1}(f_1^1) \\ &\vdots \\ \text{in}_{\leq_2}(f_s^2) &\leq_1 \text{in}_{\leq_1}(f_s^1). \end{aligned}$$

Proof. Without loss of generality we may assume that $\text{in}_{\leq_1}(f_1^1) \leq_1 \dots \leq_1 \text{in}_{\leq_1}(f_s^1)$. For $i = 1, \dots, s$ we let f_i^2 be the remainder of the division of f_i^1 with $\{f_1^2, \dots, f_{i-1}^2\}$ with respect to \leq_2 . Linear independence guarantees that $f_i^2 \neq 0$. The monomials appearing in f_i^2 must have appeared in f_1^1, \dots, f_i^1 . Thus using the chosen ordering of the polynomials f_1^1, \dots, f_s^1 we see that all monomials in f_i^2 must be less than or equal to $\text{in}_{\leq_1}(f_i^1)$ with respect to \leq_1 . In particular, $\text{in}_{\leq_2}(f_i^2) \leq_1 \text{in}_{\leq_1}(f_i^1)$. Since the monomials of $\text{in}_{\leq_2}(f_1^2), \dots, \text{in}_{\leq_2}(f_s^2)$ are different they generate the s -dimensional space $\text{in}_{\leq_2}(I)_a$. \square

Corollary 3.2.13 Let the setting be as in Definition 3.2.11. If $x^b \in M_1$ then $\phi_{\prec M_1 M_2}(x^b) \neq x^b$.

Proof. Let a be the \mathcal{A} -degree of x^b . We apply Lemma 3.2.12 with I in the lemma being $\text{in}_v(I)$. Let \leq_1 be \prec and \leq_2 be the refinement of the pre-order induced by u_1 with the reversed order of \prec . By the orientation of the graph $M_1 = \text{in}_{\leq_2}(\text{in}_v(I))$ and $M_2 = \text{in}_{\leq_1}(\text{in}_v(I))$. By multiplying elements of $\mathcal{G}_{\prec}(\text{in}_v(I))$ by monomials we can construct a k -basis f_1^1, \dots, f_s^1 of $\text{in}_v(I)_a$ with $\text{in}_{\leq_1}(f_1^1), \dots, \text{in}_{\leq_1}(f_s^1)$ being a basis of $(M_2)_a$. By the lemma there is a basis

$\text{in}_{\leq_2}(f_1^2), \dots, \text{in}_{\leq_2}(f_s^2)$ of $(M_1)_a$. Sort the list of inequalities in the lemma with $\text{in}_{\leq_2}(f_i^2)$ decreasing w.r.t. $\prec (\leq_1)$. The right hand side can now be sorted with respect to the same order without violating the inequalities. To see this use the bubble sort algorithm — when two adjacent inequalities are swapped ...

$$\begin{array}{ccc} \text{in}_{\leq_2}(f_i^2) & \leq_1 & \text{in}_{\leq_1}(f_i^1) & \quad & \text{in}_{\leq_2}(f_i^2) & \leq_1 & \text{in}_{\leq_1}(f_{i+1}^1) \\ \vee_1 & & \wedge_1 & \mapsto & \vee_1 & & \vee_1 \\ \text{in}_{\leq_2}(f_{i+1}^2) & \leq_1 & \text{in}_{\leq_1}(f_{i+1}^1) & & \text{in}_{\leq_2}(f_{i+1}^2) & \leq_1 & \text{in}_{\leq_1}(f_i^1) \end{array}$$

... the relations on the right hand side of the arrow hold by transitivity of \leq_1 .

After sorting, x^b appears somewhere on the left and $\phi_{\prec M_1 M_2}(x^b)$ on the right in the same inequality. This completes the proof. \square

Proof of Theorem 3.2.6. Suppose C_1, C_2, \dots, C_m was a path in the oriented graph with $C_1 = C_m$. Let M_1, \dots, M_m denote the initial ideals and N_1, \dots, N_m their monomials. We will prove that the bijection $\phi := \phi_{\prec M_{m-1} M_m} \circ \dots \circ \phi_{\prec M_1 M_2}$ is the identity on M_1 . Suppose it is not the identity and let x^b be the smallest element in M_1 with respect to \prec that is not fixed by ϕ . By Corollary 3.2.13, x^b is the image of a smaller element in M_1 with respect to \prec . But this element is fixed by the minimality of x^b — a contradiction. The composition being the identity implies by Corollary 3.2.13 that $\phi_{\prec M_i M_{i+1}}$ is the identity for all i . Hence $M_i = M_{i+1}$, contradicting that M_1, M_2, \dots, M_m is a path.

We claim that $C_{\prec}(I)$ is the unique sink. If v is in the relative interior of a facet of $C_{\prec}(I)$ then by Lemma 3.1.12 $\text{in}_{\prec}(\text{in}_v(I)) = \text{in}_{\prec}(I)$. By Lemma 3.2.10 this means that all edges connected to $C_{\prec}(I)$ are ingoing. Hence $C_{\prec}(I)$ is a sink.

To prove uniqueness let $C_{\prec'}(I)$ be some sink in the oriented graph. By Theorem 2.2.10, \prec has a matrix representation $(\tau_0, \dots, \tau_{n-1}) \in \mathbb{R}^{n \times n}$ such that $\tau_\varepsilon := \tau_0 + \varepsilon \tau_1 + \dots + \varepsilon^{n-1} \tau_{n-1} \in \text{int } C_{\prec}(I)$ for $\varepsilon > 0$ sufficiently small. Furthermore, for any $f \in R$, $\text{in}_{\tau_\varepsilon}(f) = \text{in}_{\prec}(f)$ for $\varepsilon > 0$ sufficiently small. If $C_{\prec'}(I)$ is a sink then according to Definition 3.2.4 there exists a complete list of inner normals $q_1 - p_1, \dots, q_r - p_r$ of $C_{\prec'}(I) \cap \mathbb{R}_{\geq 0}^n$ such that $\text{in}_{\prec}(x^{q_i} - x^{p_i}) = x^{q_i}$. Since τ_ε and \prec pick out the same initial forms on a finite set of polynomials for $\varepsilon > 0$ sufficiently small we see that $\langle \tau_\varepsilon, q_i \rangle > \langle \tau_\varepsilon, p_i \rangle$ or, equivalently, $\tau_\varepsilon \in \text{int } C_{\prec'}(I)$ for $\varepsilon > 0$ sufficiently small. We conclude that $C_{\prec'}(I) = C_{\prec}(I)$. \square

Chapter 4

Algorithms for Gröbner fans

We can find a single Gröbner cone by applying Buchberger's algorithm and Corollary 3.1.10 for some term order. Since the graph of the Gröbner fan of I is connected we may choose any graph traversal algorithm for computing the full dimensional Gröbner cones. To do the local computations we need to be able to find the edges (connecting facets) of a full dimensional cone and we need to be able to find the neighbor along an edge. We will see how to do this in the following sections.

Throughout the graph enumeration process we will represent the Gröbner cones by their marked reduced Gröbner bases, rather than by their defining inequalities, their term orders etc.. This choice is justified by the following well known theorem.

Theorem 4.0.14 *Let $I \subseteq R = k[x_1, \dots, x_n]$ be an ideal. The marked reduced Gröbner bases of I , the monomial initial ideals of I (w.r.t. positive vectors) and the full-dimensional Gröbner cones are in bijection.*

Proof. We already argued in Remark 2.2.4 that the marked reduced Gröbner bases are in bijection with the set of initial ideals $\text{in}_{\prec}(I)$ where \prec is a term order. In Remark 3.1.2 we saw that any $\text{in}_{\prec}(I)$ is of the form $\text{in}_{\omega}(I)$ for some $\omega \in \mathbb{R}_{>0}^n$. On the other hand if $\text{in}_{\omega}(I)$ is monomial with $\omega \in \mathbb{R}_{>0}^n$ then $\text{in}_{\omega}(I) = \text{in}_{\prec}(\text{in}_{\omega}(I)) = \text{in}_{\prec_{\omega}}(I)$ by Corollary 3.1.15. This proves the bijection between monomial initial ideals with respect to positive vectors and marked reduced Gröbner bases.

A full-dimensional Gröbner cone contains a positive vector ω in its interior. By Proposition 3.1.16 the interior of the Gröbner cone is the equivalence class of ω . The ideal $\text{in}_{\omega}(I)$ is homogeneous with respect to any vector by Proposition 3.1.20. This shows that $\text{in}_{\omega}(I)$ is a monomial ideal. On the other hand if for some $\omega \in \mathbb{R}_{>0}^n$, the initial ideal $\text{in}_{\omega}(I)$ is monomial then by Proposition 3.1.20, ω is an interior point of its equivalence class proving that the closure is a full-dimensional Gröbner cone. \square

An important issue when implementing the algorithms is to identify shared facets. We say that a facet is *flippable* if its relative interior contains a positive vector. The flippable facets in a Gröbner fan are always shared.

One application of Gröbner fans is in Gröbner basis conversion. We explain the well-known Gröbner walk algorithm and improve the algorithm by introducing the generic Gröbner walk ([19]) which works by making a symbolic perturbation of its target vector. In the reverse search algorithm the symbolic perturbation gives an easy way to make a choice of a search edge that will always correspond to a flippable facet. At the end of the chapter we will explain how to take advantage of symmetry in a Gröbner fan traversal.

4.1 Finding facets

Suppose that we know a marked reduced Gröbner basis $\mathcal{G}_{\prec}(I)$ with respect to some unknown term order \prec . Proposition 3.1.4 (or Corollary 3.1.10) tells us how to read off the defining inequality system for $C_{\prec}(I)$.

Since $C_{\prec}(I)$ is full-dimensional the system contains no equations but only inequalities. Some of these inequalities are equivalent in the sense that they are multiples of each other. Taking just one inequality from each equivalence class the problem is now to find irredundant facet normals of a cone — or equivalently to find the extreme rays of the dual cone. Checking if a ray is extreme can be done by linear programming.

Not all of the remaining inequalities are guaranteed to define flippable facets. One way to ensure that we only get flippable facets is by adding the constraints $e_i \cdot x \geq 0$ for $i = 1, \dots, n$ and ignoring the facets defined by these.

A more efficient method (on some examples) is to find all facets and then remove the non-flippable irredundant facet normals by explicit checks. In our implementation this is done by checking if the inequality system with the inequality in question inverted still has a positive solution.

As mentioned in [32] there is an algebraic test that helps us eliminate redundant inequalities of $C_{\prec}(I)$. Let $\alpha \in \mathbb{R}^n$ be a coefficient vector of an inequality. If α indeed is irredundant and defines a facet with a relative interior point v then Corollary 3.1.13 tells us how to compute $\mathcal{G}_{\prec}(\text{in}_v(I))$. This marked reduced Gröbner basis can be computed from $\mathcal{G}_{\prec}(I)$ as $\{\text{in}_v(g)\}_{g \in \mathcal{G}_{\prec}(I)}$ if we just know α and not necessarily v ; see the next section. A necessary condition for α to be irredundant is that the computed set $\{\text{in}_v(g)\}_{g \in \mathcal{G}_{\prec}(I)}$ indeed is a marked Gröbner basis i.e. all S-polynomials reduce to zero. This check even works for v outside the positive orthant. A quicker necessary condition that we can check is that every non-zero S-polynomial should have at least one of its terms in $\text{in}_{\prec}(I)$. For huge sets of inequalities the test works extremely well — 500 inequalities might reduce to 50 of which maybe 10 are irredundant. Our experience is that having this test as a preprocessing step can be much faster than solving the full linear programs with exact arithmetic.

4.2 Local change

Let $\mathcal{G}_{\prec}(I)$ be a known marked Gröbner basis and let F be a flippable facet of $C_{\prec}(I)$. We let $\text{flip}(\mathcal{G}_{\prec}(I), F)$ denote the unique reduced Gröbner basis different from $\mathcal{G}_{\prec}(I)$ whose Gröbner cone also has F as a facet. We will describe an

algorithm for computing $\text{flip}(\mathcal{G}_{\prec}(I), F)$ given $\mathcal{G}_{\prec}(I)$ and an inner normal vector α for F . For a marked Gröbner basis \mathcal{G} and a polynomial g we let $g^{\mathcal{G}}$ denote the normal form of g modulo \mathcal{G} and note that this form does not depend on the term order but only on the markings of \mathcal{G} .

Algorithm 4.2.1

Input: A marked reduced Gröbner basis $\mathcal{G}_{\prec}(I)$ and a v -homogeneous polynomial $m \in \text{in}_v(I)$ where v is some vector in $C_{\prec}(I)$.

Output: A polynomial $f \in I$ such that $\text{in}_v(f) = m$.

$$\left\{ \begin{array}{l} f := m - m^{\mathcal{G}_{\prec}(I)}; \end{array} \right\}$$

Proof. The polynomial m reduces to 0 modulo $\mathcal{G}_{\prec}(\text{in}_v(I))$ by the division algorithm. The Gröbner basis $\mathcal{G}_{\prec}(I)$ is the same as $\mathcal{G}_{\prec}(\text{in}_v(I))$ by Corollary 3.1.13 except that each element may have additional terms of lower v -degree. Running the division algorithm on m modulo $\mathcal{G}_{\prec}(I)$ we may make the same choices as for $\mathcal{G}_{\prec}(\text{in}_v(I))$, reducing m to 0 plus additional terms of lower v -degree. Since the remainder is unique, it does not matter how we run the algorithm – we will always end up with a unique remainder $m^{\mathcal{G}_{\prec}(I)}$ of lower v -degree than the v -degree of m . We conclude that $\text{in}_v(m - m^{\mathcal{G}_{\prec}(I)}) = m$. \square

Algorithm 4.2.2 Lift

Input: Marked reduced Gröbner bases $\mathcal{G}_{\prec}(I)$ and $\mathcal{G}_{\prec'_v}(\text{in}_v(I))$ where $v \in C_{\prec}(I)$ is an unspecified vector and \prec and \prec' are unspecified term orders. Here I must be homogeneous or $v \in \mathbb{R}_{\geq 0}^n$.

Output: The marked reduced Gröbner basis $\mathcal{G} = \mathcal{G}_{\prec'_v}(I)$.

$$\left\{ \begin{array}{l} \mathcal{G} := \{g - g^{\mathcal{G}_{\prec}(I)} : g \in \mathcal{G}_{\prec'_v}(\text{in}_v(I))\}; \\ \text{Mark the term } \text{in}_{\prec'_v}(g) \text{ in each element } g - g^{\mathcal{G}_{\prec}(I)} \text{ in } \mathcal{G}; \\ \text{Turn the minimal basis } \mathcal{G} \text{ into a reduced basis;} \end{array} \right\}$$

Proof. The elements $g - g^{\mathcal{G}_{\prec}(I)}$ are constructed according to the algorithm above and thereby satisfy $\text{in}_v(g - g^{\mathcal{G}_{\prec}(I)}) = \text{in}_v(g)$. Taking \prec' initial terms on both sides we get $\text{in}_{\prec'_v}(g - g^{\mathcal{G}_{\prec}(I)}) = \text{in}_{\prec'}(\text{in}_v(g - g^{\mathcal{G}_{\prec}(I)})) = \text{in}_{\prec'}(g) = \text{in}_{\prec'_v}(g)$. This shows that the \prec'_v initial terms of \mathcal{G} generate $\text{in}_{\prec'_v}(I) = \text{in}_{\prec'}(\text{in}_v(I)) = \text{in}_{\prec'_v}(\text{in}_v(I))$ minimally. Hence \mathcal{G} is a minimal Gröbner basis for I and making it reduced we get $\mathcal{G}_{\prec'_v}(I)$. \square

Algorithm 4.2.3 Flip

Input: A marked reduced Gröbner basis $\mathcal{G}_{\prec}(I)$ with \prec being an unknown term order and an inner normal vector α of a flippable facet F of $C_{\prec}(I)$.

Output: $\mathcal{G} = \text{flip}(\mathcal{G}_{\prec}(I), F)$.

{
 Let v be a positive vector in the relative interior of F ;
 Compute $\mathcal{G}_{\prec}(\text{in}_v(I)) = \{\text{in}_v(g) : g \in \mathcal{G}_{\prec}(I)\}$;
 Compute the marked basis $\mathcal{G}_{\prec-\alpha}(\text{in}_v(I))$ from $\mathcal{G}_{\prec}(\text{in}_v(I))$
 using Buchberger's algorithm;
 Compute $\mathcal{G} := \mathcal{G}_{(\prec-\alpha)_v}(I)$ from $\mathcal{G}_{\prec}(I)$ and $\mathcal{G}_{\prec-\alpha}(\text{in}_v(I))$
 using Algorithm 4.2.2;
 }

Proof. We may apply Algorithm 4.2.2 since $v \in C_{\prec}(I)$. It remains to be shown that we have computed a Gröbner basis for the right cone. The vector $v - \varepsilon\alpha$ is in the cone we want for $\varepsilon > 0$ sufficiently small since v is a relative interior point in F . Hence, it suffices to show that $\mathcal{G}_{(\prec-\alpha)_v}(I) = \mathcal{G}_{\prec_{v-\varepsilon\alpha}}(I)$ for $\varepsilon > 0$ sufficiently small. On a finite collection of terms the term orders $(\prec-\alpha)_v$ and $\prec_{v-\varepsilon\alpha}$ agree for $\varepsilon > 0$ sufficiently small. In particular they agree on their Gröbner bases which proves that both Gröbner bases $\mathcal{G}_{(\prec-\alpha)_v}(I)$ and $\mathcal{G}_{\prec_{v-\varepsilon\alpha}}(I)$ are Gröbner bases with respect to both term orders. By uniqueness of reduced Gröbner bases we get $\mathcal{G}_{(\prec-\alpha)_v}(I) = \mathcal{G}_{\prec_{v-\varepsilon\alpha}}(I)$. \square

Algorithm 4.2.3 is a special case of the local change procedure for a single step in the Gröbner walk [13]. See [19, Proposition 3.2] for a new treatment. Here we will add a few comments on our special case — the case where F is a facet and not a lower dimensional face:

For any vector ω in the relative interior of F , $\text{in}_{\omega}(I) = \text{in}_v(I)$ and therefore $\text{in}_{\omega}(\text{in}_v(I)) = \text{in}_{\omega}(\text{in}_{\omega}(I)) = \text{in}_{\omega}(I) = \text{in}_v(I)$ which shows that $\text{in}_v(I)$ is homogeneous with respect to the ω -grading. Since F is $(n-1)$ -dimensional, $\text{in}_v(I)$ is homogeneous with respect to all vectors inside $\text{span}_{\mathbb{R}}(\alpha)^{\perp}$. All Gröbner bases of $\text{in}_v(I)$ are homogeneous in the same way. Consequently, each of them must consist of polynomials of the form $\sum_{s=0}^t c_s x^{(a+sb)}$ where $a \in \mathbb{N}^n$ and $b \in \mathbb{Z}^n$ is parallel to α . In other words their Newton polytopes are parallel line segments. The same is true for all polynomials appearing in any run of Buchberger's algorithm starting from one of these sets. A consequence is that in order to run Buchberger's algorithm we only need to decide if we are in the situation where $x^{\gamma} \prec x^{\gamma+\alpha}$ for every $\gamma \in \mathbb{N}^n$ or in the situation where $x^{\gamma+\alpha} \prec x^{\gamma}$ for every $\gamma \in \mathbb{N}^n$. Thus specifying α or $-\alpha$ as a term order suffices — no tie-breaker is needed. The initial ideal $\text{in}_v(I)$ can have at most two reduced Gröbner bases. Both term orders are legal since $\text{in}_v(I)$ is homogeneous with respect to the strictly positive vector v .

The Gröbner basis $\mathcal{G}_{\prec}(\text{in}_v(I))$ can be read off from the marked Gröbner basis $\mathcal{G}_{\prec}(I)$ by taking initial forms of the polynomials with respect to v , see Corollary 3.1.13. Taking the initial form $\text{in}_v(g)$ of a polynomial $g \in \mathcal{G}_{\prec}(I)$ without computing v is done as follows. By Corollary 3.1.10, $\text{in}_{\prec}(\text{in}_v(g)) = \text{in}_{\prec}(g)$ and thus we already know one term of $\text{in}_v(g)$ since $\text{in}_{\prec}(g)$ is the marked term of g in $\mathcal{G}_{\prec}(I)$.

Since every ω in the relative interior of F will have $\mathcal{G}_{\prec}(\text{in}_{\omega}(I)) = \mathcal{G}_{\prec}(\text{in}_v(I))$ the remaining terms of $\text{in}_v(g)$ are exactly the terms in g with the same ω -degree as $\text{in}_{\prec}(g)$ for all ω in the relative interior of F and consequently for all ω in $\text{span}_{\mathbb{R}}(\alpha)^{\perp}$. In other words a term of g is in $\text{in}_v(g)$ if and only if its exponent vector minus the exponent of $\text{in}_{\prec}(g)$ is parallel to α . The term order \prec does not have to be known for this step, nor does it have to be known in the computation of $\mathcal{G}_{\prec-\alpha}(\text{in}_v(I))$ or in any other subsequent step. The vector v also remains unknown in the entire process.

4.3 The generic Gröbner walk

Computational experience shows that Gröbner basis computations with respect to lexicographic term orders are difficult compared to for example degree reverse lexicographic term orders. Suppose that we know a reduced Gröbner basis $\mathcal{G}_{\prec'}(I)$ and want to compute $\mathcal{G}_{\prec}(I)$ for some other term order \prec . Then we could apply Algorithm 4.2.3 a number of times until we reach the Gröbner cone $C_{\prec}(I)$ in the Gröbner fan of I . This method is known as the Gröbner walk in the literature and was first presented in [13]. In this section we shall discuss the improvements to this algorithm that were presented in [19]. In Section 4.4 we see why this is relevant for Gröbner fan traversals by reverse search.

One strategy for choosing the facets to flip through in the Gröbner walk is the following. If σ is a positive vector in the interior of $C_{\prec'}(I)$ and τ is a positive vector in the interior of $C_{\prec}(I)$ then we may consider the line segment l between σ and τ . If we are lucky this segment passes through $(n-1)$ -dimensional Gröbner cones F_1, \dots, F_s connecting $C_{\prec'}(I)$ and $C_{\prec}(I)$ by a sequence of full-dimensional cones $C_{\prec'}(I) = C_0, \dots, C_s = C_{\prec}(I)$ such that $C_{i-1} \cap C_i = F_i$ for $i = 1, \dots, s$. On the other hand, if we are unlucky l might intersect two facets of $C_{\prec'}(I)$ and leave $C_{\prec'}(I)$ through a lower dimensional face. In that case the sequence of facets to flip through is not properly defined. Another problem is that usually the vector τ is not known to us a priori.

The first problem is not a serious problem since Algorithm 4.2.3 easily generalizes to the case where F is lower dimensional. However, the homogeneity space of $\text{in}_v(I)$, where v is relative interior in F , will no longer be of codimension 1 and the Buchberger step of Algorithm 4.2.3 will be more complicated since the generators of $\text{in}_v(I)$ will no longer have 0 or 1-dimensional Newton polytopes. In fact, the worst thing that could happen is that the segment passes through the homogeneity space of I , in which case $\text{in}_v(I) = I$ and the flip algorithm amounts to running Buchberger's algorithm on I .

In [58] the two problems were solved as follows. Knowing generators of the ideal I we can apply a general bound on the degree of a reduced Gröbner basis of I . Since the Gröbner cones of I have facets normals coming from differences between exponent vectors this gives a bound on the length of the integer normal vectors of the facets. If the term order \prec' is represented by a matrix with rows $\tau_1, \dots, \tau_n \in \mathbb{N}^n$ then for a sufficiently large number $M > 0$ the vector $T_M := \sum_i M^{n-i} \tau_i$ belongs to $C_{\prec'}(I)$; see Lemma 4.3.1 for a special case. Similarly, we may choose a sufficiently large number M' depending on

the degree bound which guarantees that the line segment l does not intersect Gröbner cones of dimension less than $n - 1$. The draw back of this method is that the integer vectors $T_{M'}$ are huge. In the case of 10 variables and generators for I of degree 10 and \prec being the lexicographic Gröbner basis, the entries of $T_{M'}$ have up to 10.000 digits.

Our approach to the problems is to introduce a small $\varepsilon > 0$ and perturb the target vector τ depending on this ε . It turns out that the choices we need to make when finding the facets to flip do not depend on ε as long as ε is sufficiently small.

Lemma 4.3.1 *Let $I \subseteq R$ be an ideal and \prec the lexicographic term order with $x_1 \succ x_2 \succ \dots \succ x_n$. Define $\tau_\varepsilon = (\varepsilon^0, \varepsilon^1, \dots, \varepsilon^{n-1})$. There exists $\delta > 0$ such that $\text{in}_{\tau_\varepsilon}(I) = \text{in}_\prec(I)$ for all $\varepsilon \in (0, \delta)$.*

Proof. This follows from Lemma 3.1.9 since \prec and τ_ε agree on a finite set of polynomials for small $\varepsilon > 0$. \square

Assume for simplicity that $\sigma \in \mathbb{N}^n$. For sufficiently small $\varepsilon > 0$ the line segment

$$\omega(t) := (1 - t)\sigma + t\tau_\varepsilon \text{ with } t \in [0, 1]$$

intersects a facet of $C_{\prec'}(I)$ unless $C_{\prec'}(I) = C_\prec(I)$.

Let $\{\alpha_1, \dots, \alpha_m\}$ be the set of potential inner facet normals of $C_{\prec'}(I)$ read off from $\mathcal{G}_{\prec'}(I)$. We are only interested in the vectors α_i where $\langle \sigma, \alpha_i \rangle > 0$ and $\langle \tau_\varepsilon, \alpha_i \rangle < 0$. Let t_i denote the t -value for the intersection of the line segment and the hyperplane defined by α_i . Then

$$t_i := \frac{\langle \sigma, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle - \langle \tau_\varepsilon, \alpha_i \rangle}.$$

We wish to find i such that t_i is smallest (for small ε).

$$t_i < t_j \iff \tag{4.1}$$

$$\frac{\langle \sigma, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle - \langle \tau_\varepsilon, \alpha_i \rangle} < \frac{\langle \sigma, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle - \langle \tau_\varepsilon, \alpha_j \rangle} \iff \tag{4.2}$$

$$\frac{\langle \sigma, \alpha_i \rangle - \langle \tau_\varepsilon, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle} > \frac{\langle \sigma, \alpha_j \rangle - \langle \tau_\varepsilon, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle} \iff \tag{4.3}$$

$$\frac{\langle \tau_\varepsilon, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle} < \frac{\langle \tau_\varepsilon, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle} \iff \tag{4.4}$$

$$\langle \tau_\varepsilon, \langle \sigma, \alpha_j \rangle \alpha_i \rangle < \langle \tau_\varepsilon, \langle \sigma, \alpha_i \rangle \alpha_j \rangle \iff \tag{4.5}$$

$$x^{\langle \sigma, \alpha_j \rangle \alpha_i} \prec x^{\langle \sigma, \alpha_i \rangle \alpha_j} \tag{4.6}$$

We see that for ε sufficiently small “ $t_i < t_j$ ” does not depend on ε . Furthermore, there cannot be any ties, unless α_i and α_j represent the same hyperplane. This gives an easy method for computing the unique first facet of $C_{\prec'}(I)$ that the perturbed line intersects. We simply choose the facet defined by α_i where t_i is smallest among $\{t_1, \dots, t_m\}$ (for small $\varepsilon > 0$).

Moving a long to the next cone C_1 on the path from $C_{\prec'}(I)$ to $C_{\prec}(I)$ the same argument applies.

Recall Theorem 2.2.10 which states that any term order can be represented by a real matrix and Lemma 2.2.9 which together with Lemma 3.1.9 gives a natural generalization of Lemma 4.3.1 with $\tau_\varepsilon := \sum_{i=1}^n \varepsilon^{i-1} A_i$ for a suitable matrix A . It turns out that the argument above for independence of ε works for any such term order. This is worked out in [19]. In practise, of course, we are mainly concerned with term orders representable by integer matrices.

In [19] the perturbation strategy is taken even further and both the starting point and the target point are perturbed avoiding the need for computing a relative interior point in the starting cones. For our purposes, however, perturbing the end point suffices.

The idea of using perturbations is well-known in computational geometry. In fact perturbation techniques can be used as anti-cycling rules in the simplex algorithm for linear programming.

4.4 Computing the search edge

Let \prec be the lexicographic term order and use this term order for orienting the graph of the Gröbner fan. In Algorithm 3.2.2 the search edge $e_{C_{\prec'}(I)}$ has to be computed given $\mathcal{G}_{\prec'}(I)$ where \prec' is some unspecified term order. According to Proposition 3.2.1 the definition of search edges can be arbitrary. However, efficiently computing a search edge requires a good definition. Our search edges will always come from flippable facets.

One strategy for locally computing the search edge $e_{C_{\prec'}(I)}$ is to compute a unique representation of each flippable facet of the Gröbner cone $C_{\prec'}(I)$ and then choose the smallest of these facets to be $e_{C_{\prec'}(I)}$ in some lexicographic order. This method requires all facets to be computed every time we check if “ $e_u = (u, v)$ ” in Algorithm 3.2.2.

A better strategy is to apply a single step of the generic Gröbner walk starting in $C_{\prec'}(I)$ and heading towards $C_{\prec}(I)$. A point in the cone $C_{\prec'}(I)$ can be computed deterministically by linear programming. The generic Gröbner walk now deterministically defines the search edge. This strategy explains the name “reverse search”.

Let us consider how Algorithm 3.2.2 determines the edges of the search tree. To find the outgoing edge of a vertex, a single step in the generic Gröbner walk is performed. To find all ingoing edges to a vertex v it first computes all flippable facets of v . Then for each facet it computes the neighboring Gröbner cone u using Algorithm 4.2.3 and asks if (u, v) is the search edge of u by applying a step of the walk. If so, the edge is an ingoing edge of v , otherwise it is not.

4.5 Exploiting symmetry

In this section we explain how to take advantage of symmetry to speed up computations. The symmetric group S_n acts on polynomials and ideals of R by permuting variables and on \mathbb{R}^n by permuting coordinate entries. Let $I \subseteq R$

be an ideal. We call a subgroup $\Gamma \leq S_n$ a *symmetry group* for I if $\pi(I) = I$ for all $\pi \in \Gamma$. If we know a symmetry group for I we can enumerate the reduced Gröbner bases of I up to symmetry. Let Γ be such a symmetry group for I .

In our description all Gröbner bases will be marked and reduced. Thereby each one will uniquely represent its initial ideal and Gröbner cone. For a Gröbner basis \mathcal{G} of I we use the notation $\Gamma_{\mathcal{G}} = \{\pi(\mathcal{G})\}_{\pi \in \Gamma}$ for its orbit.

The idea is to exploit the identity $\text{flip}(\pi(\mathcal{G}), \pi(F)) = \pi(\text{flip}(\mathcal{G}, F))$ for all $\pi \in \Gamma$. In other words Γ is a group of automorphisms of the graph of the Gröbner fan of I . The *quotient graph* is defined to be the graph whose vertices are the orbits of Gröbner bases with two orbits $\Gamma_{\mathcal{G}}$ and $\Gamma_{\mathcal{G}'}$ being connected if there exists a facet F of the Gröbner cone of \mathcal{G} such that $\text{flip}(\mathcal{G}, F) \in \Gamma_{\mathcal{G}'}$. The flip graph may have loops.

The symmetry-exploiting algorithm enumerates the quotient graph by a breadth-first traversal. Orbits are represented by Gröbner basis representatives. One question that arises is how to check if two Gröbner bases \mathcal{G} and \mathcal{G}' represent the same orbit. A solution is to run through all elements $\pi \in \Gamma$ and check if $\pi(\mathcal{G})$ equals \mathcal{G}' , or even better to make a similar check for the monomial initial ideals. Although this does not seem efficient, it is still much faster in practice than redoing symmetric Gröbner basis and polyhedral computation as we would have done in the usual reverse search or breadth-first enumeration without symmetry. An example of a 100 fold speed up is given in Section 9.2. It is not clear how to combine symmetry-exploiting with reverse search.

Chapter 5

A non-regular Gröbner fan

In light of Section 3.2, and Example 3.2.3 in particular, an important fundamental question to ask is the following: *Is the Gröbner fan always the normal fan of a polytope?* The answer to this question is *no* since the Gröbner fan is not always complete. Thus we rephrase the question for restricted Gröbner fans: *Is the restricted Gröbner fan of an ideal always the normal fan of a polyhedron?*

We say that a fan is *regular* if it is the normal fan of a polyhedron. We note that the Gröbner fan being regular is stronger than the restricted Gröbner fan being so. This is because the normal fan of the Minkowski sum of two polyhedra is the common refinement of their normal fans. The claim follows since $\mathbb{R}_{\geq 0}^n$ with its proper faces is the normal fan of $\mathbb{R}_{\leq 0}^n$.

The above question is known to have a positive answer in the following three special cases:

- If the ideal is homogeneous the answer is *yes* since the Gröbner fan is the normal fan of the *state polytope* of I introduced by Bayer and Morrison in [5]. A construction of the state polytope is given in [53, Chapter 2]. We take the Minkowski sum of the state polytope with $\mathbb{R}_{\leq 0}^n$ to get a polyhedron having the restricted Gröbner fan as its normal fan.
- In the case of a *principal* ideal $I = \langle f \rangle$ the Newton polytope $\text{New}(f)$ will almost have the Gröbner fan as its normal fan since two vectors $u, v \in \mathbb{R}^n$ pick out the same initial ideal of I if and only if they are maximized on the same face of $\text{New}(f)$. The only thing that keeps $\text{New}(f)$ from having the Gröbner fan of I as its normal fan is that we have not included all equivalence classes in the Gröbner fan. However, the normal fan of the Minkowski sum of $\text{New}(f)$ and $\mathbb{R}_{\leq 0}^n$ is the restricted Gröbner fan.
- A third case where we have a similar result is for *zero-dimensional* ideals. The construction of a polytope is similar but simpler than the construction in the homogeneous case as there is only a finite number of standard monomials for each initial ideal. We claim, without proof, that the following construction works: For every term order \prec construct the vector v_{\prec} equal to the negative of the sum of all exponent vectors of all standard monomials of $\text{in}_{\prec}(I)$. Take the convex hull of all v_{\prec} as we vary the

term order. The Minkowski sum of this polytope with $\mathbb{R}_{\geq 0}^n$ is a polyhedron whose normal fan is the restricted Gröbner fan. Under a certain genericity condition this construction appeared in [46].

In contrast to the above, we have the following theorem which is the main result of [37].

Theorem 5.0.1 *The restricted Gröbner fan of the two-dimensional ideal*

$$I = \langle acd + a^2c - ab, ad^2 - c, ad^4 + ac \rangle \subset \mathbb{Q}[a, b, c, d]$$

is not the normal fan of a polyhedron.

This theorem shows that we cannot be sure, a priori, that Gröbner fans have the reverse search property by referring to an underlying polytope; see Example 3.2.3. For this reason Theorem 3.2.6 requires the non-trivial proof given in Section 3.2.1.

In the following subsections we will prove Theorem 5.0.1, explain how the example was found and investigate what happens when we homogenize the ideal.

5.1 The proof

This section contains a proof of Theorem 5.0.1. We start by deducing a necessary condition for a fan to be the normal fan of a polyhedron. We then show that the restricted Gröbner fan of the ideal in the theorem violates this condition. Finally we argue that the Gröbner fan has been computed correctly.

5.1.1 A necessary condition

Let F be a fan in \mathbb{R}^n . Suppose F is the normal fan of a polyhedron $P \subset \mathbb{R}^n$. The non-empty faces of P are in bijection with the cones in F by taking normal cones of the faces. Adjacency is preserved in the sense that two vertices of an edge of P map to cones in F having the normal cone of the edge as a common facet. Furthermore, the edge is perpendicular to the shared facet. If a set of normals of the shared facets in F are specified, then for every bounded edge the difference between its endpoints can be expressed as some scalar times the specified normal of its normal cone. The scalars are considered to be unknowns. Since the adjacency information of the vertices of P is present in F , the bounded edge graph of P can be deduced from F . A necessary condition for F to be the normal fan of P is that every combinatorial cycle in the edge graph is a geometric cycle in space. This condition gives rise to a feasible system of inequalities on the scalars dependent on F alone.

To be more specific about the inequality system, consider the adjacency graph of the n -dimensional cones in F , or equivalently the edge graph of the supposed polyhedron P . Let $V = \{1, \dots, m\}$ denote the vertices and a subset $E \subset \{(i, j) \in V \times V : i < j\}$ denote the edges in the graph. For each shared facet, choose a normal vector $d_{(i,j)} \in \mathbb{R}^n$ such that the i th cone is on the

negative side of the hyperplane with normal vector $d_{(i,j)}$ and the j th cone is on the positive side. The graph (V, E) is considered to be undirected when we define its cycles. A vector $f \in \mathbb{R}^E$ is called a *flow* in (V, E) if

$$\forall j \in V : \sum_{(i,j) \in E} f_{(i,j)} = \sum_{(j,k) \in E} f_{(j,k)}.$$

In other words the flow entering j is the same as the flow leaving j . The set of flows is a subspace of \mathbb{R}^E . We introduce a vector $s \in \mathbb{R}_{>0}^E$ of unknown scalars such that the true vector from vertex i to vertex j is $s_{(i,j)}d_{(i,j)}$. Each cycle in the graph can be represented by a flow $f \in \mathbb{R}^E$ being 0 on the edges not appearing in the cycle and ± 1 elsewhere depending on the relative orientation of the cycle and the edge. For such an f the condition that the cycle forms a loop in space can be expressed as:

$$\sum_{(i,j) \in E} f_{(i,j)}s_{(i,j)}d_{(i,j)} = 0. \quad (5.1)$$

Note that (5.1) is a system of n equations – one for each coordinate of $d_{(i,j)}$. If F is the normal fan of a polyhedron P , there exist positive scalars $s_{(i,j)}$ satisfying (5.1) for every flow f since the cycle flows span the vector space of flows. By linearity this is equivalent to having the scalars satisfy (5.1) for a basis of the vector space of flows rather than the entire space. In matrix form we may express the necessary condition as the system

$$As = 0 \quad \text{and} \quad s_{(i,j)} > 0 \quad \text{for all } (i,j) \in E \quad (5.2)$$

having a solution s where A is a suitable $nl \times |E|$ matrix with l being the dimension of the vector space of flows.

5.1.2 The certificate

Proof of Theorem 5.0.1: The restricted Gröbner fan of the ideal

$$I = \langle acd + a^2c - ab, ad^2 - c, ad^4 + ac \rangle \subset \mathbb{Q}[a, b, c, d]$$

has 81 full dimensional cones each corresponding to a monomial initial ideal. Their adjacency graph (V, E) has 163 edges, with each edge direction normal to the shared facet. The list of full-dimensional cones, reduced Gröbner bases and monomial initial ideals can be found on the webpage [35]. We present a certificate that the fan is not the normal fan of a polyhedron. Only the subgraph in Figure 5.1 is needed to describe it. Two vectors are written for each edge in the subgraph. The vector to the right is the edge direction $d_{(i,j)}$ and the vectors to the left describe four flows in the subgraph.

Let V' be the set of vertices appearing in the subgraph and E' the edges. Let f^1, f^2, f^3 and f^4 denote the flows above. Suppose the restricted Gröbner fan was the normal fan of a polyhedron P . Equality system (5.1) implies

$$\forall (r, t) \in \{1, 2, 3, 4\} \times \{1, 2, 3, 4\} : \sum_{(i,j) \in E'} f_{(i,j)}^r s_{(i,j)} d_{(i,j)_t} = 0. \quad (5.3)$$

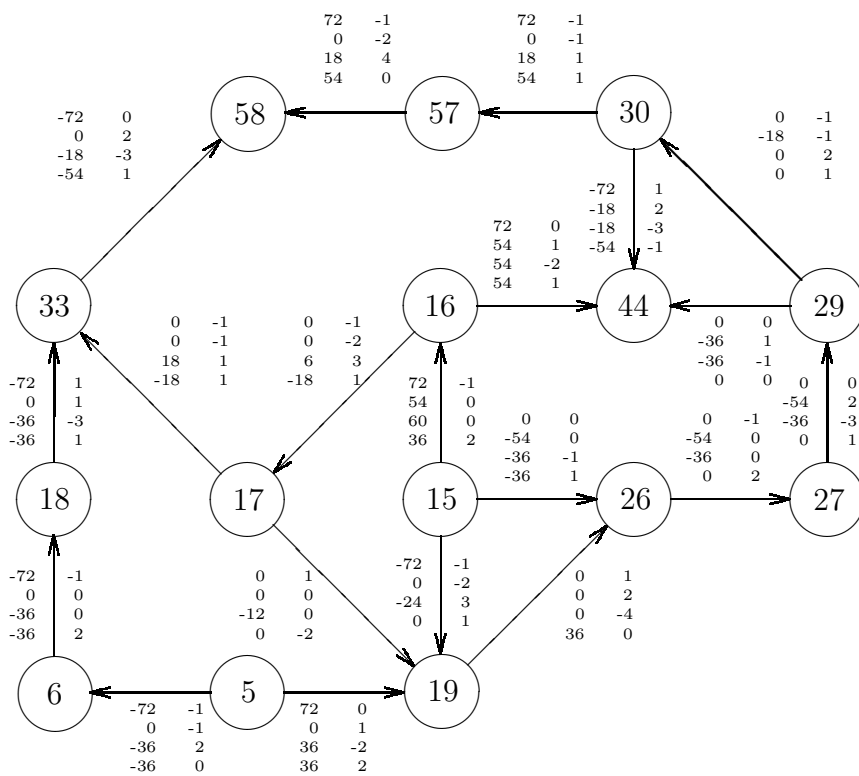


Figure 5.1: The certificate subgraph used in the proof of Theorem 5.0.1.

5	(10, 2, 5, 3)	18	(4, 3, 5, 4)	30	(15, 1, 3, 11)			
6	(14, 4, 11, 5)	19	(5, 1, 2, 2)	33	(3, 1, 2, 3)			
15	(7, 6, 5, 3)	26	(7, 1, 2, 3)	44	(7, 5, 4, 4)			
16	(7, 11, 8, 4)	27	(17, 1, 4, 9)	57	(7, 1, 2, 7)			
17	(5, 2, 3, 3)	29	(10, 1, 2, 6)	58	(7, 1, 3, 8)			
5	6	(3,1,2,1)	16	44	(5,7,5,3)	29	30	(8,1,2,5)
5	19	(8,4,5,3)	17	19	(4,1,2,2)	29	44	(9,3,3,5)
6	18	(2,1,2,1)	17	33	(6,1,3,4)	30	44	(6,5,4,4)
15	16	(6,8,6,3)	18	33	(4,1,3,4)	30	57	(13,1,3,11)
15	19	(5,3,3,2)	19	26	(10,1,3,4)	33	58	(6,1,3,7)
15	26	(9,2,3,3)	26	27	(18,2,5,9)	57	58	(10,1,3,11)
16	17	(8,15,11,5)	27	29	(13,1,3,7)			

Figure 5.2: Representative weight vectors for cones in the certificate.

In particular, the sum of the equations in (5.3) for $(r, t) = (1, 1), (2, 2), (3, 3), (4, 4)$ is zero. Therefore,

$$0 = \sum_{r=1}^4 \sum_{(i,j) \in E'} s_{(i,j)} d_{(i,j),r} f_{(i,j)}^r = \sum_{(i,j) \in E'} s_{(i,j)} \sum_{r=1}^4 d_{(i,j),r} f_{(i,j)}^r.$$

The local contribution at each edge except the edge (29,30) is zero because $d_{(i,j)} \cdot (f_{(i,j)}^1, f_{(i,j)}^2, f_{(i,j)}^3, f_{(i,j)}^4)^T = 0$ (check this in the picture). Consequently,

$$0 = s_{(29,30)} d_{(29,30)} \cdot f_{(29,30)} = 18s_{(29,30)}$$

implying $s_{(29,30)} = 0$. Hence the vertices 29 and 30 have the same coordinates which contradicts that P is a polyhedron with the required edge graph. \square

Remark 5.1.1 Another way to argue is by observing that we have applied the trivial direction of Farkas' lemma to (5.3). With A' being the 16×20 matrix representing the equalities in (5.3) a variant of Farkas' lemma says:

$$\exists y : y^T A' \geq 0 \text{ and } y^T A' \neq 0 \iff \nexists s > 0 : A' s = 0.$$

In our case $y = (1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1)^T$ where the four nonzero components correspond to the equations (1, 1), (2, 2), (3, 3) and (4, 4).

5.1.3 Correctness of the subgraph

For completeness, a positive interior point in each of the 15 maximal cones of the restricted Gröbner fan leading to the inconsistency is given in the top part of Figure 5.2. Further, a positive vector in the relative interior of every shared facet is given in the bottom part.

To verify the correctness of the certificate the following procedure is suggested: It is straightforward to check that the flows are flows and that the dot products of flows and listed directions are 0 except for the edge (29,30). The

question is how to check the correctness of the edge subgraph and the listed directions. For each of the listed edges (i, j) with $i < j$ compute the corresponding reduced Gröbner bases \mathcal{G}_i and \mathcal{G}_j and use Theorem 3.1.9 to compute their cones C_i and C_j . Check that the listed facet vector for the edge (i, j) is in the closure of both cones C_i and C_j and that the listed direction vector non-strictly separates C_i and C_j with C_j being on the non-negative side. Checking that the listed facet vector is in the relative interior of a facet of C_r completes the verification. The non-straightforward part of this test was implemented as a 230 line script in Singular [28]. The script itself is available on the internet; see [35].

5.2 Further remarks

5.2.1 Homogenizing the ideal

In [44] a complete fan in \mathbb{R}^n called the *extended* Gröbner fan is defined for any (not necessarily homogeneous) ideal $I \subseteq R$. We remind the reader that the homogenization of I is defined as $I^h := \langle f^h \rangle \subseteq k[x_0, \dots, x_n]$, where $f^h := x_0^{\deg(f)} f(\frac{x_1}{x_0}, \dots, \frac{x_n}{x_0})$ with $\deg(f)$ being the maximal degree of a term in f . Here “degree” refers to total degree, but we could also consider homogenizations with respect to ω -degree for any vector $\omega \in \mathbb{N}^n$. The *extended* Gröbner fan of I is defined as the Gröbner fan of I^h intersected with \mathbb{R}^n , where \mathbb{R}^n is embedded into \mathbb{R}^{n+1} as $\{0\} \times \mathbb{R}^n$. Every cone in the Gröbner fan is a union of cones in the extended Gröbner fan; see Proposition 5.2.3. It is clear that the extended Gröbner fan is regular since the Gröbner fan of the homogenized ideal is regular — the normal fan of the projection of the state polytope of I^h to $\{0\} \times \mathbb{R}^n$ is the extended Gröbner fan of I . Therefore our example shows that the restricted Gröbner fan of an ideal and its extended Gröbner fan need not agree in $\mathbb{R}_{>0}^n$.

In our example the procedure works as follows. We homogenize the ideal I using the variable “e” to get

$${}^h I = \langle cd^2 + ace, -c^2e + c^2d + abd, c^2e + c^3 - bce - bcd - abd + abc, -ce^2 + ad^2, \\ -c^2e + acd - abe, c^2e - bce + ac^2 - abd, c^2e + a^2c, bce + a^2b \rangle.$$

The Gröbner fan of the new ideal is a complete fan in \mathbb{R}^5 . Intersecting this fan with $\mathbb{R}^4 \times \{0\}$ we get the extended Gröbner fan, a regular fan that refines the Gröbner fan of I in the positive orthant. The extended Gröbner fan has 479 full-dimensional cones. The 81 full-dimensional Gröbner cones of I are covered by 353 such cones. Only 156 (restricted) cones are needed to cover the 81 cones in the restricted Gröbner fan of I . Considering only the non-negative orthant, 62 restricted Gröbner cones are preserved, 11 cones are subdivided into two while the remaining 8 cones are subdivided into 3, 3, 5, 5, 6, 8, 10 and 32 cones, respectively, when we pass to the extended fan. Exactly how the cones are subdivided is shown on the webpage [35]. The subgraph listed in Figure 5.1 is valid for the extended fan except that the cone at vertex 57 is divided into two.

In general, to see that the extended Gröbner fan is in fact a refinement of the Gröbner fan (where defined) we use Proposition 5.2.3 below. Recall [14, Theorem 4, page 397]:

Proposition 5.2.1 *Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal, $\omega = (1, \dots, 1) \in \mathbb{N}^n$ a vector and \prec a term order. Then $\{g^h : g \in \mathcal{G}_{\prec_\omega}(I)\}$ generates I^h .*

Remark 5.2.2 In fact the generating set of I^h is a Gröbner basis. Moreover, the proposition generalizes to an algorithm for computing I^h with respect to any grading $\omega \in \mathbb{Z}_{>0}^n$ on $k[x_1, \dots, x_n]$.

Proposition 5.2.3 below allows us to compute initial ideals with respect to arbitrary weight vectors even if the ideal I is not homogeneous. In a sense this generalizes Corollary 3.1.13.

Proposition 5.2.3 *Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal, $\omega \in \mathbb{R}^n$ a vector and \prec a term order on $k[x_1, \dots, x_n]$. Let \prec' be the term order on $k[x_0, \dots, x_n]$ defined by $x^u \prec' x^v \Leftrightarrow$*

$$\deg_{\text{total}}(x^u) < \deg_{\text{total}}(x^v) \vee (\deg_{\text{total}}(x^u) = \deg_{\text{total}}(x^v) \wedge x_{|x_0=1}^u \prec x_{|x_0=1}^v).$$

The set $\mathcal{G} := \{\text{in}_\omega(g|_{x_0=1}) : g \in \mathcal{G}_{\prec'_{(0,\omega)}}(I^h)\}$ is a Gröbner basis for $\text{in}_\omega(I)$ with respect to \prec .

Proof. It is straight forward to prove the containment $\mathcal{G} \subseteq \text{in}_\omega(I)$. It remains to be proved that $\text{in}_\prec(\text{in}_\omega(I)) \subseteq \langle \text{in}_\prec(g) : g \in \mathcal{G} \rangle$. The left hand side is generated by elements of the form $m = \text{in}_\prec(\sum_i \text{in}_\omega(f_i))$ where $f_i \in I$. We will show that any such m is on the right hand side. Without loss of generality we may assume that the f_i 's have the same ω -degree as m . Hence $m = \text{in}_\prec(\text{in}_\omega \sum_i f_i)$. Let $f = \sum_i f_i \in I$. Then $f^h \in I^h$ and the initial term $\text{in}_{\prec'_{(0,\omega)}}(f^h)$ must be divisible by the initial term $\text{in}_{\prec'_{(0,\omega)}}(g)$ of some Gröbner basis element $g \in \mathcal{G}_{\prec'_{(0,\omega)}}(I^h)$. Consequently, $(\text{in}_{\prec'_{(0,\omega)}}(g))|_{x_0=1}$ divides $(\text{in}_{\prec'_{(0,\omega)}}(f^h))|_{x_0=1}$. Observe that $(\text{in}_{\prec'_{(0,\omega)}}(g))|_{x_0=1} = \text{in}_{\prec_\omega}(g|_{x_0=1}) = \text{in}_\prec(\text{in}_\omega(g|_{x_0=1}))$ since g is homogeneous. Similarly, $(\text{in}_{\prec'_{(0,\omega)}}(f^h))|_{x_0=1} = \text{in}_{\prec_\omega}(f^h|_{x_0=1}) = \text{in}_\prec(\text{in}_\omega(f))$. This proves that $\text{in}_\prec(\text{in}_\omega(g|_{x_0=1}))$ divides $m = \text{in}_\prec(\text{in}_\omega(f))$ as desired. \square

If two generic vectors $(0, \omega)$ and $(0, \omega')$ pick out the same monomial initial ideal of I^h and thereby also the same marked Gröbner basis of I^h , then they will give the same set \mathcal{G} in the Proposition and $\text{in}_\omega(I)$ and $\text{in}_{\omega'}(I)$ must be equal. This proves that the extended Gröbner fan is a refinement of the Gröbner fan (where defined).

5.2.2 A program for finding the example

A C++ program was written for finding non-regular Gröbner fans. The input for the program is a set of generators for an ideal I and the output is either a coordinatization of a polyhedron with the restricted Gröbner fan as its normal fan or a certificate for its non-existence. The program works in two steps.

- In step 1 it calls the software package Gfan [33] being developed by the author; see Section 9.1. Using exhaustive search and the algorithms in Chapter 4, Gfan computes the full-dimensional cones and the codimension 1 cones of the Gröbner fan of I .

- From the Gröbner fan computed above the inequality system (5.2) is deduced. Linear programming methods are used for checking its feasibility. The result is either positive scalars leading to a coordinatization of the vertices of the polyhedron or a certificate for its non-existence.

The software libraries [27] and [20] were used for doing the arithmetic and solving linear programming problems, respectively.

Knowing that we should avoid homogeneous, zero-dimensional and principal ideals, it was not hard to find the example when the C++ program had been written. A practical issue is that we are restricted to ideals with not too complex Gröbner fans as the entire edge graph must be handled by the LP-code. In looking for a 3-variable example this seems to be an unfortunate restriction as nothing interesting happens in the small manageable examples we have tried. Thus it remains an open problem if a smaller example exists or if the ideal can be replaced by a prime ideal.

Chapter 6

Tropical varieties

Before we learn to add and multiply numbers we learn to compare them i.e. take their maximum. Maximum and plus are the two basic operations in tropical mathematics. Since we are usually not thinking about the max-plus operations in abstract terms we may be surprised to realize that they satisfy the distributive law. The fact that many combinatorial optimization problems can be phrased in tropical language is reason enough for a closer study of the mathematical structure. It is the purpose of the second half of this thesis to study the relation between tropical mathematics, algebraic varieties and Gröbner fans.

Tropical mathematics, being a fundamental topic, has been reinvented many times. It was given the adjective “tropical” to honor the Brazilian mathematician Imre Simon. Recently tropical mathematics has become a very active field of research. We start by giving a very general introduction to tropical mathematics. In Section 6.1 which is about tropical varieties we will leave the max-plus way of thinking and state definitions and theorems in terms of initial ideals and Gröbner cones.

In the *tropical semi-ring* $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)$ the tropical sum $a \oplus b$ of two numbers is their maximum and the tropical product $a \odot b$ is the usual sum $a + b$. For example $(-4) \oplus (-7) \odot 4 = -3$. We notice that \oplus and \odot are associative and commutative compositions. Furthermore, the distributive law holds:

$$a \odot (b \oplus c) = a + \max(b, c) = \max(a + b, a + c) = a \odot b \oplus a \odot c$$

The additive neutral element is $-\infty$. Since $-\infty$ is the only element having an additive inverse the word “semi-ring” is suitable.

A tropical polynomial of degree n in one variable is of the following form

$$F = c_0 \oplus c_1 \odot x \oplus c_2 \odot x \odot x \oplus \cdots \oplus c_n \odot x \odot \cdots \odot x$$

where $c_n \neq -\infty$. If some coefficient c_i equals the additive neutral element $-\infty$ we may consider its term as being non-present. A tropical polynomial F defines the tropical polynomial function $F(x)$. As the following example shows this function is piecewise linear.

Example 6.0.4 Let $F = (-1) \odot x \odot x \oplus 1 \odot x \oplus 2$. Figure 6.1 shows the graph of $F(x) = \max(2x - 1, x + 1, 2)$. Notice that F factors as $F = (-1) \odot (0 \odot x \oplus 1) \odot (0 \odot x \oplus 2)$.

We wish to define the “zero set” of a tropical polynomial. A tropical polynomial function $F(x)$ only attains the additive neutral value $-\infty$ if either

- $x = -\infty$ and the constant coefficient in F is $-\infty$, or
- F contains no terms.

For this reason, defining the zero set of F to be the set of points where $F(x) = -\infty$ would be uninteresting. Instead we choose the following definition:

Definition 6.0.5 Let F be a tropical polynomial. A point $x \in \mathbb{R}$ is a zero of F if the value $F(x)$ is attained by at least two terms of F when evaluating F tropically in x . The *zero set* $\mathbb{T}(F)$ is the set of all zeros of F .

In other words $\mathbb{T}(F)$ is the set of points where $F(x)$ is not differentiable. In Example 6.0.4, $\mathbb{T}(F) = \{1, 2\}$. The definition is reasonable since $\mathbb{T}(F_1 \odot F_2) = \mathbb{T}(F_1) \cup \mathbb{T}(F_2)$. Notice that from each tropical linear term in the factorization of F in Example 6.0.4 we can read off one zero. The generalization to several variables is straight forward. Formally the tropical polynomial (semi-)ring in n variables is the group (semi-)ring $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)[\mathbb{N}^n]$ over the monoid \mathbb{N}^n .

Example 6.0.6 Let $F = 0 \oplus (-2) \odot x \oplus (-1) \odot y \oplus (-4) \odot x \odot y \oplus (-6) \odot x \odot x \oplus (-4) \odot y \odot y$. Evaluating the polynomial gives a piece-wise linear function. The zero set $\mathbb{T}(F)$ is a *tropical curve* which divides \mathbb{R}^2 into six regions, one for each term. In each region the maximum is attained by just one term. See Figure 6.2. Depending on the coefficients one or more of these regions may be missing.

We mention a few properties of zero-sets in the plane. A polynomial of degree 1 with three terms defines a tropical line which is the union of three half-lines meeting in a point. The coordinates of the point depends on the coefficients of the defining polynomial. The half-lines go off in direction south, west and north-east. An interesting property of the plane is that any two generic points are contained in a unique line and any two generic lines intersect in a single point. This also holds for tropical lines, see Figure 6.3.




Figure 6.1: The graph of $F(x) = \max(2, x + 1, 2x - 1)$ in Example 6.0.4

The tropical semi-ring can be generalized to polytopes in \mathbb{R}^n . The tropical sum $P \oplus Q$ of two polytopes will be the convex hull of their union and the tropical product $P \odot Q$ will be their usual Minkowski sum $P + Q$. Observe that the tropical product of the Newton polytopes of two polynomials is the Newton polytope of the product of the polynomials. Many known problems and algorithms can be stated in tropical language.

- Evaluating a tropical polynomial with all coefficients 0 is the same as maximizing a linear form over its Newton polytope.
- The maximal matching in a bipartite graph with $2 \times n$ vertices equals the tropical permanent of the $n \times n$ matrix of edge weights.
- The negated $n \times n$ matrix of shortest distances in a directed graph with n vertices can be computed tropically as A^n where A is the $n \times n$ matrix of negated edge lengths.

In tropical mathematics definitions are made in such a way that the classical results have tropical analogs. The endless list of constructions and theorems in tropical mathematics includes the tropical projective plane [48], a tropical Bezout theorem [48], a notion of tropical convexity [15], a group law for tropical elliptic curves [59] etc.. Some of these examples seem to suggest that tropical geometry is a pure abstract piece of mathematics with no applications. This is not the case. Tropical mathematics is closely related to Bernstein's Theorem which bounds the number of solutions in the algebraic torus of a polynomial system of n polynomial equations in n variables based on the geometry of the Newton polytopes of the polynomials. Related to this is the polyhedral numerical homotopy for finding the solutions. In [30] progress was made on Smale's sixth problem by showing that the number of relative equilibria of four bodies satisfying Newton's laws of gravity is finite. A key step here was a computation of the mixed faces of a Minkowski sum and an investigation of the corresponding initial ideals. Other applications of tropical mathematics are in algebraic statistics where polytope propagation has been used as a method for parametric inference for hidden Markov models in computational biology; see [47]. Applications of tropical ideas to mathematics itself have so far mainly been in complex and real enumerative geometry; see [24, Section 3].




Figure 6.2: The zero set of the quadratic curve in Example 6.0.6 which is equal to the tropical variety in Example 6.3.4.




Figure 6.3: The intersection point of two tropical lines. The tropical line containing two specified points.

Our primary interest is in tropical varieties, which we will define soon, and in structure that will lead to algorithms for their computation. Although attempts are made to put tropical mathematics in an abstract setting we shall be very concrete and study tropicalizations of algebraic varieties with an embedding. This will work best for our computational purposes. One way to make things more abstract is by for example considering tropical curves as abstract balanced graphs without an embedding, see [24]. More ambitious attempts can be found in Mikhalkin's book [43] which is work in progress. Here a definition of for example a tropical schemes can be found.

Remark 6.0.7 We have chosen to use maximum as the tropical addition. An equally good choice would have been minimum. The choice of maximum is compatible with the choice of maximizing the degree when taking initial forms and considering outer normal cones rather than inner. However, in Section 6.1 our choice forces us to introduce a minus sign when taking the valuation of a variety.

6.1 Tropical varieties

Tropical varieties are images of varieties in $(\mathbb{C}\{\{t\}\}^*)^n$ under the valuation map. We will explain this in the following.

Definition 6.1.1 The *Puiseux series field* $\mathbb{C}\{\{t\}\}$ is the set of (possibly finite) formal series of the form

$$p = c_1 t^{v_1} + c_2 t^{v_2} + \dots,$$

where $c_i \in \mathbb{C}^*$ and $v_i \in \mathbb{Q}$ are increasing with a common denominator $N \in \mathbb{N} \setminus \{0\}$. The field comes with a valuation:

$$\text{val} : \mathbb{C}\{\{t\}\} \rightarrow \mathbb{Q} \cup \{\infty\}$$

taking a series to the exponent of the first term. For the series above $\text{val}(p) = v_1$. We define $\text{val}(0) = \infty$.

An important property of $\mathbb{C}\{\{t\}\}$ is that it is algebraically closed. The proof of this is constructive, namely, it follows from the Newton-Puiseux algorithm which produces all solutions, up to any desired degree, of a polynomial equation $f = 0$ with $f \in \mathbb{C}\{\{t\}\}[x] \setminus \{0\}$. The Newton-Puiseux algorithm works by




Figure 6.4: The Newton polygons of f and f' in Example 6.1.2. A normal vector of the lower edge in the left picture is $(-1, \frac{1}{2})$, where the first coordinate is the t -component and the second coordinate is the x component.

considering the Newton polygon of f (the convex hull of exponent vectors of f) which may be unbounded in \mathbb{R}^2 . Letting the t -axis be vertical, the negative slopes of the lower-edges of the Newton polygon of f give the various possibilities for the first exponent v_1 in a Puiseux series solution. The coefficient is found by finding roots of a certain initial form in $\mathbb{C}[x]$. This gives the possible first terms of a solution. Repeating the process we get the succeeding terms. The process is best illustrated by an example. We will not explain the details but just give the flavor of the algorithm. It is the close connection to Newton polygons that ties Puiseux series solutions and initial ideals together.

Example 6.1.2 To find the two roots of $f = (1+t) + t^2x + tx^2$ we draw its Newton polygon; see Figure 6.4 (left). We are searching for a solution of the form $x = c_1t^{v_1} + x'$ where x' consists of higher order terms. We define $b = c_1t^{v_1}$. Substituting $x = b + x'$ into f we get $f = (1+t) + t^2(b+x') + t(b^2 + 2bx' + x'^2)$. Using that the degree of b is lower than the degree of x' we see that the lowest degree terms in f (before canceling out terms) are among 1 , t^2b and tb^2 . These terms appear as points in our picture. If x is a root the lowest degree terms must cancel. This can only happen if the exponent v_1 is $-\frac{1}{2}$. The value $v_1 = -\frac{1}{2}$ can be read off from our picture by observing that $(-1, -v_1)$ is a normal to the lower edge in our picture. The choice of v_1 leads to the constraint $0 = 1 + t(c_1(t^{-\frac{1}{2}}))^2$ or, equivalently, $0 = 1 + c_1^2$ on the coefficient c_1 . We have the choice of letting c_1 be $+i$ or $-i$. Let us choose $c_1 = i$. Now x must have the form $x = it^{-\frac{1}{2}} + x'$. Substituting this into f we get a polynomial $f' = (t + it^{\frac{3}{2}}) + (2it^{\frac{1}{2}} + t^2)x' + tx'^2$. We may now compute the next term of x , the first term of x' , by solving $f' = 0$. The Newton polygon of f' (Figure 6.4 (right)) shows that there is only one choice $v_2 = \frac{1}{2}$ for the next exponent. (The choice $-\frac{1}{2}$ is not valid since $-\frac{1}{2} \not\prec v_1$.) Continuing this way we get the Puiseux series $x = it^{-\frac{1}{2}} + \frac{1}{2}it^{\frac{1}{2}} - \frac{1}{2}t + \dots$

Remark 6.1.3 The difficult part of the proof of the Newton-Puiseux algorithm is to show that the exponents in the produced series have a common denominator. It is important that \mathbb{C} has characteristic 0. The Puiseux series field over an algebraically closed field is not algebraically closed in general.

If a polynomial in $\mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ happens to be in $\mathbb{C}[x_1, \dots, x_n]$ we say that the polynomial has *constant coefficients*.

Let $V \subseteq \mathbb{C}\{\{t\}\}^n$ be a variety defined by an ideal $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$. We let val denote the coordinatewise valuation $\text{val} : \mathbb{C}\{\{t\}\}^n \rightarrow (\mathbb{Q} \cup \{\infty\})^n$. We may consider the usual topological closure of $-\text{val}(V) \cap \mathbb{Q}^n$ in \mathbb{R}^n . This set is the *tropical variety* of I or the *tropicalization* of V . This is the natural way to define tropical varieties from a Puiseux series view-point. However, for the purpose of this thesis it is much more natural to define the tropical varieties in terms of initial ideals and Gröbner fans. A definition in terms of initial ideals was first presented in [51] and [54]. Our definition requires a few steps and goes as follows.

Definition 6.1.4 For $\omega \in \mathbb{R}^n$ the t - ω -degree of a term ct^ax^v with $c \in \mathbb{C}^*$, $a \in \mathbb{Q}$ and $v \in \mathbb{Z}^n$ is defined as $-\text{val}(ct^a) + \omega \cdot v = -a + \omega \cdot v$. The t -initial form $\text{t-in}_\omega(f) \in \mathbb{C}[x_1, \dots, x_n]$ of a polynomial $f \in \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ is the sum of all terms in f of maximal t - ω -weight but with 1 substituted for t .

Remark 6.1.5 Notice that since t has t - ω -degree -1 , the maximal t - ω -weight is attained by a term if the polynomial is non-zero. Furthermore, only a finite number of terms attain the maximum. Therefore, it makes sense to substitute $t = 1$ and consider the finite sum of terms as a polynomial in $\mathbb{C}[x_1, \dots, x_n]$.

Example 6.1.6 Consider $f = (1 + t) + t^2x + tx^2 \in \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ from Example 6.1.2. Let $\omega = (\frac{1}{2}) \in \mathbb{R}^1$. Then $\text{t-in}_\omega(f) = 1 + x^2$. This happens to be the polynomial defining the coefficient $c_1 \in \mathbb{C}$ in the example. For any other choice of ω the t -initial form is a monomial.

Definition 6.1.7 Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ and $\omega \in \mathbb{R}^n$. The t -initial ideal of I with respect to ω is defined as:

$$\text{t-in}_\omega(I) := \langle \text{t-in}_\omega(f) : f \in I \rangle \subseteq \mathbb{C}[x_1, \dots, x_n].$$

Remark 6.1.8 Sometimes we will take t -initial ideals of ideals that are not in ideals in $\mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ without further definition. In these cases it should be clear from the context what is meant.

Definition 6.1.9 Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be an ideal. The *tropical variety* of I is the set

$$\mathcal{T}(I) := \{\omega \in \mathbb{R}^n : \text{t-in}_\omega(I) \text{ is monomial-free}\}.$$

Here monomial-free means that the ideal does not contain a monomial.

Lemma 6.1.10 Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be an ideal and $\omega \in \mathbb{R}^n$. Then $\omega \in \mathcal{T}(I)$ if and only if for all $f \in I$, $\text{t-in}_\omega(f)$ is not a monomial.

Proof. The “only if” direction is clear. For the other implication, suppose that $\omega \notin \mathcal{T}(I)$. Then there exists a monomial in $\text{t-in}_\omega(I)$. After a few rewritings we may assume that the monomial has the form $\sum_i \text{t-in}_\omega(f_i)$ where $f_i \in I$. Multiplying the f_i 's by suitable t -powers, they all get the same maximal t - ω -degree. Taking the sum of the new polynomials the monomials with this maximal t - ω -degree will cancel to give a single monomial. This proves that I has a polynomial f with $\text{t-in}_\omega(f)$ being a monomial. \square

Lemma 6.1.11 *For a principal ideal $\langle f \rangle \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ we have*

$$\mathcal{T}(\langle f \rangle) = \{\omega \in \mathbb{R}^n : \text{t-in}_\omega(f) \text{ is not a monomial}\}.$$

Proof. The inclusion “ \subseteq ” is clear. For the other inclusion, suppose that $\text{t-in}_\omega(f)$ is not a monomial. According to Lemma 6.1.10 we need to check that $\text{t-in}_\omega(h)$ is not a monomial for all $h \in \langle f \rangle$. Such an h is of the form gf with $g \in \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$. Observe that $\text{t-in}_\omega(gf) = \text{t-in}_\omega(g)\text{t-in}_\omega(f)$. This cannot be a monomial since $\text{t-in}_\omega(f)$ has more than one term. \square

Example 6.1.12 In Example 6.1.2 the tropical variety of $\langle f \rangle$ is $\{\frac{1}{2}\}$ while $\mathcal{T}(\langle f' \rangle) = \{-\frac{1}{2}, \frac{1}{2}\}$.

The first part of this thesis has been about usual initial ideals where the valuation is not taken into account when taking initial forms. The following lemma justifies Definition 6.1.14 as an equivalent definition for ideals generated by constant-coefficient polynomials. The proof will be given later in this section.

Lemma 6.1.13 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be an ideal and $J = \langle I \rangle_{\mathbb{C}\{\{t\}\}[x_1, \dots, x_n]} \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be the ideal it generates. Then $\text{in}_\omega(I) = \text{t-in}_\omega(J)$ for all $\omega \in \mathbb{R}^n$.*

Definition 6.1.14 Let $I \subseteq \mathbb{Q}[x_1, \dots, x_n]$ (or $I \subseteq \mathbb{C}[x_1, \dots, x_n]$). The *tropical variety* of I is the set

$$\mathcal{T}(I) := \{\omega \in \mathbb{R}^n : \text{in}_\omega(I) \text{ is monomial-free}\}.$$

We could have stated the above definition for any field k . However, if $k = \mathbb{C}\{\{t\}\}$ the definition would not be consistent with our previous definition. For this reason we stick to \mathbb{Q} and \mathbb{C} .

In Section 6.3 we will explain why it for most purposes suffices only to consider tropical varieties defined by Definition 6.1.14. Basically, the other tropical varieties (Definition 6.1.9) are gotten by intersecting with a hyperplane.

The first computational question for tropical varieties we will address is how to compute t-initial ideals. We restrict ourselves to ideals generated by elements of $\mathbb{C}[t, x_1, \dots, x_n]$ since it is here Gröbner bases work best. Lemma 6.1.13 will follow as an easy corollary.

Proposition 6.1.15 [42, Proposition 7.3] *Let $I \subseteq \mathbb{C}[t, x_1, \dots, x_n]$ be an ideal, $J = \langle I \rangle_{\mathbb{C}\{\{t\}\}[x_1, \dots, x_n]}$ and $\omega \in \mathbb{R}^n$. Then $\text{t-in}_\omega(I) = \text{t-in}_\omega(J)$.*

Proof. We need to prove the inclusion $\text{t-in}_\omega(I) \supseteq \text{t-in}_\omega(J)$. The other inclusion is clear since $I \subseteq J$. The right hand side is generated by elements of the form $f = \text{t-in}_\omega(g)$ where $g \in J$. Consider such f and g . The polynomial g must be of the form $g = \sum_i c_i \cdot g_i$ where $g_i \in I$ and $c_i \in \mathbb{C}\{\{t\}\}$. Let d be the $(-1, \omega)$ -degree of $\text{in}_\omega(g)$. The degrees of terms in g_i are bounded. Terms $\alpha \cdot t^\beta$ in c_i of large enough t -degree will make the $(-1, \omega)$ -degree of $\alpha \cdot t^\beta \cdot g_i$ drop below d since the degree of t is negative. Consequently, these terms can simply be ignored since they cannot affect the initial form of $g = \sum_i c_i \cdot g_i$. Without loss

of generality we may assume that these terms have already been removed from g . Renaming and possibly repeating some g_i 's we may write g as a finite sum $g = \sum_i c'_i \cdot g_i$ where $c'_i = \alpha_i \cdot t^{\beta_i}$ and $g_i \in I$ with $\alpha_i \in \mathbb{C}$ and $\beta_i \in \mathbb{Q}$. We will split the sum into subsums grouping together the c'_i 's that have the same t -exponent modulo \mathbb{Z} . For suitable index sets A_j we let $g = \sum_j G_j$ where $G_j = \sum_{i \in A_j} c'_i \cdot g_i$. Notice that all t -exponents in a G_j are congruent modulo \mathbb{Z} while t -exponents from different G_j 's are not. In particular there is no cancellation in the sum $g = \sum_j G_j$. As a consequence $\text{in}_\omega(g) = \sum_{j \in S} \text{in}_\omega(G_j)$ for a suitable subset S . We also have $\text{t-in}_\omega(g) = \sum_{j \in S} \text{t-in}_\omega(G_j)$. We wish to show that each $\text{t-in}_\omega(G_j)$ is in $\text{t-in}(I)$. We can write $t^{\gamma_j} \cdot G_j = \sum_{i \in A_j} t^{\gamma_j} \cdot c'_i \cdot g_i$ for suitable $\gamma_j \in \mathbb{Q}$ such that $t^{\gamma_j} \cdot c' \in \mathbb{C}[t]$ for all $i \in A_j$. Observe that

$$\text{t-in}_\omega(G_j) = \text{t-in}_\omega(t^{\gamma_j} \cdot G_j) = \text{t-in}_\omega\left(\sum_{i \in A_j} t^{\gamma_j} \cdot c'_i \cdot g_i\right) \in \text{t-in}_\omega(I).$$

Applying $\text{t-in}_\omega(g) = \sum_{j \in S} \text{t-in}_\omega(G_j)$ we see that $f = \text{t-in}_\omega(g) \in \text{t-in}_\omega(I)$. \square

Remark 6.1.16 For $f \in \mathbb{C}[t, x_1, \dots, x_n]$ we have $\text{t-in}_\omega(f) = (\text{in}_{(-1, \omega)}(f))|_{t=1}$. Consequently, for $I \subseteq \mathbb{C}[t, x_1, \dots, x_n]$ we have $\text{t-in}_\omega(I) = (\text{in}_{(-1, \omega)}(I))|_{t=1}$. In order to compute $\text{t-in}_\omega(I)$ we may simply compute the initial ideal $\text{in}_{(-1, \omega)}(I)$ using Corollary 3.1.13 or Proposition 5.2.3 and substitute.

Proof of Lemma 6.1.13. Consider $I' = \langle I \rangle_{\mathbb{C}[t, x_1, \dots, x_n]} \subseteq \mathbb{C}[t, x_1, \dots, x_n]$. Proposition 6.1.15 shows that $\text{t-in}_\omega(J) = \text{t-in}_\omega(I')$ and according to Remark 6.1.16 $\text{t-in}_\omega(I') = \text{in}_{(-1, \omega)}(I')|_{t=1}$. Since the reduced Gröbner bases of I and I' are the same, not involving t , we get $\text{in}_{(-1, \omega)}(I')|_{t=1} = \text{in}_\omega(I)$. \square

We finish this section by stating the theorem referred to in the beginning. We consider this the most important theorem in tropical geometry since it shows that the initial ideal definition and the alternative valuation definition of tropical varieties are equivalent.

Theorem 6.1.17 *Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be an ideal. Then*

$$\mathbb{Q}^n \cap \mathcal{T}(I) = -\text{val}(V(I) \cap (\mathbb{C}\{\{t\}\}^*)^n).$$

We provide a proof for the trivial “ \supseteq ” inclusion and for the special case where the ideal is zero-dimensional; see [42, Theorem 3.1]. The second half of the proof uses techniques from Section 6.4 below.

Proof. Let $p \in V(I) \cap (\mathbb{C}\{\{t\}\}^*)^n$. According to Lemma 6.1.10 it suffices to show that for all $f \in I$ we have $-\text{val}(p) \in \mathcal{T}(\langle f \rangle)$. Let $\omega = -\text{val}(p)$. When we insert p into f we get 0. This means that all terms in f cancel out. In particular the terms of lowest t -degree in the sum $f(p)$ before cancellation must cancel. These arise from the lowest t - ω -degree terms of f which happen to be the terms in $\text{t-in}_\omega(f)$ before substituting $t = 1$. Consequently $\text{t-in}_\omega(f)$ cannot be a monomial. This proves that $\omega = -\text{val}(p) \in \mathcal{T}(I)$.

Suppose I is a zero-dimensional ideal and consider a minimal primary decomposition $I = \bigcap_i Q_i$. According to Proposition 6.4.1 and Proposition 6.4.2

$\mathcal{T}(I) = \mathcal{T}(\bigcap_i Q_i) = \bigcup_i \mathcal{T}(Q_i) = \bigcup_i \mathcal{T}(\sqrt{Q_i})$. However, the prime ideal $\sqrt{Q_i}$ of dimension 0 is maximal and must have form $\langle x_1 - p_1, \dots, x_n - p_n \rangle$ for some $p_j \in \mathbb{C}\{\{t\}\}$ since $\mathbb{C}\{\{t\}\}$ is algebraically closed. If $\omega \in \mathcal{T}(I)$ then $\omega \in \mathcal{T}(\sqrt{Q_i})$ for some i . This ω must pick out both terms in every binomial $x_j - p_j$. Hence $\omega = -\text{val}(p)$. \square

As mentioned, the initial ideal definition of tropical varieties is due to Speyer and Sturmfels and they present a proof of Theorem 6.1.17 in [51] which seems to have a gap however. Other proofs can be found in [39, Theorem 5.2.2] and [16, Theorem 4.2]. In [42] a constructive proof is presented which will take a point in the tropical variety and lift it to a point in $(\mathbb{C}\{\{t\}\})^*$ with the right valuation.

In view of the original Puiseux valuation definition of tropical varieties it seems more natural to define tropical varieties for ideals in the Laurent polynomial ring $\mathbb{C}\{\{t\}\}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. However as we will apply Gröbner basis techniques in the following, to us it makes more sense to consider ideals in the polynomial ring rather than the Laurent polynomial ring. For example, in the Laurent polynomial ring the division algorithm does not have to terminate since any two monomials are divisible by each other.

6.2 Examples and the basic structure

In this section we will take a closer look at the structure of tropical varieties. We will restrict ourselves to varieties of ideals in $\mathbb{C}[x_1, \dots, x_n]$ defined using usual initial ideals; see Definition 6.1.14. Such tropical varieties are invariant under scaling by a positive scalar $s \in \mathbb{R}_{>0}$. In fact, if the defining ideal is homogeneous (with respect to some positive grading) the Gröbner fan has full support and the tropical variety can be covered by a finite set of Gröbner cones. Sometimes it is useful to give the tropical variety a polyhedral structure. *In the homogeneous case we also define $\mathcal{T}(I)$ to be the set of Gröbner cones $C_\omega(I)$ for which $\text{in}_\omega(I)$ does not contain a monomial.* This is a closed condition in the sense that if $\text{in}_\omega(J)$ is monomial-free then so is J . Hence, $\mathcal{T}(I)$ is a polyhedral fan. It should lead to no confusion that we use $\mathcal{T}(I)$ to denote the fan as well as the support of the fan.

The tropical variety defined by a principal ideal $f \in \langle f \rangle$ is called a *tropical hypersurface*. Since $\text{in}_\omega(gf) = \text{in}_\omega(g)\text{in}_\omega(f)$ for $f, g \in \mathbb{C}[x_1, \dots, x_n]$ we have $\text{in}_\omega(\langle f \rangle) = \langle \text{in}_\omega(f) \rangle$. This shows that $\mathcal{T}(\langle f \rangle)$ consist of all normal cones of the Newton polytope of f of dimension less than n . For short we denote the tropical hypersurface by $\mathcal{T}(f)$.

Example 6.2.1 Let $I = \langle x + y + z \rangle \subseteq \mathbb{C}[x, y, z]$. The Newton polytope of $x + y + z$ has 3 vertices, 3 edges and one 2-dimensional face. The normal cones of the edges and the 2-dimensional face are three halfplanes and a line. These are the cones in $\mathcal{T}(I)$.

Example 6.2.2 By a *monomial difference* we mean a binomial of the form $x^u - x^v$ with $u, v \in \mathbb{N}^n$. The tropical variety of a homogeneous ideal I generated




Figure 6.5: The tropical variety of the ideal in Example 6.2.3 drawn projectively.

by monomial differences is the homogeneity space of I . To see this observe that any reduced Gröbner basis of I must consist of only binomials. Any initial ideal $\text{in}_v(I)$ can be computed from a reduced Gröbner basis by applying Corollary 3.1.13. However, for any v that does not pick out both terms in all binomials the initial ideal will contain a monomial. Thus any monomial-free initial ideal of I is I itself. In particular, the tropical variety of a toric ideal is just a subspace.

Example 6.2.3 This example has become the standard example of a tropical variety. The 10 2×2 minors of a 2×5 matrix

$$\begin{pmatrix} x_{11} & x_{12} & x_{13} & x_{14} & x_{15} \\ x_{21} & x_{22} & x_{23} & x_{24} & x_{25} \end{pmatrix}$$

are

$$\begin{aligned} a &= x_{11}x_{22} - x_{12}x_{21} \\ b &= \dots \\ &\dots \end{aligned}$$

They satisfy the Grassmann-Plücker relations:

$$0 = bf - ah - ce = bg - ai - de = cg - aj - df = ci - bj - dh = fi - ej - gh.$$

The ideal generated by these relations is called the Grassmann-Plücker ideal $I \subseteq \mathbb{Q}[a, \dots, j]$. The ideal is prime and has dimension 7 and the tropical variety of I is pure of dimension 7. The homogeneity space of I has dimension 5 and all cones in the tropical variety contain this subspace. Besides the homogeneity space the tropical variety consists of 10 6-dimensional cones and 15 7-dimensional cones. The f-vector of the complex is $(1, 10, 15)$. Modulo the homogeneity space the tropical variety is 2-dimensional. Projectively we may draw its combinatorics as the Petersen graph; see Figure 6.5. The homogeneity space is the center of the projection. The Gröbner fan of which the tropical variety is a subfan has f-vector $(1, 20, 120, 300, 330, 132)$.

Example 6.2.4 Let $I \subseteq \mathbb{Q}[a, b, \dots, o]$ be the ideal generated by the 3×3 minors of the matrix

$$\begin{pmatrix} a & b & c & d & e \\ f & g & h & i & j \\ k & l & m & n & o \end{pmatrix}.$$

The f-vectors of the tropical variety and the Gröbner fan are (1, 45, 315, 930, 1260, 630) and (1, 45, 585, 3390, 10710, 19890, 21750, 12960, 3240), respectively. For example there is 1 cone of dimension 7 in both fans, namely the homogeneity space $C_0(I)$. The Krull dimension of $\mathbb{Q}[a, b, \dots, o]/I$ is $d = 12$. The tropical variety is pure. It only has cones up to dimension 12 while the Gröbner fan contains cones all the way up to dimension 15.

If the ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ is not homogeneous we may homogenize it according to the definition in Section 5.2.1. As we saw in that section the Gröbner fan structure may change when we homogenize. However as a subset of \mathbb{R}^n the tropical variety is preserved under homogenization if we restrict the homogenized tropical variety to $\{0\} \times \mathbb{R}^n$:

Lemma 6.2.5 *Let $I = \langle f_1, \dots, f_m \rangle \subseteq k[x_1, \dots, x_n]$ be an ideal. For $\omega \in \mathbb{R}^n$ the initial ideal $\text{in}_\omega(I)$ contains a monomial if and only if $\text{in}_{(0) \times \omega}(\langle f_1^h, \dots, f_m^h \rangle)$ contains a monomial.*

Proof. Suppose $x^u \in \text{in}_\omega(I)$. Then there exists an $f \in I$ such that $\text{in}_\omega(f) = x^u$. For a suitable choice of $a_j \in \mathbb{N}^n$ and indices i_j we have $f = \sum_j x^{a_j} f_{i_j}$. Consider $f' = \sum_j x^{a'_j} f_{i_j}^h$ where $a'_j \in \mathbb{N}^{n+1}$ is $a_j \in \mathbb{N}^n$ with an additional coordinate adjusted so that the same cancellations happen as before. Taking the $(0, \omega)$ -initial form of f' the terms keep their old ω degrees as new $(0, \omega)$ -degrees. For this reason $\text{in}_{(0, \omega)}(f')$ has just one term and $\text{in}_{(0, \omega)}(\langle f_1^h, \dots, f_m^h \rangle)$ contains a monomial.

For the other inclusion, suppose that $x^u \in \text{in}_{(0, \omega)}(\langle f_1^h, \dots, f_m^h \rangle)$ for $u \in \mathbb{N}^{n+1}$. Then $x^u = \text{in}_{(0, \omega)}(f)$ for some $f \in \langle f_1^h, \dots, f_m^h \rangle$. Notice that $f|_{x_0=1} \in \langle f_1, \dots, f_m \rangle_{k[x_1, \dots, x_n]} = I$. Since $\text{in}_\omega(f|_{x_0=1}) = x^u|_{x_0=1}$ we have $x^u|_{x_0=1} \in \text{in}_\omega(I)$. \square

The lemma says that we need not compute the true homogenization of I but may homogenize any generating set. Usually such a homogenization could lead to too many components but as tropical varieties ignore the coordinate hyperplanes of $\mathbb{C}\{\{t\}\}^n$ this is not a problem.

Notice that the above proof will go through even if we homogenize with respect to another grading of $k[x_1, \dots, x_n]$. However, the polyhedral structure of the tropical variety depends on the choice of grading:

Example 6.2.6 Consider $I = \langle ab^3c + c^2 + bc + b^2 + abc^3 + ab^2c^2, a^2bc + c + b + a + abc^2 + ab^2c \rangle \subseteq k[a, b, c]$. The f-vector of $\mathcal{T}(I^h)$ is (1, 8, 6) if we homogenize with respect to the grading (1, 1, 1) but (1, 10, 8) if we choose the grading (2, 1, 1).

Applying a $GL_n(\mathbb{Z})$ linear transformation to a tropical variety we again get a tropical variety. This is easiest seen if we consider tropical varieties defined by ideals in $R = \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. For a matrix $A \in GL_n(\mathbb{Z})$ we define the

\mathbb{C} -algebra isomorphism $\varphi_A : R \rightarrow R$ by sending x^v to x^{Av} . We refer to this substitution as a *multiplicative change of variables*. Below initial ideals are considered as ideals in the Laurent polynomial ring.

Proposition 6.2.7 *Let A be as above. For an ideal $I \subseteq \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ and a vector $\omega \in \mathbb{R}^n$ we have*

$$\text{in}_{A^T\omega}(I) = \varphi_A^{-1}(\text{in}_\omega(\varphi_A(I))).$$

Proof. To prove the proposition observe that for $f \in I$ we have $\text{in}_{A^T\omega}(f) = \varphi_A^{-1}(\text{in}_\omega(\varphi_A(f)))$, since $(A^T\omega)^T v = \omega^T(Av)$, and use that φ is an isomorphism of Laurent polynomial rings. \square

Corollary 6.2.8 *Let A be as above. For an ideal $I \subseteq \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$*

$$\mathcal{T}(I) = A^T(\mathcal{T}(\varphi_A(I)))$$

where the right hand side is a linear transformation of $\mathcal{T}(\varphi_A(I))$.

If a tropical variety is defined by a homogeneous ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ it is also the tropical variety of an ideal $J \subseteq R$ and thus applying the linear transformation to the variety we get a new variety – also defined by some ideal in $\mathbb{C}[x_1, \dots, x_n]$. However, it is not clear how the polyhedral structure induced by the Gröbner fan changes during this procedure.

6.3 Reduction to $\mathbb{Q}[x_1, \dots, x_n]$

Since we are seeking algorithms that can be implemented on computers it is not possible to consider ideals generated by arbitrary elements of $\mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$. The reason is that the coefficient field is uncountable. The uncountability comes partly from \mathbb{C} and partly from the infinite series in t . The field $\mathbb{C}(t)$ of rational functions in t is a subfield of $\mathbb{C}\{\{t\}\}$ with a compatible valuation and the field of algebraic numbers $\overline{\mathbb{Q}} \subseteq \mathbb{C}$ is countable. Restricting ourselves to polynomials generated by elements in $\overline{\mathbb{Q}}(t)[x_1, \dots, x_n]$ we no longer have the problem with uncountability. It will follow from Lemma 6.3.1 and Lemma 6.3.5 that it is no further restriction to have our algorithms work for ideals in $I \subseteq \mathbb{Q}[x_1, \dots, x_n]$ only.

We start by considering an ideal $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ generated by elements of $\mathbb{C}(t)[x_1, \dots, x_n]$. We may clear denominators of each generator by multiplying it with a suitable polynomial in $\mathbb{C}[t]$ which is a unit in $\mathbb{C}(t)$. Hence the ideal is generated by polynomials in $\mathbb{C}[t, x_1, \dots, x_n]$. Let $J = I \cap \mathbb{C}[t, x_1, \dots, x_n]$. According to Proposition 6.1.15 and Remark 6.1.16, $\text{t-in}_\omega(I) = \text{t-in}_\omega(J) = \text{in}_{(-1, \omega)}(J)|_{t=1}$. The ideal $\text{in}_{(-1, \omega)}(J)$ is homogeneous in the $(-1, \omega)$ -grading, implying that $\text{in}_{(-1, \omega)}(J)$ is monomial-free if and only if $\text{in}_{(-1, \omega)}(J)|_{t=1}$ is monomial-free; see Lemma 6.3.2 below. Thus, to find the ω 's for which $\text{t-in}_\omega(I)$ is monomial-free we need to find all ω 's for which $\text{in}_{(-1, \omega)}(J)$ is monomial-free. We summarize this in the following Lemma.

Lemma 6.3.1 *Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be an ideal generated by elements of $\mathbb{C}[t, x_1, \dots, x_n]$. Then*

$$\{-1\} \times \mathcal{T}(I) = \mathcal{T}(I \cap \mathbb{C}[t, x_1, \dots, x_n]) \cap (\{-1\} \times \mathbb{R}^n).$$

Lemma 6.3.2 *Let $I \subseteq k[x_1, \dots, x_n]$ be an ideal in the polynomial ring over a field k . Suppose I is homogeneous with respect to a grading $(-1, \omega) \in \mathbb{R}^n$. Then I is monomial-free if and only if $I_{|x_1|=1}$ is monomial-free.*

Proof. (Assuming the axiom of choice.) Without loss of generality we may assume that $\omega \in \mathbb{Q}^{n-1}$ since the homogeneity space of I is defined by rational equations. Choose $N \in \mathbb{N}_{>0}$ such that $N\omega \in \mathbb{Z}^n$. Clearly, if I contains a monomial then so does $I_{|x_1|=1}$. For the other direction first assume that k is algebraically closed. If I does not contain a monomial then according to Hilbert's Nullstellensatz $x_1 \cdots x_n \notin \sqrt{I} = I(V(I))$. That $x_1 \cdots x_n$ does not vanish on $V(I)$ means that there must exist $a \in (k^*)^n \cap V(I)$. For $s \in k^*$ consider the points $\alpha_s = (a_1 s^{-N}, a_2 s^{N\omega_1}, \dots, a_n s^{N\omega_{n-1}})$ which are in $V(I) \cap (k^*)^n$ since I is $(-N, N\omega)$ -homogeneous. Choose s such that $s^N = a_1$. Now $\alpha_s \in V(I) \cap (\{1\} \times (k^*)^n)$ and therefore $((\alpha_s)_2, \dots, (\alpha_s)_n)$ is in $V(I_{|x_1|=1}) \cap (k^*)^{n-1}$. This proves that $I_{|x_1|=1}$ is monomial-free. If k is not algebraically closed we consider the ideal I generates in $\bar{k}[x_1, \dots, x_n]$ where \bar{k} is the algebraic closure. Deciding if the ideals contain a monomial can be done with Gröbner basis techniques; see Algorithm 7.1.6. Hence the answer does not depend of the field and we may apply the lemma for the ideals in the polynomial ring over the algebraically closed field \bar{k} . \square

Remark 6.3.3 Notice that by substituting t by $t^{\frac{1}{N}}$ the same argument shows that the situation where I is generated by polynomials in $\mathbb{C}(t^{\frac{1}{N}})[x_1, \dots, x_n]$ can also be reduced to the $\mathbb{C}[t, x_1, \dots, x_n]$ case.

Example 6.3.4 Let $f = 1 + t^2x + ty + t^4xy + t^6x^2 + t^4y^2$. The tropical variety $\mathcal{T}(\langle f \rangle_{\mathbb{C}[t,x,y]}) \subseteq \mathbb{R}^3$ is a union of 12 two-dimensional polyhedral cones. Nine of these intersect the $t = -1$ plane. The intersection is shown in Figure 6.2. It is equal to $\{-1\} \times \mathcal{T}(\langle f \rangle_{\mathbb{C}\{\{t\}\}[x,y]})$ which coincides with the tropical curve in Example 6.0.6.

A well-known strategy for doing computations in \mathbb{C} is to consider an algebraic field extension $\mathbb{Q}(\alpha)$ of \mathbb{Q} where α is an algebraic number. Here $\mathbb{Q}(\alpha)$ denotes the smallest field containing \mathbb{Q} and α . The field $\mathbb{Q}(\alpha)$ is isomorphic to $\mathbb{Q}[a]/M$ where M is the ideal generated by the minimal polynomial $m \in \mathbb{Q}[a]$ of α . The minimal polynomial is the monic polynomial in $\mathbb{Q}[a]$ of smallest degree having α as a root. The polynomial m is irreducible and M is a prime ideal, but also a maximal ideal.

We will see that tropical varieties behave nicely with respect to this kind of field extension. This is useful for practical reasons since the implementations need not know of any other fields than \mathbb{Q} . Let φ be the homomorphism $\mathbb{Q}[a] \rightarrow \mathbb{Q}[a]/M$ taking an element to its coset. This homomorphism extends to $\varphi : \mathbb{Q}[a, x_1, \dots, x_n] \rightarrow (\mathbb{Q}[a]/M)[x_1, \dots, x_n]$. We compare the tropical varieties of $I \subseteq (\mathbb{Q}[a]/M)[x_1, \dots, x_n]$ and of $\varphi^{-1}(I)$.

Lemma 6.3.5 [42, Lemma 3.12] *Let k be a field and $M = \langle m \rangle \subseteq k[a]$ a maximal ideal where m is not a monomial. Let $I \subseteq (k[a]/M)[x_1, \dots, x_n]$ be an ideal. For $\omega \in \mathbb{R}^n$ we have*

$$\text{in}_\omega(I) \text{ contains a monomial} \iff \text{in}_{(0,\omega)}(\varphi^{-1}(I)) \text{ contains a monomial}$$

where $\varphi : k[a, x_1, \dots, x_n] \rightarrow (k[a]/M)[x_1, \dots, x_n]$ is the homomorphism taking elements to their cosets.

Proof. \Leftarrow : Suppose $\text{in}_{(0,\omega)}\varphi^{-1}(I)$ contains a monomial. This means that there exists an $f \in \varphi^{-1}(I)$ with $\text{in}_{(0,\omega)}(f)$ being a monomial. The polynomial $\varphi(f)$ is in I . When applying φ the monomial $\text{in}_{(0,\omega)}(f)$ maps to a monomial whose coefficient in $k[a]/M$ has a representative $h \in k[a]$ with just one term. The representative h cannot be 0 modulo M since M does not contain a monomial. As the Newton polytope of $\varphi(f)$ is clearly contained in the projection of the Newton polytope of f to n dimensions, $\text{in}_\omega(\varphi(f))$ is a monomial.

\Rightarrow : Suppose $\text{in}_\omega(I)$ contains a monomial. This means that there exists an $f \in I$ with $\text{in}_\omega(f)$ being a monomial. Let g be in $\varphi^{-1}(I)$ such that g maps to f under the surjection φ and with the further condition that the support of g projected to n dimensions equals the support of f . The initial form $\text{in}_{(0,\omega)}(g)$ is a polynomial with all exponent vectors having the same x_1, \dots, x_n parts as $\text{in}_\omega(f)$ does. We must show that $\text{in}_{(0,\omega)}(\varphi^{-1}(I))$ contains a monomial. This is the same thing as showing that $(\text{in}_{(0,\omega)}(\varphi^{-1}(I)) : x_1 \cdots x_n^\infty)$ contains a monomial. Let g' be $\text{in}_{(0,\omega)}(g)$ with the common x -part removed from the monomials, that is $g' \in k[a]$. Notice, $\varphi(g') \neq 0$. We now have $g' \notin M$ and hence $\langle g' \rangle + M = k[a]$ since M is maximal. Now, M is a subset of, and g' is contained in $(\text{in}_{(0,\omega)}(\varphi^{-1}(I)) : x_1 \cdots x_n^\infty)$, implying that $(\text{in}_{(0,\omega)}(\varphi^{-1}(I)) : x_1 \cdots x_n^\infty) \supseteq k[a]$. This shows that $\text{in}_{(0,\omega)}(\varphi^{-1}(I))$ contains a monomial. \square

Remark 6.3.6 The preimage $\varphi^{-1}(I)$ contains m which has at least two terms, none of them containing the variables x_1, \dots, x_n . The tropical variety $\mathcal{T}(M)$ is just the coordinate hyperplane and therefore $\mathcal{T}(\varphi^{-1}(I)) \subseteq \{0\} \times \mathbb{R}^n$ and the above lemma completely describes $\mathcal{T}(\varphi^{-1}(I))$.

6.4 Decomposing tropical varieties

An algebraic variety can be decomposed by making a primary decomposition of its defining ideal. The same holds for tropical varieties. The union of two varieties $V(I)$ and $V(J)$ in $\mathbb{C}\{\{t\}\}^n$ is again a variety. Similarly, the union of the “shadows” under the negative valuation map, $\mathcal{T}(I) \cup \mathcal{T}(J)$, is a tropical variety. It is possible to give a simple proof of this that does not use Theorem 6.1.17 and is independent of properties of the field.

Proposition 6.4.1 *Let $I, J \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be ideals. For $\omega \in \mathbb{R}^n$ the initial ideal $\text{t-in}_\omega(I \cap J)$ contains a monomial if and only if $\text{t-in}_\omega(I)$ and $\text{t-in}_\omega(J)$ contain monomials. In particular, $\mathcal{T}(I) \cup \mathcal{T}(J) = \mathcal{T}(I \cap J)$.*

Proof. Suppose $\text{t-in}_\omega(I)$ and $\text{t-in}_\omega(J)$ both contains monomials. Lemma 6.1.10 shows that there exists $f \in I$ and $g \in J$ such that $\text{t-in}_\omega(f)$ and $\text{t-in}_\omega(g)$ are monomials. The initial form of the product $\text{t-in}_\omega(fg) = \text{t-in}_\omega(f)\text{t-in}_\omega(g)$ is a monomial in $\text{t-in}_\omega(I \cap J)$ since $fg \in I \cap J$. The other inclusion is clear since $I \cap J \subseteq I$ implies $\text{t-in}_\omega(I \cap J) \subseteq \text{t-in}_\omega(I)$ and $I \cap J \subseteq J$ implies $\text{t-in}_\omega(I \cap J) \subseteq \text{t-in}_\omega(J)$. \square

Similarly taking the radical of an ideal does not affect the tropical variety it defines:

Proposition 6.4.2 *Let $I \subseteq \mathbb{C}\{\{t\}\}[x_1, \dots, x_n]$ be an ideal. For $\omega \in \mathbb{R}^n$ the initial ideal $\text{t-in}_\omega(I)$ contains a monomial if and only if $\text{t-in}_\omega(\sqrt{I})$ contains a monomial. In particular, $\mathcal{T}(I) = \mathcal{T}(\sqrt{I})$.*

Proof. The inclusion $I \subseteq \sqrt{I}$ implies $\mathcal{T}(I) \supseteq \mathcal{T}(\sqrt{I})$. On the other hand, if $\omega \notin \mathcal{T}(\sqrt{I})$ then by Lemma 6.1.10 there must exist $f \in \sqrt{I}$ such that $\text{t-in}_\omega(f)$ is a monomial. For some $m \in \mathbb{N}$, $f^m \in I$. Hence $\text{t-in}_\omega(f)^m = \text{t-in}_\omega(f^m) \in \text{t-in}_\omega(I)$ and $\omega \notin \mathcal{T}(I)$. \square

Proposition 6.4.1 and Proposition 6.4.2 imply that, as a set, the tropical variety of an ideal I is the union of the tropical varieties defined by the minimal associated primes of I .

The second most important theorem of tropical geometry is the Bieri-Groves theorem, originally stated in [8] as a theorem about valuations. Rephrasing it in our initial ideal language it goes as follows; see also [54, Theorem 9.6].

Theorem 6.4.3 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a homogeneous (with respect to some positive grading) monomial-free prime polynomial ideal of dimension d . The polyhedral fan $\mathcal{T}(I)$ is a pure d -dimensional complex.*

In Chapter 8 we give a proof of the theorem.

It is our goal to develop an algorithm for computing the tropical variety of a homogeneous prime ideal as a polyhedral complex which is more efficient than computing the entire Gröbner fan. Combining Lemma 6.2.5, Proposition 6.4.1 and Proposition 6.4.2 and a method for computing primary decompositions of ideals in $\mathbb{C}[x_1, \dots, x_n]$ this gives an algorithm for computing tropical varieties of arbitrary ideals in $\mathbb{C}[x_1, \dots, x_n]$ as sets.

6.4.1 Saturated initial ideals and decomposition

It is not possible to define a polyhedral structure on $\mathcal{T}(I)$ by considering initial ideals in $\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ instead of in $\mathbb{C}[x_1, \dots, x_n]$ — or equivalently considering $x_1 \cdots x_n$ saturated initial ideals. However considering *saturated initial ideals* ($\text{in}_\omega(I) : x_1 \cdots x_n^\infty$) sometimes gives interesting information about the components of the algebraic variety as the following example shows.

Example 6.4.4 To construct a tropical variety which does not have an obvious polyhedral structure, we consider the hyperplane $\mathcal{T}(abc+1)$ and the line $\mathcal{T}((a+$




Figure 6.6: A three-dimensional drawing of the tropical variety $\mathcal{T}(I \cap J)$ modulo the homogeneity space; see Example 6.4.4. The support is the union of a plane and a line.

$b + c, ab + bc + ca$). As we prefer ideals to be homogeneous we introduce the homogenizing variable d and the ideals to consider are $I = \langle abc + d^3 \rangle$ and $J = \langle a + b + c, ab + bc + ca \rangle$. The following is a generating set for $I \cap J$:

$$\begin{aligned} & \{ab^3cd + c^2d^4 + bcd^4 + b^2d^4 + abc^3d + ab^2c^2d, \\ & \quad a^2bcd + cd^4 + bd^4 + ad^4 + abc^2d + ab^2cd, \\ & \quad a^2b^2c + bcd^3 + acd^3 + abd^3 + ab^2c^2 + a^2bc^2\} \end{aligned}$$

The f-vector of the Gröbner fan of $I \cap J$ is $(1, 8, 18, 12)$. The homogeneity space has dimension 1. Six of the rays (two-dimensional) are contained in the hyperplane while the two remaining ones are sticking out to form the line. The hyperplane is divided into 6 three-dimensional cones and the remaining 12 three-dimensional cones connects the two groups of rays.

As a complex $\mathcal{T}(I \cap J)$ consists of the rays and the 6 three-dimensional cones in the hyperplane; see Figure 6.6. Given the support of $T(I \cap J) = T(I) \cup T(J)$ it is not clear at all that this is a natural choice of a polyhedral structure. One could hope that the non-monomial initial ideals in the Laurent polynomial ring define a structure. Unfortunately this is not the case. There are only 3 saturated initial ideals for $I \cap J$ different from $\langle 1 \rangle$.

- For the two rays ($\omega = \pm(1, 1, 1, 0)$) we have $(\text{in}_\omega(I \cap J) : x_1 \cdots x_n^\infty) = J$.
- For any point in the hyperplane not in $C_0(I \cap J)$ we get I .
- For $\omega = 0$ we get the saturation of the original ideal $I \cap J$.

In particular, the equivalence classes induced by taking saturated initial ideals are not necessarily convex.

The example seems to suggest that we might be able to do primary decomposition by computing tropical varieties. Unfortunately it is easy to come up with examples showing that this does not always work.

Example 6.4.5 Consider the ideal $\langle a + bx + cy + dxy \rangle \subseteq \mathbb{C}[x, y]$ where $a, b, c, d \in \mathbb{C}$. The tropical variety as a set is the union of the coordinate axes which is the union of two tropical varieties. However the polynomial only factors when $ad = bc$.

The example shows that the tropical variety can be reducible even if the algebraic variety is not. On the other hand it is possible that several components of the algebraic variety have the same tropical “shadow”:

Example 6.4.6 Let $\langle (x - 1)(x - 2) \rangle \subseteq \mathbb{C}[x]$. The tropical variety is just the origin which is irreducible. However, the algebraic variety is not irreducible.

Chapter 7

Tropical algorithms

In this chapter we describe the tropical algorithms presented in [10].

7.1 Tropical bases

Proposition 6.4.1 says that the union of two tropical varieties is again a tropical variety. For algebraic varieties the same is true for intersections. However, the intersection of two tropical varieties is not necessarily a tropical variety.

Example 7.1.1 Let $f_1 = x^2 + xy + 1, f_2 = x + 1 \in \mathbb{C}[x, y]$. The tropical variety $\mathcal{T}(f_1)$ is the union of three halflines and $\mathcal{T}(f_2)$ is a line. The intersection $\mathcal{T}(f_1) \cap \mathcal{T}(f_2)$ is the half line $\{0\} \times (-\infty, 0]$. Suppose this is the tropical variety of some ideal. Due to the invariance under positive scaling this cannot be decomposed further and thus must be defined by a prime ideal and according to Theorem 6.4.3 this prime ideal must have dimension 1. A prime ideal of dimension 1 in $\mathbb{C}[x, y]$ is a principal ideal. However, $\{0\} \times (-\infty, 0]$ is not a tropical hypersurface. This is a contradiction.

A finite intersection of tropical hypersurfaces is called a *tropical prevariety*. An algebraic variety is the intersection of the hypersurfaces defined by the elements of any finite generating set for a defining ideal of the variety. This is not the case for tropical varieties. A finite subset $\{f_1, \dots, f_t\} \subseteq I$ is called a *tropical basis* of I if $\langle f_1, \dots, f_t \rangle = I$ and $\bigcap_i \mathcal{T}(f_i) = \mathcal{T}(I)$.

Let $\text{in}_\omega(I)$ be some initial ideal of I . If $C_\omega(I) \not\subseteq \mathcal{T}(I)$ then $\text{in}_\omega(I)$ contains a monomial x^m and there must exist a polynomial $f \in I$ such that $x^m = \text{in}_\omega(f)$. Such an f is called a *witness* for $C_\omega(I)$ if the following holds:

$$\text{rel int}(C_\omega(I)) \cap \mathcal{T}(f) = \emptyset.$$

The witness shows that $C_\omega(I)$ is not in the tropical variety.

Proposition 7.1.2 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a homogeneous ideal and $\omega \in \mathbb{R}^n$. If $\text{in}_\omega(I)$ contains a monomial then $C_\omega(I)$ has a witness.*

Proof. Let $x^m \in \text{in}_\omega(I)$ and \prec be a term order. According to Algorithm 4.2.1 and Lemma 3.1.14 the polynomial $f := x^m - (x^m)^{\mathcal{G}_{\prec, \omega}(I)}$ satisfies $f \in I$ and

$\text{in}_\omega(f) = x^m$. To prove that f is a witness we simply observe that any ω' in $\text{rel int}(C_\omega(I))$ would have given us the same f since $\text{in}_\omega(I) = \text{in}_{\omega'}(I)$ implies $\text{in}_{\prec_\omega}(I) = \text{in}_{\prec}(\text{in}_\omega(I)) = \text{in}_{\prec}(\text{in}_{\omega'}(I)) = \text{in}_{\prec_{\omega'}}(I)$ and $\mathcal{G}_{\prec_\omega}(I) = \mathcal{G}_{\prec_{\omega'}}(I)$. \square

Remark 7.1.3 Notice that the process described in the above proof will produce a homogeneous witness.

Corollary 7.1.4 [10, Theorem 11] *Every polynomial ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ has a finite tropical basis.*

Proof. We start by considering the case where I is homogeneous. For any $\text{in}_\omega(I)$ containing a monomial we take a homogeneous witness. The finite set of these witnesses together with a finite homogeneous generating set of I forms a tropical basis. If I is not homogeneous we may consider its homogenization $I^h \subseteq \mathbb{C}[x_0, \dots, x_n]$ which has a homogeneous tropical basis B . Substituting $x_0 = 1$ in the basis no terms will cancel and the new tropical hypersurfaces will cut out $\mathcal{T}(I)$ since

$$\begin{aligned} \{0\} \times \mathcal{T}(I) &=_{\text{Lemma 6.2.5}} \mathcal{T}(I^h) \cap (\{0\} \times \mathbb{R}^n) = \left(\bigcap_{f \in B} \mathcal{T}(f) \right) \cap (\{0\} \times \mathbb{R}^n) = \\ & \bigcap_{f \in B} (\mathcal{T}(f) \cap (\{0\} \times \mathbb{R}^n)) =_{\text{Lemma 6.2.5}} \bigcap_{f \in B} (\{0\} \times \mathcal{T}(f|_{x_0=1})) = \{0\} \times \bigcap_{f \in B} \mathcal{T}(f|_{x_0=1}). \end{aligned}$$

Hence we have a tropical basis. \square

In other words a tropical variety is a tropical prevariety. We argue that the proof of Corollary 7.1.4 is constructive. Checking if a homogeneous ideal contains a monomial can be done with the following lemma and repeatedly use of Algorithm 2.5.2 for computing saturations as explained in Section 2.5. Let k be a field.

Lemma 7.1.5 *The ideal $I \subseteq k[x_1, \dots, x_n]$ contains a monomial if and only if $(I : x_1 \cdots x_n^\infty) = \langle 1 \rangle$.*

This leads to the following algorithm for finding a monomial in an ideal:

Algorithm 7.1.6 [10, Algorithm 3] Monomial in Ideal

Input: A set of generators for an ideal $I \subseteq k[x_1, \dots, x_n]$.

Output: A monomial $m \in I$ if one exists, **no** otherwise.

```
{
  If  $((I : x_1 \cdots x_n^\infty) \neq \langle 1 \rangle)$  return no;
   $m := x_1 \cdots x_n$ ;
  While  $(m \notin I)$   $m := m \cdot x_1 \cdots x_n$ ;
  Return  $m$ ;
}
```

Even if I is not homogeneous the test $(I : x_1 \cdots x_n^\infty) \neq \langle 1 \rangle$ can still be performed with Algorithm 2.5.2 by homogenizing any generating set of I . This follows from Lemma 6.2.5 with $\omega = 0$.

If the monomial produced by Algorithm 7.1.6 has small exponents we can hope for the lifted witness of Proposition 7.1.2 not to have too high degree or too many terms. One option for finding a small monomial is to use [50, Algorithm 4.2.2] which computes the ideal generated by all multi-homogeneous polynomials in an ideal. The multigrading is given by a matrix $A \in \mathbb{Z}^{d \times n}$. The i th variable gets the i th column as its degree. If we choose A to be the identity matrix this algorithm will produce the ideal generated by all monomials in the ideal. Picking a small minimal generator is a good choice of m . In practise we shall be less careful with the choice. The following recursive algorithm computes the $x_1 < x_2 < \dots$ lexicographically smallest monomial in an ideal I .

Algorithm 7.1.7 Lexicographically smallest monomial in ideal

Input: An index i and a (homogeneous) ideal $I \subseteq k[x_1, \dots, x_n]$ containing a monomial in the variables x_i, \dots, x_n .

Output: The $x_1 < x_2 < \dots$ lexicographically smallest monomial in $I \cap k[x_i, \dots, x_n]$.

```

{
  If ( $i > n$ )
    Return 1;
  else
    Compute  $J := (I : x_i^\infty)$ ;
    Apply the algorithm recursively to find the  $x_1 < x_2 < \dots$ 
      lexicographically smallest monomial  $m \in J \cap k[x_{i+1}, \dots, x_n]$ ;
    Keep multiplying  $m$  by  $x_i$  until  $m \in I$ ;
    Return  $m$ ;
}

```

Proof. If $i > n$ then by the assumption on the input $1 \in I$ and this is the smallest monomial in I . Let us consider the case $i \leq n$. The ideal $J = (I : x_i^\infty)$ contains a monomial in $k[x_{i+1}, \dots, x_n]$ since I contains a monomial in $k[x_i, \dots, x_n]$ by assumption. Recursively we compute the lexicographically smallest such monomial. For a monomial to be the lexicographically smallest monomial in $I \cap k[x_i, \dots, x_n]$ it must be as small as possible on the variables x_{i+1}, \dots, x_n as these are the most important ones in our lexicographic ordering. For this reason the constructed m is the smallest monomial in $I \cap k[x_i, \dots, x_n]$. \square

Computing every cone of the Gröbner fan of a homogeneous ideal and checking if the initial ideals contain a monomial we have an algorithm for computing the tropical variety and Proposition 7.1.2 gives a tropical basis. In Section 7.2.1 we present a much better algorithm for computing both of these in the case of a tropical curve and in Section 7.4 we give a better algorithm for computing tropical varieties defined by prime ideals.

7.2 Computing tropical prevarieties

In this section we will discuss the problem of going from a tropical basis to a polyhedral complex covering its tropical variety, or, more generally, given a list

of polynomials f_1, \dots, f_m with constant coefficients computing the common refinement $\bigwedge_i \mathcal{T}(f)$. This computation only depends on the Newton polytopes of the polynomials. For this reason we introduce the notation $\mathcal{T}(P)$ for the subfan of the normal fan of a polytope $P \subseteq \mathbb{R}^n$ consisting of all cones of dimension less than n . Let P_i be the Newton polytope of f_i for $i = 1, \dots, m$.

Definition 7.2.1 Let P_1, \dots, P_m be polytopes in \mathbb{R}^n and Q their Minkowski sum. Notice that every face $\text{face}_\omega(Q)$ of Q can be written uniquely as a sum of faces of P_1, \dots, P_m :

$$\text{face}_\omega(Q) = \text{face}_\omega(P_1) + \dots + \text{face}_\omega(P_m)$$

where $\omega \in \mathbb{R}^n$. The face $\text{face}_\omega(Q)$ is *mixed* if $\dim(\text{face}_\omega(P_i)) > 0$ for all i .

Lemma 7.2.2 Let P_1, \dots, P_m and Q be as above. The set of normal cones of the mixed faces of Q equals $\bigwedge_i \mathcal{T}(P_i)$.

Lemma 7.2.2 shows that we are interested in computing the mixed faces of a Minkowski sum. Typically the number of mixed faces is much lower than the total number of faces in the sum.

Example 7.2.3 Consider the Newton polytopes of the 20 3×3 minors of the matrix in Example 6.2.4. Each of these 4-dimensional polytopes have f-vector $(6, 15, 18, 9, 1)$. Their Minkowski sum is 8-dimensional and has f-vector

$$(3240, 12960, 21750, 19890, 10710, 3390, 585, 45, 1).$$

Of the 19890 three-dimensional faces only 630 are mixed. All other mixed faces are of higher dimension. The set of mixed faces is not a polyhedral complex. The set of normal cones of the mixed faces on the other hand is, as we have seen in Lemma 7.2.2, a polyhedral complex. Its f-vector is $(1, 45, 315, 930, 1260, 630)$. Comparing the f-vectors to those in Example 6.2.4 it seems most likely that the normal fan of the Minkowski sum equals the Gröbner fan and that the set of normal cones of the mixed faces is the tropical variety – implying that the 20 minors is a tropical basis. Usually these pairs of fans do not coincide.

The number of mixed faces can be exponential n . See [57] for a proof of NP-hardness of an associated decision problem. The mixed faces do not seem to be connected in any way that is useful for an efficient enumeration.

The case where $n = m + 1$ and the n th coordinate of the vertices of P_1, \dots, P_m is generic is of special interest in the polyhedral homotopy method for solving zero-dimensional polynomial systems numerically; see [31] and [54]. In this case the problem is well-studied and goes by the name “mixed cell computation” (of a subdivision) or the more misleading name “mixed volume computation”; see [23]. Only the lower mixed faces are of interest in this case.

In the homotopy case the generic last coordinate and the number of polytopes guarantee that all mixed lower faces are of dimension $n - 1$. In our general case we can have mixed faces of different dimensions.

Example 7.2.4 Let $P_1 = [0, 1]^3 \subset \mathbb{R}^3$ be the 3-dimensional cube and $P_2 = \text{conv}((0, 1), (0, -1), (1, 0), (-1, 0)) \times [0, 1]$ be the rotated cube. The f-vectors of $\mathcal{T}(P_1)$ and $\mathcal{T}(P_2)$ are $(1, 6, 12)$. The refinement $\mathcal{T}(P_1) \wedge \mathcal{T}(P_2)$ has f-vector $(1, 10, 8)$. Hence $P_1 + P_2$ has 1 mixed face of dimension 3, 10 of dimension 2 and 8 of dimension 1.

Each edge E_{ij} in one of our polytopes P_i gives rise to a set of linear equations and inequalities L_{ij} describing the normal cone of E_{ij} . In the generic homotopy case finding a normal of a mixed lower facet amounts to finding a choice c_i of edges E_{c_i} , one from each polytope, such that the system $\bigcup_i L_{ic_i}$ has a feasible solution with the n th coordinate negative. Thus finding all mixed faces amounts to a combinatorial search and feasibility checking of many linear programming problems. How to cleverly organized the combinatorial search taking advantage of the inner workings of the simplex algorithm is described in [23].

If the polytopes are not in generic position the above method might compute a set of edges whose Minkowski sum is not a mixed face, namely, the edges may be part of larger faces of the P_i 's whose Minkowski sum is a mixed face.

Example 7.2.5 Consider the Minkowski sum of the square $P_1 = ([0, 1] \times [0, 1]) \times \{0\}$ and the diagonal edge $P_2 = \text{conv}((0, 0, 0), (1, 1, 0))$. For all 4 possible pairs of edges from the two polytopes the combined linear systems are feasible with the solution $(0, 0, -1)$. However, the Minkowski sum of any edge of the square and the diagonal edge is not a face of $P_1 + P_2$. One way to interpret this is to say that the mixed face $\text{face}_{(0,0,-1)\mathcal{T}}(P_1 + P_2)$ is computed four times.

The above example shows that we need to be more careful with the combinatorics. In [10] we suggested the following straight forward algorithm for computing common refinements. We say that a subset of cones $S \subseteq \mathcal{F}$ is a *representation* of a fan \mathcal{F} if $\bigcup_{C \in S} C$ is the support of \mathcal{F} .

Algorithm 7.2.6 Common Refinement

Input: Representations S_1 and S_2 of two fans \mathcal{F}_1 and \mathcal{F}_2 .

Output: A representation S of $\mathcal{F}_1 \wedge \mathcal{F}_2$.

$$\left\{ \begin{array}{l} S := \emptyset; \\ \text{For every pair } (C_1, C_2) \in S_1 \times S_2 \\ \quad S := S \cup \{C_1 \cap C_2\}; \end{array} \right\}$$

An implementation of the above algorithm requires a method for detecting duplicates of cones. This can be achieved by writing the cones in a unique form.

Algorithm 7.2.6 also has the disadvantage that the same cone may be computed several times. In fact no matter which cones we choose to intersect we will always get some cone. In particular, the cone $\{0\}$ can be computed many times. We solve this problem by introducing strict inequalities as explained below.

We may write the support of each $\mathcal{T}(P_i)$ as a disjoint union of half open cones. Here a *half open cone* is a finite intersection of subspaces, closed half




Figure 7.1: The three disjoint unions in Example 7.2.7. Four half open cones of which one includes the origin. Three half open cones of which one includes the origin. All possible intersections of two cones, one from each of the two unions.

spaces and open half spaces. We apply the same strategy as in Algorithm 7.2.6 above to the disjoint unions of half cones. The advantage is that intersections of half open cones can be empty – in which case the intersection can safely be ignored. Furthermore, the output is a disjoint union of half cones covering $\bigcap_i \mathcal{T}(P_i)$.

To get the support of $\mathcal{T}(P_i)$ written as a disjoint union of half open cones we may take the set of half open cones to be the relative interiors of the cones in $\mathcal{T}(P_i)$. In practise it is an advantage to use as few cones as possible. It is not difficult to see that it is possible to choose just one half open cone for each maximal cone in $\mathcal{T}(P_i)$ and that this is optimal.

Example 7.2.7 Consider the polynomials $f_1 = xy + x + y + 1$ and $f_2 = x + y + x^2y$. The supports of $\mathcal{T}(f_1)$ and $\mathcal{T}(f_2)$ can be written as disjoint unions as shown in Figure 7.2.7. We get the support of $\mathcal{T}(f_1) \wedge \mathcal{T}(f_2)$ written as a disjoint union by forming all intersections of pairs of half open cones from the two fans.

Typically the refinement algorithm is applied several times to compute refinements of more than just two fans, or an algorithm is written that will handle more refinement in a single step. Depending on the purpose of the computation we may be interested in extracting the polyhedral structure of $\bigwedge_i \mathcal{T}(P_i)$ by computing all faces of the half open cones in the intersection.

It would be an interesting research project to see how experience from the generic homotopy case can be applied to our non-generic setting. It seems that at least all implementations for the homotopy method rely on the genericity.

7.2.1 Tropical curves

We will define tropical curves in the case of constant coefficients. The tropical curve in Example 6.0.6 and Figure 6.2 was defined for non-constant coefficients and is not an example of the curves we will study in this section.

Definition 7.2.8 An ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ is said to define a *tropical curve* if $\dim(I) = \text{homog}(I) + 1$ and I has a monomial-free initial ideal $\text{in}_\omega(I)$ with $\text{homog}(\text{in}_\omega(I)) = \dim(I)$.

Lemma 7.2.9 *Let I be an ideal defining a tropical curve. As a polyhedral complex the tropical variety $\mathcal{T}(I)$ consists of the homogeneity space $C_0(I)$ of dimension $\text{homog}(I)$ and a finite number of $(\text{homog}(I) + 1)$ -dimensional cones.*

Proof. The ideal I has a minimal primary decomposition $I = \bigcap_i Q_i$ and the support of $\mathcal{T}(I)$ is $\bigcap_i \text{supp}(\mathcal{T}(\sqrt{Q_i}))$. By Theorem 6.4.3 each $\mathcal{T}(\sqrt{Q_i})$ is pure of dimension $\dim(\sqrt{Q_i})$ which is at most $\dim(I)$ since $\sqrt{Q_i}$ is a prime containing I . In fact this dimension is attained for at least one associated prime $\sqrt{Q_i}$. This proves that $\mathcal{T}(I)$ is a complex of dimension $\dim(I)$. The complex $\mathcal{T}(I)$ contains the cone $C_0(I)$. This proves the claim. \square

In other words, modulo $C_0(I)$ the tropical variety is just a union of half-lines meeting in a point. This is the important fact that makes the following algorithm work.

Algorithm 7.2.10 [10, Algorithm 5] Tropical Basis of Curve

Input: A set of homogeneous generators \mathcal{G} for an ideal I defining a tropical curve.

Output: A homogeneous tropical basis \mathcal{G}' of I .

```

{
  Compute a finite set of cones  $S$  covering  $\text{supp}(\bigwedge_{g \in \mathcal{G}} \mathcal{T}(\langle g \rangle))$ 
  such that every  $C \in S$  has  $C_0(I)$  as a face;
  For every  $C \in S$ 
  {
    Let  $w$  be a generic relative interior point in  $C$ ;
    If  $(\text{in}_w(I)$  contains a monomial) then
      add a homogeneous witness for  $C_w(I)$  to  $\mathcal{G}$  and restart the algorithm;
    }
  }
   $\mathcal{G}' := \mathcal{G}$ ;
}

```

We will assume that the input is homogeneous with respect to any grading in the homogeneity space $C_0(I)$. A consequence is that we can choose all cones in the algorithm to be invariant under translation by vectors in $C_0(I)$.

Proof. The algorithm terminates because I has only finitely many Gröbner cones and the number of Gröbner cones with relative interior intersecting the tropical prevariety $\text{supp}(\bigwedge_{g \in \mathcal{G}} \mathcal{T}(\langle g \rangle))$ decreases strictly in every iteration. If a vector w passes the monomial test (which verifies $w \in \mathcal{T}(I)$) then C has dimension 0 or 1 modulo the homogeneity space since we are looking at a curve and w is generic in C . Any other relative interior point w' of C would also have passed the monomial test since modulo $C_0(I)$ the point w' is a scaled version of w . This property fails if $\mathcal{T}(I)$ is not a tropical curve since the dimension of C might be higher than 1 modulo the homogeneity space. Hence, when we terminate only points in the tropical variety are covered by S and \mathcal{G}' is a tropical basis. \square

Remark 7.2.11 It is not clear how the generic vector w in Algorithm 7.2.10 can be computed. Here we will explain how to compute its initial ideal $\text{in}_w(I)$

using the perturbation techniques from Section 4.3. We must clarify what we mean by generic. The following condition will suffice for our purpose: w is generic in C if w is not in any $(\dim(C) - 1)$ -dimensional Gröbner cone of I . First we compute a positive vector v_1 (not necessarily generic) in the relative interior of C which exists since I is homogeneous. Let $V = \{v_1, \dots, v_s\}$ be a basis for the space spanned by C . For sufficiently small $\varepsilon > 0$ the vector $w_\varepsilon = v_1 + \varepsilon v_2 + \dots + \varepsilon^{s-1} v_s$ stays in one Gröbner cone of I of dimension at least s . To see this, extend V to a basis V' for \mathbb{R}^n and write the rows in a matrix A which now defines a term order \prec_A by Definition 2.2.8. By an argument similar to the proof of Lemma 2.2.9 we have that for small $\varepsilon > 0$ the vector w_ε satisfies the finitely many conditions of Corollary 3.1.10 to be in $C_{\prec_A}(I)$. Hence by Corollary 3.1.13 we can compute $\text{in}_{w_\varepsilon}(I)$ by taking w_ε initial forms of $\mathcal{G}_{\prec_A}(I)$. To do this we need to decide for each vector in a finite set of exponent vector differences if the vector is perpendicular to w_ε . For sufficiently small $\varepsilon > 0$ this only happens if the vector is perpendicular to all vectors in V . This shows that $\text{in}_{w_\varepsilon}(I)$ does not depend on ε for $\varepsilon > 0$ sufficiently small. The Gröbner cone $C_{w_\varepsilon}(I)$ has dimension at least s as wanted since $\text{in}_{w_\varepsilon}(I)$ is homogeneous with respect to any vector in V . Any vector w in the relative interior of $C_{w_\varepsilon}(I)$ is generic and has initial ideal $\text{in}_w(I) = \text{in}_{w_\varepsilon}(I)$ which we have shown how to compute.

A representation of the tropical curve is gotten by applying Algorithm 7.2.6 to the tropical hypersurfaces defined by the output of Algorithm 7.2.10.

7.3 Connectedness of tropical varieties of prime ideals

Definition 7.3.1 Let \mathcal{F} be a polyhedral complex. A *ridge path* between two polyhedra $S, T \in \mathcal{F}$ of the same dimension is a sequence of polyhedra in \mathcal{F}

$$S = P_0, \dots, P_m = T$$

of some length $m + 1$ such that for all $i = 1, \dots, m$ the intersection $P_{i-1} \cap P_i$ is a facet of P_{i-1} and of P_i .

A ridge path consists of polyhedra of the same dimension and we use the word ridge since $P_{i-1} \cap P_i$ is ridge in \mathcal{F} .

In [10] we presented the following theorem which we shall refer to as Speyer's Theorem in this thesis. It is an extension to the Bieri-Groves Theorem 6.4.3.

Theorem 7.3.2 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a homogeneous monomial-free prime ideal. Then any two maximal cones in $\mathcal{T}(I)$ are connected by a ridge path in $\mathcal{T}(I)$.*

When any two maximal cones are connected by ridge paths we also say that the complex is *connected in codimension one*. Since $\mathcal{T}(I)$ is only defined when I is homogeneous, the Theorem has only been stated for homogeneous ideals. However, being connected in codimension one is really a property of the

support of $\mathcal{T}(I)$, so by Lemma 6.2.5 we know that the theorem also holds for inhomogeneous ideals.

A proof of Theorem 7.3.2 is presented in [10]. It goes by induction on the dimension of the ideal. The base case relies on a result in [17]. The induction step uses a multiplicative change of variables, the transverse intersection lemma also described in [10] and the Kleinman-Bertini theorem. See [10] for details.

Remark 7.3.3 Similar to the proof that the Puiseux series field is algebraically closed it is important here that the field \mathbb{C} has characteristic 0. This is needed when applying the Kleinman-Bertini theorem.

7.4 Traversing tropical varieties

By Theorem 7.3.2 we know that the maximal cones in the tropical variety of a d -dimensional homogeneous prime ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ are connected by ridge paths. To traverse the variety we just need to traverse its maximal (d -dimensional) cones since the variety is pure. The key step is to go from one maximal cone to the next. Let $C_w(I)$ be a Gröbner cone of dimension d contained in $\mathcal{T}(I)$ and let u be a vector in the relative interior of one of its facets $C_u(I)$. The structure of $\mathcal{T}(I)$ around this facet is revealed by the initial ideal $\text{in}_u(I)$. According to Proposition 3.1.20

$$\text{in}_{u+\varepsilon v}(I) = \text{in}_v(\text{in}_u(I))$$

for $\varepsilon > 0$ sufficiently small. Observe that $\text{in}_u(I)$ has a $(d-1)$ -dimensional homogeneity space by Remark 3.1.21. The Krull dimension, however, stays fixed when we move to an initial ideal. Therefore $\text{in}_u(I)$ defines a tropical curve and we may apply Algorithm 7.2.10 to get a vector v in the relative interior of each d -dimensional cone in $\mathcal{T}(\text{in}_u(I))$. Using the above formula $\text{in}_{u+\varepsilon v}(I)$ can be computed. In the following we will explain how to compute the Gröbner cone $C_{u+\varepsilon v}(I)$.

Let \prec be a term order and keep it fixed for the rest of this section. We will represent any Gröbner cone $C_w(I)$ as a pair of marked reduced Gröbner bases $(\mathcal{G}_{\prec_w}(\text{in}_w(I)), \mathcal{G}_{\prec_w}(I))$. Notice that this representation does not depend on w but only on $C_w(I)$ since $\text{in}_{\prec_w}(I) = \text{in}_{\prec}(\text{in}_w(I))$ by Corollary 3.1.15. Notice also that $C_w(I)$ is a face of $C_{\prec_w}(I)$. As in Example 3.1.5 the defining inequality system of the polyhedral cone $C_w(I)$ can be deduced from the pair of bases. Applying the lifting step of the Gröbner walk (Algorithm 4.2.2) we can compute a pair of Gröbner bases representing $C_{u+\varepsilon v}(I)$. The following algorithm computes all neighboring cones in $\mathcal{T}(I)$ of a specified cone. The local step is illustrated in Figure 7.2.

Algorithm 7.4.1 [10, Algorithm 7] Neighbors

Input: A pair $(\mathcal{G}_{\prec_w}(\text{in}_w(I)), \mathcal{G}_{\prec_w}(I))$ representing a d -dimensional Gröbner cone $C_w(I) \in \mathcal{T}(I)$. Here I must be a homogeneous ideal of dimension d .

Output: A collection N of all pairs of the form $(\mathcal{G}_{\prec_{w'}}(\text{in}_{w'}(I)), \mathcal{G}_{\prec_{w'}}(I))$ representing Gröbner cones $C_{w'}(I) \in \mathcal{T}(I) \setminus \{C_w(I)\}$ which have a facet in common with $C_w(I)$.

```

{
  N := ∅;
  Compute the set  $\mathcal{F}$  of facets of  $C_w(I)$ ;
  For each  $F \in \mathcal{F}$ 
  {
    Compute the initial ideal  $J := \text{in}_u(I)$  where  $u \in \text{rel int}(F)$ ;
    Use Algorithm 7.2.10 and Algorithm 7.2.6 to produce a relative
      interior point  $v$  in each maximal cone of  $\mathcal{T}(J)$ ;
    For each such  $v$ 
    {
      Compute  $(\mathcal{G}_{\prec_v}(\text{in}_v(J)), \mathcal{G}_{\prec_v}(J)) = (\mathcal{G}_{(\prec_v)_u}(\text{in}_v(J)), \mathcal{G}_{(\prec_v)_u}(J))$ ;
      Apply Algorithm 4.2.2 to  $\mathcal{G}_{\prec_w}(I)$  and  $\mathcal{G}_{(\prec_v)_u}(J)$  to get  $\mathcal{G}_{(\prec_v)_u}(I)$ ;
       $N := N \cup \{(\mathcal{G}_{(\prec_v)_u}(\text{in}_v(J)), \mathcal{G}_{(\prec_v)_u}(I))\}$ ;
    }
  }
}

```

Proof. Proposition 3.1.4 gives a description of $C_w(I)$ in terms of linear inequalities and equations. Computing the facets of $C_w(I)$ is a problem in computational polyhedral geometry. Interior points can be computed by solving linear programming problems.

The identity $(\mathcal{G}_{\prec_v}(\text{in}_v(J)), \mathcal{G}_{\prec_v}(J)) = (\mathcal{G}_{(\prec_v)_u}(\text{in}_v(J)), \mathcal{G}_{(\prec_v)_u}(J))$ follows since J is u -homogeneous. We may apply Algorithm 4.2.2 since $u \in F \subset C_w(I) \subseteq C_{\prec_w}(I)$. We need to argue that $\{(\mathcal{G}_{(\prec_v)_u}(\text{in}_v(J)), \mathcal{G}_{(\prec_v)_u}(I))\}$ represents the Gröbner cone $C_{w'}(I)$ where $w' = u + \varepsilon v$ for $\varepsilon > 0$ sufficiently small. By Proposition 3.1.20 $\text{in}_{w'}(I) = \text{in}_v(\text{in}_u(I))$ for $\varepsilon > 0$ sufficiently small. So indeed $\mathcal{G}_{(\prec_v)_u}(\text{in}_v(J)) = \mathcal{G}_{\prec_v}(\text{in}_v(J)) = \mathcal{G}_{\prec}(\text{in}_v(J)) = \mathcal{G}_{\prec}(\text{in}_{w'}(I))$. It remains to argue that $\mathcal{G}_{(\prec_v)_u}(I) = \mathcal{G}_{\prec_{u+\varepsilon v}}(I)$. For $\varepsilon > 0$ sufficiently small the two term orders $(\prec_v)_u$ and $\prec_{u+\varepsilon v}$ agree on a finite set of monomials. In particular they give the same reduced Gröbner basis for I . \square

Applying the above method we get an exhaustive traversal algorithm for tropical varieties of prime ideals. If the ideal is not prime, the algorithm will compute all cones connected to the starting cone by ridge paths. Thus the algorithm computes a codimension-one-connected component of dimension d . The algorithm can be implemented in such a way that it does its enumeration up to symmetry with the same methods as for the symmetric Gröbner fan traversal explained in Section 4.5. It seems unlikely that it would be possible to find a good reverse search rule for tropical varieties.

To begin a traversal of a tropical variety we must first find a starting cone. The situation in Example 6.2.4 is typical. In the lower dimensions the Gröbner




Figure 7.2: The situation in Algorithm 7.4.1. The tropical variety of I is shown to the left and the tropical variety of $\text{in}_u(I)$ is shown to the right.

fan and the tropical variety almost agree while many cones are missing in the tropical variety in higher dimensions. In [10] we suggested the following method for computing a starting cone.

Algorithm 7.4.2 [10, Algorithm 9] Starting Cone

Input: Generators for a homogeneous ideal I whose tropical variety is pure of dimension d .

Output: A pair $(\mathcal{G}_{\prec_{w'}}(\text{in}_{w'}(I)), \mathcal{G}_{\prec_{w'}}(I))$ representing a d -dimensional Gröbner cone in $\mathcal{T}(I)$.

```

{
  If  $\text{homog}(I) = d$  then
    Return  $(\mathcal{G}_{\prec}(I), \mathcal{G}_{\prec}(I))$ ;
  If not
    {
      Guess a point  $w \in \text{supp}(\mathcal{T}(I)) \setminus C_0(I)$ ;
      Compute  $\mathcal{G}_{\prec_w}(I)$ ;
       $(\mathcal{G}_{\text{Init}}, \mathcal{G}_{\text{Full}}) := \text{Starting Cone}(\text{in}_w(I))$ ;
      Apply Algorithm 4.2.2 to  $\mathcal{G}_{\prec_w}(I)$  and  $\mathcal{G}_{\text{Full}}$ 
        to get a marked reduced Gröbner basis  $\mathcal{G}'$  for  $I$ ;
      Return  $(\mathcal{G}_{\text{Init}}, \mathcal{G}')$ ;
    }
}

```

Here the part of guessing a point $w \in \text{supp}(\mathcal{T}(I)) \setminus C_0(I)$ is unspecified.

Proof. If $\text{homog}(I) = d$ then the homogeneity space $C_0(I)$ is the only possible choice of cone. It is represented by $(\mathcal{G}_{\prec}(I), \mathcal{G}_{\prec}(I))$. If $\text{homog}(I) < d$ then by the hypothesis $\text{supp}(\mathcal{T}(I)) \setminus C_0(I) \neq \emptyset$. Let $w \in \text{supp}(\mathcal{T}(I)) \setminus C_0(I)$. The recursive call finds a d -dimensional cone in the pure d -dimensional complex $\mathcal{T}(\text{in}_w(I))$. Let u be a relative interior point in that cone. Then $(\mathcal{G}_{\text{Init}}, \mathcal{G}_{\text{Full}}) = (\mathcal{G}_{\prec_u}(\text{in}_u(\text{in}_w(I))), \mathcal{G}_{\prec_u}(\text{in}_w(I)))$. By Proposition 3.1.20 the initial ideal $\text{in}_u(\text{in}_w(I))$ is an initial ideal of I with respect to $w + \varepsilon u$ for $\varepsilon > 0$ sufficiently small. We already have the first basis in the pair representing $C_{w+\varepsilon u}(I)$,

namely, it is $\mathcal{G}_{\prec_{w+\varepsilon u}}(\text{in}_u(\text{in}_w(I))) = \mathcal{G}_{\prec_u}(\text{in}_u(\text{in}_w(I))) = \mathcal{G}_{\text{init}}$. Since $w \in C_{\prec_w}(I)$ we may apply Algorithm 4.2.2 to $\mathcal{G}_{\prec_w}(I)$ and $\mathcal{G}_{\prec_u}(\text{in}_w(I))$ to get $\mathcal{G}_{(\prec_u)_w}(I) = \mathcal{G}'$. We argue that $\mathcal{G}' = \mathcal{G}_{(\prec_u)_w}(I)$ equals $\mathcal{G}_{\prec_{w+\varepsilon u}}(I)$ for $\varepsilon > 0$ sufficiently small. It suffices to show that $\prec_{w+\varepsilon u}$ picks out the marked terms of $\mathcal{G}_{(\prec_u)_w}(I)$. However, for a fixed pair of monomials to compare the term orders agree for $\varepsilon > 0$ sufficiently small. Since there are only finitely many terms in a Gröbner basis the conclusion follows. \square

One strategy for guessing a vector $w \in \text{supp}(\mathcal{T}(I)) \setminus C_0(I)$ in Algorithm 7.4.2 is to compute all rays of some full-dimensional Gröbner and check if any of those are in the tropical variety. If not, we may try another Gröbner cone, or try a vector in the tropical prevariety defined by the given set of generators for I if the generating set is not too complex. This strategy works reasonable well in practise. In the worst case, however, using such strategies may be as complicated as computing the entire Gröbner fan.

Chapter 8

A proof of the Bieri Groves theorem

In this chapter we give a proof of Bieri Groves theorem based on [42, Corollary 6.13] (Theorem 8.2.1 below) and Lemma 8.1.3 from elimination theory. Theorem 8.2.1 can be proved using Krull's principal ideal theorem. We shall restrict ourselves to the case of constant coefficients and prove Bieri Groves theorem as it is formulated in Theorem 6.4.3 in terms of Gröbner cones. Notice that our proof only relies on the trivial inclusion in Theorem 6.1.17. The proof presented here has some similarities with a proof given in [54, Chapter 9]. Our proof is almost selfcontained with more details.

8.1 An equivalent theorem

We start by proving that the following theorem is equivalent to Theorem 6.4.3:

Theorem 8.1.1 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a monomial-free prime ideal with $C_0(I) \cap \mathbb{R}_{>0}^n \neq \emptyset$ and $\omega \in \mathbb{R}^n$. If $\text{in}_\omega(I)$ is monomial-free then $\dim(\text{in}_\omega(I) : x_1 \cdots x_n^\infty) = \dim(I)$.*

For a prime ideal I in the Laurent polynomial ring $\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, the theorem says that an initial ideal of I is either $\langle 1 \rangle$ or has dimension $\dim(I)$.

For the proof we need Lemma 8.1.2 and Proposition 8.1.6 below.

Lemma 8.1.2 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a monomial-free ideal of dimension d with an h -dimensional homogeneity space. Then $d \geq h$.*

Proof. Let $A \in \mathbb{Z}^{h \times n}$ be a matrix whose rows form a basis for the homogeneity space of I . By Hilbert's Nullstellensatz there exists a point $p \in V(I) \cap (\mathbb{C}^*)^n$ since I is monomial-free. Let a_1, \dots, a_n be the columns of A . For $t = (t_1, \dots, t_h) \in (\mathbb{C}^*)^h$ the point $\phi(t) := (t^{a_1} p_1, \dots, t^{a_n} p_n)$ belongs to $V(I)$ since I is homogeneous with respect to the gradings given by the rows of A . The toric ideal $I_A = \langle y^u - y^v : u, v \in \mathbb{Z}^n \text{ and } Au = Av \rangle \subseteq \mathbb{C}[y_1, \dots, y_n]$ is the ideal of the image of $\mu(t) := (t^{a_1}, \dots, t^{a_n})$ where $t = (t_1, \dots, t_h) \in (\mathbb{C}^*)^h$. The toric ideal I_A has dimension $\text{rank}(A) = h$ and is the image of the ideal

of $\phi((\mathbb{C}^*)^h)$ under the \mathbb{C} -algebra isomorphism $\Phi : \mathbb{C}[x_1, \dots, x_n] \rightarrow \mathbb{C}[y_1, \dots, y_n]$ taking x_i to $p_i^{-1}y_i$. Hence, $I \subseteq \Phi^{-1}(I_A)$ and I has dimension at least h . \square

Lemma 8.1.3 ([4, Theorem 1.22]) *Let $I \subseteq k[x_1, \dots, x_n]$ where k is algebraically closed. The projection $\{x_1 : x \in V(I)\}$ of $V(I)$ to the first axis is either k , the empty set, a finite set, or the complement of a finite set. The same holds for the projection of $V(I) \cap (k^*)^n$ to the first axis.*

The lemma is a special case of [4, Theorem 1.22] which states that the projection of a *constructable* set in k^n is constructable. The constructable sets in k^n are those obtained by doing a finite number of boolean operations on algebraic sets in k^n . A constructable set in k^1 is either a finite set or the complement of a finite set.

Recall that an algebraically closed field k is always infinite and that for a field extension $k \subsetneq k'$ the field k' is a k -vector space and thus $k' \setminus k$ is infinite as well.

Corollary 8.1.4 *Let $I \subseteq k[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be an ideal with dimension larger than 0 and k an algebraically closed field. Let k' be a strictly larger algebraically closed field containing k . Then there exists $p \in V_{k'^*}(I) \setminus V_{k^*}(I)$.*

Proof. The variety $V_{k^*}(I)$ is infinite. Hence the projection of $V_{k^*}(I)$ to some axis must be infinite. The projection of $V_{k'^*}(I)$ contains the projection of $V_{k^*}(I)$ and is thus infinite. By Lemma 8.1.3 the projection of $V_{k'^*}(I)$ is k' minus a finite set of points. Since $k' \setminus k$ is infinite this proves that $V_{k^*}(I) \neq V_{k'^*}(I)$. \square

Lemma 8.1.5 *If $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ and $p \in V_{\mathbb{C}\{\{t\}\}}(I)$ with $\text{val}(p) \in \mathbb{R}_{>0} \times \mathbb{R}^{n-1}$ (and thereby $\mathcal{T}(I) \cap (\mathbb{R}_{<0} \times \mathbb{R}^{n-1}) \neq \emptyset$) then $\mathcal{T}(I) \cap (\mathbb{R}_{>0} \times \mathbb{R}^{n-1}) \neq \emptyset$.*

Proof. Since $\text{in}_{-\text{val}(p)}(I)$ is monomial-free, $I \cap \mathbb{C}[x_1]$ cannot contain a non-zero polynomial. Hence $I \cap \mathbb{C}[x_1] = \langle 0 \rangle$. This shows that the Zariski closure of the projection of $V_{\mathbb{C}\{\{t\}\}}(I)$ to the x_1 -axis is the whole axis. In particular the image of the projection is not finite. By Lemma 8.1.3 it is the axis minus a finite set of points. Since there are infinitely many points with negative valuation in $\mathbb{C}\{\{t\}\}$ some of them are in the projection. Hence there exists $q \in V_{\mathbb{C}\{\{t\}\}}(I)$ with $\text{val}(q_1) < 0$. Consequently, according to the weak direction of Theorem 6.1.17 $-\text{val}(q) \in \mathcal{T}(I)$. \square

The homogeneity space $C_0(I)$ of an ideal in the Laurent polynomial ring and its dimension $\text{homog}(I)$ are defined in the natural way.

Proposition 8.1.6 *Let $I \subseteq \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be a d -dimensional ideal with $d > \text{homog}(I)$ then there exist an $\omega \in \mathbb{R}^n \setminus C_0(I)$ with $\text{in}_\omega(I)$ being monomial-free.*

Proof. We start by proving the case where $\text{homog}(I) = 0$. That is, we show that there exists $\omega \in \mathbb{R}^n \setminus \{0\}$ such that $\text{in}_\omega(I)$ is monomial-free. Let $J = I \cap \mathbb{C}[x_1, \dots, x_n]$. According to Corollary 8.1.4 there exists $p \in V_{\mathbb{C}\{\{t\}\}}(I) \setminus V_{\mathbb{C}^*}(I)$.

Suppose $\text{in}_\omega(I)$ contains a monomial for all $\omega \neq 0$. Then we must have $\text{val}(p) = 0$. Consequently, p has the form

$$(p_1, \dots, p_n) = (c_1 t^0 + \text{h.o.t.}, \dots, c_n t^0 + \text{h.o.t.})$$

with $c_i \in \mathbb{C}^*$ and where the higher order terms (h.o.t.) are different from 0 for at least one coordinate. With out loss of generality we may assume that this is the first one and call these higher order terms $q_1 := p_1 - c_1 t^0$. We make the substitution $x_1 = c_1 t^0 + y_1$ which gives us an ideal $J' \subseteq \mathbb{C}[y_1, x_2, \dots, x_n]$. The substitution is an isomorphism between the two polynomial rings taking J to J' . Notice that J' has the zero (q_1, p_2, \dots, p_n) . This shows that $\text{in}_{-\text{val}(q_1, p_2, \dots, p_n)}(J') = \text{in}_{-e_1}(J')$ is monomial-free. Lemma 8.1.5 shows that there exists an $\omega \in \mathcal{T}(J')$ with $\omega_1 > 0$. This means that for all $f \in J'$ the initial form $\text{in}_\omega(f)$ is not a monomial. Since $\omega_1 > 0$, taking f back to the original polynomial ring the initial form is preserved (with x_1 instead of y_1 , of course). Hence, $\omega \in \mathcal{T}(J) \setminus \{0\} = \mathcal{T}(I) \setminus \{0\}$.

Reduction from the general case where $\text{homog}(I) > 0$ is done by applying a multiplicative change of variables represented by a matrix $A \in GL_n(\mathbb{Z})$ as in Corollary 6.2.8. With the right choice of A , the transformed ideal $\varphi_A(I)$ can be generated by polynomials only involving $x_1, \dots, x_{n-\text{homog}(I)}$. Hence $C_0(\varphi_A(I)) = \text{span}(e_{n-\text{homog}(I)+1}, \dots, e_n)$. Deciding if there exists $\omega \in \mathbb{R}^n \setminus C_0(I)$ with $\text{in}_\omega(I)$ monomial-free is now the same as deciding if there exists $\omega' \in \mathbb{R}^{n-\text{homog}(I)} \setminus C_0(K)$ with $\text{in}_{\omega'}(K)$ monomial-free where $K := \varphi_A(I) \cap \mathbb{C}[x_1, \dots, x_{n-\text{homog}(I)}]$. The first part of the proof applies to K which has $\dim(K) = \dim(I) - \text{homog}(I) > \text{homog}(K) = 0$. \square

Proof of the equivalence of Theorem 6.4.3 and Theorem 8.1.1. In order to prove that Theorem 6.4.3 implies Theorem 8.1.1 notice that Theorem 6.4.3 and Proposition 3.1.20 imply that $\text{in}_\omega(I)$ has a d -dimensional tropical variety and thereby a monomial-free initial ideal J with a d -dimensional homogeneity space. We now have:

$$\begin{aligned} \dim(I) &= \dim(\text{in}_\omega(I)) \geq \dim(\text{in}_\omega(I) : x_1 \cdots x_n^\infty) \geq \dim(J : x_1 \cdots x_n^\infty) \\ &\geq \text{homog}(J : x_1 \cdots x_n^\infty) \geq \text{homog}(J) = \dim(I) \end{aligned}$$

where the first two inequalities follow from the fact that the dimension cannot increase when saturating. The third inequality follows from Lemma 8.1.2, the fourth from the fact that an ideal can only become more homogeneous when saturating and the last one from the choice of J . This proves that $\dim(\text{in}_\omega(I) : x_1, \dots, x_n^\infty) = \dim(I)$.

To prove that Theorem 8.1.1 implies Theorem 6.4.3 it suffices, by Proposition 3.1.20, to prove that for $\omega \in \mathcal{T}(I)$ the initial ideal $\text{in}_\omega(I)$ has a tropical variety containing a d -dimensional cone, or, equivalently, that $\text{in}_\omega(I)$ has a monomial-free initial ideal with a d -dimensional homogeneity space.

According to Theorem 8.1.1 we have $\dim(\text{in}_\omega(I) : x_1 \cdots x_n^\infty) = \dim(I)$. If $\text{homog}(\text{in}_\omega(I) : x_1 \cdots x_n^\infty) < d$ then we may apply Proposition 8.1.6 to find a monomial-free initial ideal of $\text{in}_\omega(I)$ with a larger dimensional homogeneity space. By Proposition 3.1.20 this initial ideal is also an initial ideal of I . Let us

call it $\text{in}_{\omega'}(I)$. If $\text{homog}(\text{in}_{\omega'}(I)) < d$ we may repeat the process until we find an initial ideal with a homogeneity space of the right dimension. \square

8.2 A proof of the equivalent theorem

Having proved the equivalence of the two theorems it remains to prove Theorem 8.1.1. For this we need the following variant of [42, Corollary 6.13].

Theorem 8.2.1 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a prime ideal, $\omega \in \mathbb{R}^n$ with $\text{in}_{\omega}(I)$ monomial-free and $P \in \text{minAss}(\text{in}_{\omega}(I))$ then we have $\dim(P) = \dim(I)$.*

We will present a proof of the theorem at the end of this section. Here is a corollary.

Corollary 8.2.2 *Let $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ be a prime ideal of dimension d and $\omega \in \mathbb{R}^n$. If $\text{in}_{\omega}(I)$ is monomial-free then $\text{in}_{\omega}(I)$ has an associated prime of dimension d which does not contain a monomial.*

Proof. Consider a minimal primary decomposition of $\text{in}_{\omega}(I) = \cap_i Q_i$. Suppose the corollary was wrong then all associated primes of dimension d would contain a monomial. Theorem 8.2.1 implies that all minimal associated primes contain a monomial. For this reason also all associated primes contain a monomial. Since the associated primes are the radicals of the Q_i 's, every Q_i contains a monomial and so does their intersection $\text{in}_{\omega}(I)$. This is a contradiction. \square

Proof of Theorem 8.1.1. We know that $\dim(\text{in}_{\omega}(I)) = \dim(I)$. The only problem might be if the dimension of $\text{in}_{\omega}(I)$ drops when we saturate with respect to $x_1 \cdots x_n$. When we saturate we are removing components in the coordinate hyperplanes of \mathbb{C}^n . Corollary 8.2.2 tells us that there is a d -dimensional component whose ideal does not contain a monomial and thus the component is not contained in a coordinate hyperplane and cannot disappear when we saturate. For this reason the saturation $(\text{in}_{\omega}(I) : x_1 \cdots x_n^{\infty})$ has dimension d . \square

For completeness we provide a proof of Theorem 8.2.1.

Proof of Theorem 8.2.1. With our knowledge of the Gröbner fan we may assume, without loss of generality, that $\omega \in \mathbb{Q}^n$. Let $N \in \mathbb{N}$ be an integer such that $N\omega \in \mathbb{Z}^n$. The ideal $I' = \langle I \rangle_{\mathbb{C}[t^{\pm \frac{1}{N}}, x_1, \dots, x_n]} \subseteq \mathbb{C}[t^{\pm \frac{1}{N}}, x_1, \dots, x_n]$ is a monomial-free prime ideal with $\dim(I') = \dim(I) + 1$. Consider the \mathbb{C} -algebra isomorphism

$$\begin{aligned} \Phi_{\omega} : \mathbb{C}[t^{\pm \frac{1}{N}}, x_1, \dots, x_n] &\rightarrow \mathbb{C}[t^{\pm \frac{1}{N}}, x_1, \dots, x_n] \\ x_i &\mapsto t^{-\omega_i} x_i \quad \text{and} \quad t \mapsto t. \end{aligned}$$

The image $\Phi_{\omega}(I')$ is a monomial-free prime ideal. Observe that $\text{t-in}_0(\Phi_{\omega}(I')) = \text{t-in}_{\omega}(I') = \text{in}_{\omega}(I)$. Let $J := \Phi_{\omega}(I') \cap \mathbb{C}[t^{\frac{1}{N}}, x_1, \dots, x_n]$. Then $\dim(J) = \dim(I) + 1$. Observe that

$$\text{t-in}_0(J) = (J + \langle t^{\frac{1}{N}} \rangle) \cap \mathbb{C}[x_1, \dots, x_n].$$

Let $P \in \min\text{Ass}(\text{in}_\omega(I)) = \min\text{Ass}((J + \langle t^{\frac{1}{N}} \rangle) \cap \mathbb{C}[x_1, \dots, x_n])$. The minimal associated primes of $J + \langle t^{\frac{1}{N}} \rangle$ and of $(J + \langle t^{\frac{1}{N}} \rangle) \cap \mathbb{C}[x_1, \dots, x_n]$ are in bijection where ideals correspond to ideals of the same dimension. Let $P' \in \min\text{Ass}(J + \langle t^{\frac{1}{N}} \rangle)$ be the ideal corresponding to P . We will apply Krull's principal ideal theorem [1, Corollary 11.17] to $\mathbb{C}[t^{\frac{1}{N}}, x_1, \dots, x_n]/J$. Notice that $t^{\frac{1}{N}}$ is not zero and not a zero-divisor in $\mathbb{C}[t^{\frac{1}{N}}, x_1, \dots, x_n]/J$ since J is monomial-free and prime. Furthermore, $t^{\frac{1}{N}}$ is not a unit in $\mathbb{C}[t^{\frac{1}{N}}, x_1, \dots, x_n]/J$ since in that case there would exist $f \in \mathbb{C}[t^{\frac{1}{N}}, x_1, \dots, x_n]$ such that $t^{\frac{1}{N}}f - 1 \in J$ — contradicting that $\text{t-in}_0(J) = \text{in}_\omega(I)$ is monomial-free. Hence Krull's principal ideal theorem implies that $\dim(P) = \dim(P') = \dim(J) - 1 = \dim(I)$. \square

Chapter 9

Software and examples

In this chapter we describe the software Gfan for computing Gröbner fans and tropical varieties and we present examples computed with the software.

9.1 Gfan

The software package Gfan ([33]) was developed by the author and contains implementations of all algorithms described in this thesis. A presentation of Gfan is given in [36]. Gfan is a C++ program ([52]) which uses the libraries cddlib ([20]) and GMP ([27]) for polyhedral computations and exact arithmetic, respectively.

We give an example of how the program can be used to compute the tropical variety of the Grassmann-Plücker ideal $I \subseteq \mathbb{Q}[a, \dots, j]$ in Example 6.2.3. The ideal generated by I in $\mathbb{C}[a, \dots, j]$ is prime and by the Bieri Groves Theorem 6.4.3 its tropical variety is a pure complex and thereby $\mathcal{T}(I)$ is also pure. Gfan is a collection of command line tools. The command `gfan_tropicalstartingcone` is an implementation of Algorithm 7.4.2. If we run it on the input

```
{  
bf-ah-ce,  
bg-ai-de,  
cg-aj-df,  
ci-bj-dh,  
fi-ej-gh  
}
```

which is a generating set for I we get the pair of Gröbner bases

```
{  
f*i-e*j,  
d*h-c*i,  
d*f+a*j,  
d*e+a*i,  
c*e+a*h}
```

```
{
f*i-g*h-e*j,
d*h-c*i+b*j,
d*f-c*g+a*j,
d*e-b*g+a*i,
c*e-b*f+a*h}
```

representing a maximal cone in $\mathcal{T}(I)$. Speyer's Theorem 7.3.2 states that $\mathcal{T}(I)$ is connected in codimension one. We wish to traverse $\mathcal{T}(I)$ by repeatedly applying Algorithm 7.4.1. This can be done with the command `gfan.tropicaltraverse`. The input is the pair above and the output is a list of pairs representing all the maximal cones in $\mathcal{T}(I)$. It is also possible to get a representation of the combinatorics of $\mathcal{T}(I)$. In Gfan version 0.3, which has not been released at the time of writing, the representation looks like this:

```
_application PolyhedralFan
_version 2.2
_type PolyhedralFan                PURE
                                     1

AMBIENT_DIM                          F_VECTOR
10                                    10 15

DIM                                    MAXIMAL_CONES
7                                     {0 4}
                                     {4 5}
                                     {1 4}
LINEALITY_DIM                          {2 5}
5                                     {3 5}
                                     {0 6}
RAYS                                    {1 7}
-1 0 0 0 0 0 0 0 0 0 0              {2 6}
0 -1 0 0 0 0 0 0 0 0 0              {6 7}
1 1 1 0 0 0 0 0 0 0 0              {3 7}
0 0 -1 0 0 0 0 0 0 0 0              {0 8}
-1 -1 0 0 -1 0 0 0 0 0 0            {1 9}
0 0 0 0 -1 0 0 0 0 0 0              {2 9}
1 1 1 0 1 1 0 0 0 0 0              {8 9}
1 0 0 0 1 1 0 0 0 0 0              {3 8}
-1 0 -1 0 0 -1 0 0 0 0 0
0 0 0 0 0 -1 0 0 0 0 0

N_RAYS                                  MULTIPLICITIES
10                                      1
                                     1
                                     1
LINEALITY_SPACE                         1
0 0 0 0 1 1 1 1 1 1 1              1
0 0 0 1 0 0 1 0 1 1 1              1
0 0 1 0 0 1 0 1 0 1 1              1
0 1 0 0 0 -1 -1 0 0 -1 1          1
1 0 0 0 0 0 0 -1 -1 -1 1          1
                                     1
ORTH_LINEALITY_SPACE                    1
0 0 0 0 0 1 -1 -1 1 0 1            1
0 0 0 0 1 0 -1 -1 0 1 1            1
0 0 1 -1 0 0 0 -1 1 0 1            1
0 1 0 -1 0 0 0 -1 0 1 1            1
1 0 0 -1 0 0 -1 -1 1 1 1          1
```


The above output is in a Polymake compatible format ([25]). This format is also used by the tropical software TrIm ([56]). The file describes a polyhedral fan. The lowest dimensional cone is called the *lineality space* which equals the homogeneity space in case of a tropical variety. Its dimension, a generating set and a generating set for its orthogonal complement are listed. The ten 6-dimensional rays are listed with a relative interior point for each. The rays are indexed by numbers $0, \dots, 9$. Each maximal cones is generated by a subset of these rays. The complete list of the 15 maximal cones is given. Listed is also the dimension of the ambient space, the dimension of the complex and the (truncated) f-vector.

Associated to each maximal cone is its *multiplicity* which we have not defined in this thesis. The multiplicity is a number in $\mathbb{N}_{>0}$ which depends on the saturated initial ideal of the cone. See [55, Definition 3.1] for a definition. Let ω be a relative interior point in a $d - 1$ dimensional cone of $\mathcal{T}(I)$. As we have seen the initial ideal $\text{in}_\omega(I)$ defines a tropical curve. The multiplicities of $\mathcal{T}(I)$ satisfy the *balancing condition* meaning that the rays of $\mathcal{T}(\text{in}_\omega(I))$ sum up to zero if weighted by their multiplicities. Here each ray is thought of as a certain suitable vector in the d -dimensional cone of $\mathcal{T}(\text{in}_\omega(I))$. See [55, Definition 3.3] for a precise definition of the balancing condition.

Most examples in this thesis were computed by Gfan. In particular Gfan drew Figure 1.1 and Figure 3.3. In the next section we present some more examples.

9.2 Computational results and examples

We list some families of ideals that the Gfan software has been applied to. Gröbner fans of these ideals have been computed for the parameters listed in Table 9.1. Results for the tropical varieties can be found in Table 9.2. These computations first appeared in [21] and [10] respectively. The coefficient field is always \mathbb{Q} . The columns of the tables are to be interpreted as follows. In each row, the first column contains the name of the ideal (to be explained below). The second column lists n , the number of variables in the ideal. The third column lists h , the dimension of the homogeneity space. The quantity “deg” is the lowest total degree of any reduced Gröbner basis of the ideal and “DEG” is the highest. The dimension of the ideal and tropical variety is denoted d . The f-vector of the complexes are also listed.

Example 9.2.1 Let $\text{Det}_{t,m,n}$ denote the ideal in the polynomial ring in mn variables generated by the $t \times t$ minors of the matrix:

$$\begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1n} \\ x_{21} & x_{22} & \cdots & x_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mn} \end{pmatrix}.$$

Example	n	h	deg	DEG	f -vector
Det _{3,3,4}	12	6	3	3	(1,12,66,204,342,288,96)
Det _{3,3,5}	15	7	3	3	(1,45,585,3390,10710,19890,21750,12960,3240)
Det _{3,4,4}	16	7	3	5	(1,?,?,?,?,?,?,163032)
Detsym _{3,4}	10	4	3	8	(1,518,5412,20505,36024,29808,9395)
Grass _{2,5}	10	5	2	3	(1,20,120,300,330,132)
Cyclic ₅	5	0	8	15	(1,?,?,?,?,55320)

Table 9.1: Statistics for the Gröbner fans computed using Gfan.

Example	n	h	d	f -vector
Det _{3,3,5}	15	7	12	(1,45,315,930,1260,630)
Det _{3,4,4}	16	7	12	(1,50,360,1128,1680,936)
Detsym _{3,4}	10	4	7	(1,20,75,75)
Detsym _{3,5}	15	5	9	(1,75,495,1155,855)
Grass _{2,5}	10	5	7	(1,10,15)
ComSymMat ₃	12	2	9	(1,66,705,3246,7932,10888,8184,2745)

Table 9.2: Statistics for the Tropical varieties computed using Gfan.

Example 9.2.2 Let $\text{Grass}_{d,n}$ denote the ideal in the polynomial ring in $\binom{n}{d}$ variables generated by the relations on the $d \times d$ minors of a $d \times n$ matrix.

Example 9.2.3 Let $\text{Detsym}_{t,n}$ denote the ideal in the polynomial ring in $\frac{n(n+1)}{2}$ variables generated by the $t \times t$ minors of a symmetric matrix of variables. For example, $\text{Detsym}_{3,4}$ is generated by the 3×3 minors of the following matrix:

$$\begin{pmatrix} a & b & c & d \\ b & e & f & g \\ c & f & h & i \\ d & g & i & j \end{pmatrix}.$$

Example 9.2.4 Let Cyclic_5 denote the ideal $\langle a + b + c + d + e, ab + bc + cd + de + ae, abc + bcd + cde + dea + eab, abcd + abce + abde + acde + bcde, abcde - 1 \rangle \subseteq k[a, b, c, d, e]$. In general, Cyclic_n stands for the generalization of this polynomial system to n variables [9]. These polynomial systems have become benchmarks for computer algebra packages and their lexicographic Gröbner bases are notoriously hard to compute.

Example 9.2.5 Let A and B be symmetric $n \times n$ matrices of variables. The ideal generated by the relations $AB - BA = 0$ is denoted ComSymMat_n . It is an ideal in the polynomial ring with $(n + 1)n$ variables.

Extracting the f -vector from the full-dimensional Gröbner cones produced in the enumeration process was the most time-consuming part of the computation of Gröbner fan examples these. In example $\text{Det}_{3,4,4}$ this extraction was not possible to complete within reasonable time with the current software package. For

this particular example the 163032 full-dimensional Gröbner cones were computed up to the action of a symmetry group of order 576. The full-dimensional cones come in 289 orbits. The computation of the full dimensional cones up to symmetry took 7 minutes on a 2.4 GHz Intel Pentium processor. Using reverse search without symmetry the same computation would take approximately 14 hours. The maximal cones of the tropical variety were traversed in two minutes exploiting symmetry.

The f-vector extraction routine in Gfan only works for complete fans. This is why the f-vector for the Cyclic₅ example is not shown.

Bibliography

- [1] M. F. Atiyah and I. G. Macdonald. *Introduction to commutative algebra*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969.
- [2] David Avis and Komei Fukuda. A basis enumeration algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete Computational Geometry*, 8:295–313, 1992.
- [3] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Appl. Math.*, 65(1-3):21–46, 1996. First International Colloquium on Graphs and Optimization (GOI), 1992 (Grimentz).
- [4] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, second edition, 2006.
- [5] David Bayer and Ian Morrison. Standard bases and geometric invariant theory. I. Initial ideals and state polytopes. *J. Symbolic Comput.*, 6(2-3):209–217, 1988. Computational aspects of commutative algebra.
- [6] Thomas Becker and Volker Weispfenning. *Gröbner bases*, volume 141 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1993. A computational approach to commutative algebra, In cooperation with Heinz Kredel.
- [7] George M. Bergman. The logarithmic limit-set of an algebraic variety. *Trans. Amer. Math. Soc.*, 157:459–469, 1971.
- [8] 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.
- [9] Göran Björck and Ralf Fröberg. A faster way to count the solutions of inhomogeneous systems of algebraic equations, with applications to cyclic n -roots. *J. Symb. Comput.*, 12(3/4):329–336, 1991.
- [10] T. Bogart, A. N. Jensen, D. Speyer, B. Sturmfels, and R. R. Thomas. Computing tropical varieties. *J. Symbolic Comput.*, 42(1-2):54–73, 2007, math.AG/0507563.

- [11] Tristram Bogart, Anders N. Jensen, and Rekha R. Thomas. The circuit ideal of a vector configuration. *J. Algebra*, 309(2):518–542, 2007, math.AC/0508628.
- [12] Bruno Buchberger. An algorithm for finding a basis for the residue class ring of a zero-dimensional polynomial ideal, 1965. PhD-thesis (German).
- [13] S. Collart, M. Kalkbrener, and D. Mall. Converting bases with the Gröbner walk. *J. Symbolic Comput.*, 24(3-4):465–469, 1997. Computational algebra and number theory (London, 1993).
- [14] David Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, second edition, 1997. An introduction to computational algebraic geometry and commutative algebra.
- [15] Mike Develin and Bernd Sturmfels. Erratum for: “Tropical convexity” [Doc. Math. **9** (2004), 1–27 (electronic); mr2054977]. *Doc. Math.*, 9:205–206 (electronic), 2004.
- [16] Jan Draisma. A tropical approach to secant dimensions, 2006, math.AG/0605345v3.
- [17] Manfred Einsiedler, Mikhail Kapranov, and Douglas Lind. Non-Archimedean amoebas and tropical varieties. *J. Reine Angew. Math.*, 601:139–157, 2006.
- [18] J. C. Faugère, P. Gianni, D. Lazard, and T. Mora. Efficient computation of zero-dimensional Gröbner bases by change of ordering. *J. Symbolic Comput.*, 16(4):329–344, 1993.
- [19] K. Fukuda, A. N. Jensen, N. Lauritzen, and R. Thomas. The generic Gröbner walk. *J. Symbolic Comput.*, 42(3):298–312, 2007, math.AC/0501345.
- [20] Komei Fukuda. *cddlib reference manual, cddlib Version 094b*. Swiss Federal Institute of Technology, Lausanne and Zürich, Switzerland, 2005. http://www.ifor.math.ethz.ch/~fukuda/cdd_home/cdd.html.
- [21] Komei Fukuda, Anders Jensen, and Rekha Thomas. Computing Gröbner fans. *Mathematics of Computation*, to appear, 2007, math.AC/0509544.
- [22] 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.
- [23] Tangan Gao and T. Y. Li. Mixed volume computation via linear programming. *Taiwanese J. Math.*, 4(4):599–619, 2000.
- [24] Andreas Gathmann. Tropical algebraic geometry. 2006, math.AG/0601322.

- [25] Evgenij Gawrilow and Michael Joswig. `polymake`: a framework for analyzing convex polytopes. In Gil Kalai and Günter M. Ziegler, editors, *Polytopes — Combinatorics and Computation*, pages 43–74. Birkhäuser, 2000.
- [26] Patrizia Gianni, Barry Trager, and Gail Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *J. Symbolic Comput.*, 6(2-3):149–167, 1988. Computational aspects of commutative algebra.
- [27] Torbjörn Granlund et al. GNU multiple precision arithmetic library 4.1.2, December 2002. <http://swox.com/gmp/>.
- [28] G.-M. Greuel, G. Pfister, and H. Schönemann. SINGULAR 2.0.5. A Computer Algebra System for Polynomial Computations, Centre for Computer Algebra, University of Kaiserslautern, 2004. <http://www.singular.uni-kl.de>.
- [29] Gert-Martin Greuel and Gerhard Pfister. *A Singular introduction to commutative algebra*. Springer-Verlag, Berlin, 2002. With contributions by Olaf Bachmann, Christoph Lossen and Hans Schönemann, With 1 CD-ROM (Windows, Macintosh, and UNIX).
- [30] Marshall Hampton and Richard Moeckel. Finiteness of relative equilibria of the four-body problem. *Invent. Math.*, 163(2):289–312, 2006.
- [31] Birkett Huber and Bernd Sturmfels. A polyhedral method for solving sparse polynomial systems. *Math. Comp.*, 64(212):1541–1555, 1995.
- [32] Birkett Huber and Rekha R. Thomas. Computing Gröbner fans of toric ideals. *Experimental Mathematics*, 9(3/4):321–331, 2000.
- [33] Anders Nedergaard Jensen. Gfan, a software system for Gröbner fans. Available at <http://home.imf.au.dk/ajensen/software/gfan/gfan.html>.
- [34] Anders Nedergaard Jensen. Computing Gröbner fans of toric ideals, 2002. Master thesis, University of Aarhus.
- [35] Anders Nedergaard Jensen. A Singular script for verifying a Gröbner fan computation., 2005. <http://home.imf.au.dk/ajensen/papers/singularscript.html>.
- [36] Anders Nedergaard Jensen. A presentation of the Gfan software. In Iglesias and Takayama, editors, *Mathematical Software - ICMS 2006*, pages 222–224. Springer, 2006.
- [37] Anders Nedergaard Jensen. A non-regular Gröbner fan. *Discrete Comput. Geom.*, 37(3):443–453, 2007, math.CO/0501352.
- [38] Michael Kalkbrener. On the complexity of Gröbner bases conversion. *J. Symbolic Comput.*, 28(1-2):265–273, 1999. Polynomial elimination—algorithms and applications.

- [39] Eric Katz. A tropical toolkit, 2006, math.AG/0610878v2.
- [40] Niels Lauritzen. *Concrete abstract algebra*. Cambridge University Press, Cambridge, 2003. From numbers to Gröbner bases.
- [41] Niels Lauritzen. Truncated Gröbner fans and lattice ideals. 2005, math.AG/0509247.
- [42] Hannah Markwig, Thomas Markwig, and Anders Jensen. An algorithm for lifting points in a tropical variety. 2007, math.AG/0705.2441.
- [43] Grigory Mikhalkin. Tropical geometry, 2007. Work in progress, <http://www.math.toronto.edu/mikha/TG-project.html>.
- [44] Teo Mora and Lorenzo Robbiano. The Gröbner fan of an ideal. *J. Symbolic Comput.*, 6(2-3):183–208, 1988. Computational aspects of commutative algebra.
- [45] T. S. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall. The double description method. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 51–73. Princeton University Press, Princeton, N. J., 1953.
- [46] Shmuel Onn and Bernd Sturmfels. Cutting corners. *Advances in Applied Mathematics*, 23:29–48, 1999.
- [47] Lior Pachter and Bernd Sturmfels. *Algebraic Statistics for Computational Biology*. Cambridge Univ. Press, New York, 2005.
- [48] Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald. First steps in tropical geometry. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 2005, math.AG/0306366.
- [49] Lorenzo Robbiano. Term orderings on the polynomial ring. In *EUROCAL '85, Vol. 2 (Linz, 1985)*, volume 204 of *Lecture Notes in Comput. Sci.*, pages 513–517. Springer, Berlin, 1985.
- [50] Mutsumi Saito, Bernd Sturmfels, and Nobuki Takayama. *Gröbner deformations of hypergeometric differential equations*, volume 6 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2000.
- [51] David Speyer and Bernd Sturmfels. The tropical Grassmannian. *Adv. Geom.*, 4(3):389–411, 2004, math.AG/0304218.
- [52] Bjarne Stoustrup. *The C++ Programming Language*. Addison-Wesley, third edition, 1997.
- [53] Bernd Sturmfels. *Gröbner bases and convex polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1996.

- [54] Bernd Sturmfels. *Solving systems of polynomial equations*, volume 97 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2002.
- [55] Bernd Sturmfels and Jenia Tevelev. Elimination theory for tropical varieties. 2007, math.AG/0704.3471v1.
- [56] Bernd Sturmfels and Josephine Yu. Tropical implicitization and mixed fiber polytopes, 2007, cs.SC/0706.0564v1.
- [57] Thorsten Theobald. On the frontiers of polynomial computations in tropical geometry. *J. Symbolic Comput.*, 41(12):1360–1375, 2006, math.CO/0411012.
- [58] Quoc-Nam Tran. A fast algorithm for Gröbner basis conversion and its applications. *J. Symbolic Comput.*, 30(4):451–467, 2000.
- [59] Magnus Dehli Vigeland. The group law on a tropical elliptic curve. 2004, math.AG/0411485.
- [60] Christophe Weibel and Komei Fukuda. Computing faces up to k dimensions of a minkowski sum of polytopes. In *CCCG*, pages 256–259, 2005.
- [61] Günter Ziegler. *Lectures on Polytopes*, volume 152 of *GTM*. Springer-Verlag, 1994.

Index

- Ass(I), 20
- face $_{\omega}(P)$, 15
- $\mathcal{G}_{\prec}(I)$, 18
- homog(I), 27
- in $_{\omega}(f)$, 18
- in $_{\omega}(I)$, 18
- in $_{\prec}(f)$, 18
- in $_{\prec}(I)$, 18
- int(P), 17
- minAss(I), 20
- New(f), 18
- ω -degree, 18
- rel int(P), 17
- supp(\mathcal{C}), 16
- $\mathcal{T}(f)$, 67
- $\mathcal{T}(I)$, 67
- $\mathcal{T}(P)$, 80

- associated prime, 20

- Buchberger's algorithm, 18

- common refinement, 17
- complete fan, 16
- cone, 15

- degree, 18, 39
- dimension, 15, 21
- division algorithm, 19

- embedded prime, 21
- extended Gröbner fan, 56

- f-vector, 16
- face, 15
- facet, 15
- fan, 16
- Farkas' lemma, 55
- fiber, 41
- flip, 44, 46
- flippable, 43

- flow, 53

- generic Gröbner walk, 47
- Gfan, 8, 57, 95
- grading, 39
- graph of fan, 36
- Gröbner basis, 18
- Gröbner cone, 29
- Gröbner fan, 29
- Gröbner region, 29
- Gröbner walk, 46

- Hilbert function, 39
- homogeneity space, 27
- homogeneous, 19, 39
- homogenization, 56

- initial form, 18
- initial ideal, 18
- initial term, 18
- irreducible variety, 20

- Krull dimension, 21

- lift, 45
- local change, 44

- marked, 18
- marked Gröbner basis, 18
- matrix term order, 20
- minimal associated prime, 20
- minimal Gröbner basis, 18
- minimal primary decomposition, 21
- minimal prime, 21
- Minkowski sum, 17
- mixed face, 80
- monomial-free, 64

- Newton polytope, 18
- Newton-Puiseux algorithm, 62
- non-regular, 51

- normal cone, 16
- normal fan, 16
- normal form, 18

- polyhedral complex, 15
- polyhedral cone, 15
- polyhedral fan, 16
- polyhedron, 15
- polytope, 15
- primary decomposition, 21
- primary ideal, 21
- pure fan, 16

- quotient graph, 50

- radical, 21
- reduced Gröbner basis, 18
- refinement, 17
- regular, 51
- relative interior, 17
- restricted Gröbner fan, 29
- reverse lexicographic, 20
- reverse search, 36
- reverse search property, 37
- ridge path, 84

- S-polynomial, 19
- saturated, 23
- saturated initial ideal, 73
- saturation, 23
- search edge, 36
- search tree, 36
- shadow, 9, 72
- Singular, 21, 56
- Speyer's theorem, 84
- standard monomial, 18
- state polytope, 38
- support, 16
- symmetry, 49
- symmetry group, 50

- t - ω -degree, 64
- t -initial form, 64
- t -initial ideal, 64
- term, 17
- term order, 17
- toric ideal, 23
- tropical basis, 77
- tropical curve, 82
- tropical curve, non-constant coefficients, 60
- tropical hypersurface, 67
- tropical polynomial, 59, 60
- tropical prevariety, 77
- tropical semi-ring, 59
- tropical variety, 64, 65, 67
- tropicalization, 64

- witness, 77

- zero set, 60