UNIVERSITY OF AARHUS
DEPARTMENT OF MATHEMATICS

ISSN: 1397–4076

COMPUTING GRÖBNER FANS

by Komei Fukuda, Anders N. Jensen and Rekha R. Thomas

Preprint Series No.: 15

September 2005

2005/09/30

Ny Munkegade, Bldg. 530
DK-8000 Aarhus C, Denmark

http://www.imf.au.dk
institut@imf.au.dk

Computing Gröbner fans

Komei Fukuda[*], Anders N. Jensen[†], and Rekha R. Thomas[‡]

September 29, 2005

Abstract

This paper presents algorithms for computing the Gröbner fan of an arbitrary polynomial ideal. The computation involves enumeration of all reduced Gröbner bases of the ideal. Our algorithms are based on a uniform definition of the Gröbner fan that applies to both homogeneous and non-homogeneous ideals and a proof that this object is a polyhedral complex. We show that the cells of a Gröbner fan can easily be oriented acyclically and with a unique sink, allowing their enumeration by the memory-less reverse search procedure. The significance of this follows from the fact that Gröbner fans are not always normal fans of polyhedra in which case reverse search applies automatically. Computational results using our implementation of these algorithms in the software package Gfan are included.

1 Introduction

The *Gröbner fan* of an ideal $I \subseteq k[x_1, \ldots, x_n]$ was defined by Mora and Robbiano in [16]. It is a fan of polyhedral cones indexing *initial ideals* of I. The full-dimensional cones are in bijection with the distinct monomial initial ideals with respect to term orders or equivalently, the reduced Gröbner bases of the ideal. In this paper we will describe algorithms for computing Gröbner fans of arbitrary polynomial ideals allowing us to study their structure in detail. Our algorithms are implemented in the software package Gfan [12].

The computation of the Gröbner fan of I in terms of reduced Gröbner bases yields *a universal Gröbner basis* of I, a set of polynomials which is a Gröbner basis of I with respect to every term order. The Gröbner fan also plays an important role in *Gröbner basis conversion* [5] and the emerging field of *tropical mathematics* as it contains the *tropical variety* of I as a subfan [4]. Many of the well-known theoretical applications of Gröbner bases rely on the existence of a Gröbner basis or

[*]Research supported by the Swiss National Science Foundation Project 200021-105202, "Polytopes, Matroids and Polynomial Systems".

[†]Research partially supported by the Faculty of Science, University of Aarhus, Danish Research Training Council (Forskeruddannelsesrådet, FUR) , Institute for Operations Research ETH, the Swiss National Science Foundation Project 200021-105202, grants DMS 0222452 and DMS 0100141 of the U.S. National Science Foundation and the American Institute of Mathematics.

[‡]Research supported by grant DMS 0401047 of the U.S. National Science Foundation.

initial ideal for an ideal with prescribed properties such as a particular complexity (a specified degree or squarefree-ness) or homological properties (Cohen-Macaulay, Gorenstein etc). The (even partial) computation of the Gröbner fan makes such experimentations possible. No software package for the computation of Gröbner fans is available at present.

In the literature a distinction is often made between the case of I being homogeneous where the Gröbner fan is a complete fan in $\mathbb{R}^n$, and the case of I being non-homogeneous, where the *restricted* Gröbner fan with support equal to $\mathbb{R}^n_{\geq 0}$ is considered. In this paper we avoid this distinction by giving a new uniform definition of the Gröbner fan and we prove that this actually defines a fan in the sense of a *polyhedral complex* — a proof that was left out in [16] but was proven for the special case of homogeneous ideals in [18]. See Section 2.

In [18, Algorithm 3.2 and 3.6] and [16, Section 6] methods for computing the Gröbner fan of a polynomial ideal were given. In this paper we shall study [18, Algorithm 3.6] in detail. This algorithm traverses the maximal cones of the Gröbner fan. In the special case of *toric* ideals the traversal algorithm was already studied and implemented in [11]. The traversal is graph-like — given a maximal cone we need to be able to find its facets and we need to be able to walk through a facet to the neighboring maximal cone. Algorithms for performing these local computations are discussed in Section 4. These amount to solving linear programming problems and using the local basis change procedure due to [5]. We explain how to apply these methods to our case.

The Gröbner fan of a homogeneous ideal I is known to be the normal fan of a polytope, the *state polytope* of I ([18, Theorem 2.5]). In this homogeneous case traversal of the maximal cones in the Gröbner fan by walking through facets is equivalent to traversal of the edge graph of the state polytope. In [1] the memoryless *reverse search* procedure for traversing the edge graph of a polytope was given. This procedure easily applies to Gröbner fans of homogeneous ideals. However, the question is what happens if the ideal is not homogeneous. In [13] a *non-regular* Gröbner fan was presented — a fan that is not the normal fan of any polyhedron. In light of this example it is not clear that the reverse search technique applies to Gröbner fans in general. In Section 3 we prove that all Gröbner fans have what we shall call the *reverse search property*, allowing them to be traversed efficiently.

Gröbner fans are often computed for ideals that possess a great deal of symmetry. In Section 4.4 we describe how to take advantage of symmetry in the computations. The methods used here are similar to those in Rambau's software package TOPCOM [17] for traversing the secondary fan of a point configuration up to symmetry.

In Section 5 we discuss the complexity of our enumeration algorithm and in Section 6 we present several examples of Gröbner fans computed using Gfan. This software package uses the GNU multi-precision library [10] for exact arithmetics and Cddlib [7] for solving linear programming problems.

2 The Gröbner fan of a polynomial ideal

Let $R = k[x_1, \ldots, x_n]$ be the polynomial ring in n variables over a field k and let $I \subseteq R$ be an ideal. The *Gröbner fan* and the *restricted Gröbner fan* of I are n-

dimensional polyhedral fans defined in [16]. 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* of a polyhedron 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 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.

The *support* of $\mathcal{C}$ is the union of its members. A polyhedral complex is a *fan* if it only consists of cones. A fan is *pure* if all its maximal cones have the same dimension.

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

Definition 2.2 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.

If the support of a fan is $\mathbb{R}^n$, the fan is said to be *complete*. It is clear that the normal fan of a polytope is complete. Not all fans arise as the normal fan of a polyhedron [9, page 25].

For $\alpha \in \mathbb{N}^n$ we use the notation $\mathbf{x}^\alpha := x_1^{\alpha_1} \dots x_n^{\alpha_n}$ for a monomial in R. By a *term order* on R we mean a total ordering on all monomials in R such that:

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

By a *term* we mean a monomial together with its coefficient. Term orders are used for ordering terms, ignoring the coefficients. For a vector $\omega \in \mathbb{R}^n_{\geq 0}$ and a term order $\prec$ we define the new term order $\prec_\omega$ as follows:

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

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 whose exponents maximize $\langle \omega, \cdot \rangle$. The polynomial f is *ω-homogeneous* if $\text{in}_\omega(f) = f$. The *ω-degree* of a term $c\mathbf{x}^\alpha$ is $\langle \omega, \alpha \rangle$ and the *ω-degree* of a non-zero polynomial f is the

common ω-degree of the terms of $\mathrm{in}_\omega(f)$. The *initial ideals* of an ideal I with respect to $\prec$ and ω are defined as

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

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

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

Definition 2.3 Let $I \subseteq R$ be an ideal and $\prec$ a term order on R. A generating set $\mathcal{G} = \{g_1, \ldots, g_m\}$ for I is called a *Gröbner basis* for I with respect to $\prec$ if

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

The Gröbner basis $\mathcal{G}$ is *minimal* if $\{\mathrm{in}_\prec(g_1), \ldots, \mathrm{in}_\prec(g_m)\}$ generates $\mathrm{in}_\prec(I)$ 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 $\mathrm{in}_\prec(I)$.

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.

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 $\mathcal{G}_\prec(I)$. For two term orders $\prec$ and $\prec'$, if $\mathrm{in}_\prec(I) = \mathrm{in}_{\prec'}(I)$ then $\mathcal{G}_\prec(I) = \mathcal{G}_{\prec'}(I)$. Conversely, given a marked Gröbner basis $\mathcal{G}_\prec(I)$, $\mathrm{in}_\prec(I)$ can be easily read off.

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

$$u \sim v \iff \mathrm{in}_u(I) = \mathrm{in}_v(I). \tag{1}$$

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

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

and

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

Remark 2.4 It is well known that for a fixed ideal I there are only finitely many sets $C_\prec(I)$ and they cover $\mathbb{R}^n_{\geq 0}$, see [16]. Secondly, every initial ideal $\mathrm{in}_\prec(I)$ is of the form $\mathrm{in}_\omega(I)$ for some $\omega \in \mathbb{R}^n_{>0}$, see [18, Proposition 1.11]. 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 2.5 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 $\mathrm{in}_u(I) = \mathrm{in}_v(I) = \langle 1 \rangle$ but $\mathrm{in}_{\frac{1}{2}(u+v)}(I) = \langle x, y \rangle$.

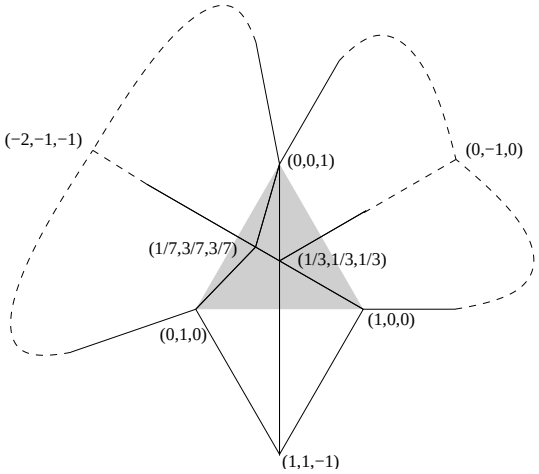

Figure 1: The Gröbner fan of the ideal in Example 2.7 has 7 three-dimensional, 14 two-dimensional and 8 one-dimensional cones. The intersections of the two-dimensional cones with the hyperplane $x+y+z=1$ are drawn as lines. The dotted part of the figure shows the combinatorial structure outside the hyperplane. The gray triangle indicates the positive orthant.

Proposition 2.6 *Let $\prec$ be a term order and $v \in C_\prec(I)$. For $u \in \mathbb{R}^n$*

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

This proposition is a little more general than Proposition 2.3 in [18] as it allows the vectors u and v to have negative components. A proof is given in the next section. For fixed $\prec$ and v as in Proposition 2.6, 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 $\mathrm{in}_u(g) = \mathrm{in}_v(g)$ which is equivalent to having u satisfy a set of linear equations and strict linear inequalities, see Example 2.7. The closure is obtained by making the strict inequalities non-strict. Under the assumptions of Proposition 2.6 we may write this in the following way:

$$u \in C_v(I) \iff \forall g \in \mathcal{G}_\prec(I),\ \mathrm{in}_v(\mathrm{in}_u(g)) = \mathrm{in}_v(g). \tag{2}$$

As we saw in Example 2.5, 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. In this case, there must exist a vector $p \in \mathbb{R}^n_{>0}$ in the interior of $C_v(I)$. Then $\mathrm{in}_p(I) = \mathrm{in}_v(I)$ and, by Lemma 2.15, $p \in C_{\prec_p}(I)$ for any $\prec$. Hence the equivalence class of v is of the form required in Proposition 2.6.

Example 2.7 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 $\mathcal{G}_\prec(I) = \{\underline{y^2}+x-x^3y-x^4, \underline{z}+y+x\}$. If $v = (1,4,5)$ then $\mathrm{in}_v(I) = \mathrm{in}_\prec(I) = \langle y^2, z\rangle$ and $C_v(I) = C_\prec(I)$. By Proposition 2.6, $\mathrm{in}_u(I) = \mathrm{in}_v(I)$ if and only if the following two equations are satisfied:

$$\mathrm{in}_u(z+y+x) = z \ (\Leftrightarrow\ u_z > \max\{u_x, u_y\}),$$

and

$$\mathrm{in}_u(y^2 + x - x^3y - x^4) = y^2 \ (\Leftrightarrow \ 2u_y > \max\{u_x, 3u_x + u_y, 4u_x\}).$$

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. Since $(-2,-1,0)$ is in $C_\prec(I)$ a description of vectors u in $C_{(-2,-1,0)}(I)$ is given by:

$$\mathrm{in}_{(-2,-1,0)}(\mathrm{in}_u(z + y + x)) = z \ (\Leftrightarrow \ u_z \geq \max\{u_x, u_y\}),$$

and

$$\mathrm{in}_{(-2,-1,0)}(\mathrm{in}_u(y^2 + x - x^3y - x^4)) = y^2 + x \ (\Leftrightarrow \ 2u_y = u_x \geq \max\{3u_x + u_y, 4u_x\}).$$

Definition 2.8 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.

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 2.19). The support of the Gröbner fan of I is called the *Gröbner region* of I. Recall that 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. We define the *restricted* Gröbner fan of an ideal to be the common refinement of the Gröbner fan and the faces of the non-negative orthant. The support of the restricted Gröbner fan is $\mathbb{R}^n_{\geq 0}$. The *Newton polytope* of a polynomial is the convex hull of its exponent vectors.

Example 2.9 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 2. 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 equivalence classes of the equivalence relation (1).

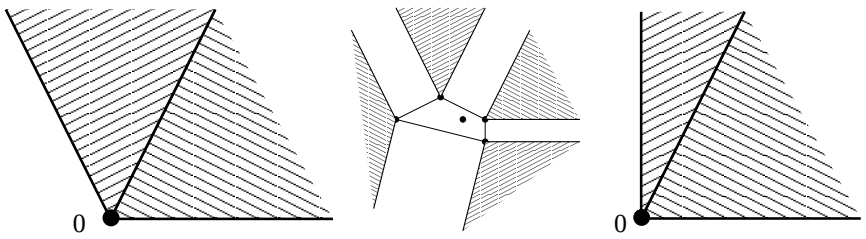

Figure 2: The Gröbner fan of the ideal in Example 2.9 is shown on the left. The restricted Gröbner fan is on the right. In the middle the Newton polytope of the generator is drawn with the shape of its normal fan indicated.

2.1 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 $\mathbb{R}^n_{\geq 0}$. In [16] 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 [18]. 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 that the Gröbner fan of a homogeneous ideal is a polyhedral complex [18]. 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 2.6 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 2.16) and, finally, that the intersection of two cones in the fan is a face of both (Proposition 2.18).

To prove Proposition 2.6 we start by proving a similar statement for the equivalence classes arising from initial ideals with respect to term orders.

Lemma 2.10 *Let $\prec$ be a term order. For $v \in \mathbb{R}^n$,*

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

Proof. $\Rightarrow$: Let $g \in \mathcal{G}_\prec(I)$. Since $\mathcal{G}_\prec(I)$ is reduced, only one term from g, $\mathrm{in}_\prec(g)$, can be in $\mathrm{in}_\prec(I) = \mathrm{in}_v(I)$. The initial ideal $\mathrm{in}_v(I)$ is a monomial ideal, implying that all terms of an element in the ideal must be in the ideal too. Hence, the initial form $\mathrm{in}_v(g) \in \mathrm{in}_v(I)$ has to be equal to $\mathrm{in}_\prec(g)$.
$\Leftarrow$: We must show that $\mathrm{in}_v(I) = \mathrm{in}_\prec(I)$ where $\mathrm{in}_\prec(I) = \langle \mathrm{in}_\prec(g) \rangle_{g \in \mathcal{G}_\prec(I)}$. The "$\supseteq$" inclusion is clear since $\mathrm{in}_\prec(g) = \mathrm{in}_v(g) \in \mathrm{in}_v(I)$ for all $g \in \mathcal{G}_\prec(I)$.

To prove the "$\subseteq$" inclusion, since $\mathrm{in}_v(I) = \langle \mathrm{in}_v(f),\ f \in I \rangle$, it suffices to show that $\mathrm{in}_v(f) \in \mathrm{in}_\prec(I)$ for all $f \in I$. Pick $f \in I$ and reduce it to zero using the division algorithm (e.g. [6, Chapter 2]) with $\mathcal{G}_\prec(I)$ and $\prec$. We may write

$$f = m_1 g_{i_1} + \cdots + m_r g_{i_r} \tag{3}$$

where m_j is a monomial and g_{i_j} is an element from $\mathcal{G}_\prec(I)$. The division algorithm guarantees that $\mathrm{in}_\prec(f) \geq m_j \mathrm{in}_\prec(g_{i_j})$ 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) has v-degree less than or equal to the v-degree of the left hand side. Consequently,

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

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

By Lemma 2.10 the equivalence class of $\mathrm{in}_\prec(I)$ is open. Since $\mathrm{in}_\prec(I)$ is of the form $\mathrm{in}_v(I)$ for some v (see Remark 2.4), the equivalence class of $\mathrm{in}_\prec(I)$ 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 2.11 *Let $\prec$ be a term order and $v \in \mathbb{R}^n$. Then*

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

Proof. Lemma 2.10 tells us that v lies in the interior of $C_\prec(I)$ if and only if $\mathrm{in}_v(g) = \mathrm{in}_\prec(g)$ 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 $\mathrm{in}_\prec(\mathrm{in}_v(g)) = \mathrm{in}_\prec(g)$ for all g in $\mathcal{G}_\prec(I)$. □

Lemma 2.12 *A polynomial $f \in \mathrm{in}_v(I)$ can be written in the form $f = \sum_i \mathrm{in}_v(c_i)$ where $c_i \in I$ and all summands in the sum have different v-degrees.*

Proof. The initial ideal $\mathrm{in}_v(I)$ is generated by v-homogeneous polynomials, implying that all v-homogeneous components of f are in $\mathrm{in}_v(I)$. 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 $\mathrm{in}_v(a_1) + \cdots + \mathrm{in}_v(a_s)$ for some polynomials $a_1, \ldots, a_s$ in I. Since h is v-homogeneous we can rewrite h as the sum $\sum_{j\in J} \mathrm{in}_v(a_j)$ of forms having the same v-degree as h. We pull out the initial form and get $h = \mathrm{in}_v(\sum_{j\in J} a_j)$. □

Lemma 2.13 *Let $\prec$ be a term order. If $v \in C_\prec(I)$ then $\mathrm{in}_\prec(\mathrm{in}_v(I)) = \mathrm{in}_\prec(I)$.*

Proof. Let $g \in \mathcal{G}_\prec(I)$. Since $v \in C_\prec(I)$, by Corollary 2.11, $\mathrm{in}_\prec(g) = \mathrm{in}_\prec(\mathrm{in}_v(g))$ and hence $\mathrm{in}_\prec(I) = \langle \mathrm{in}_\prec(g) \rangle_{g\in\mathcal{G}_\prec(I)} \subseteq \mathrm{in}_\prec(\mathrm{in}_v(I))$.

We now prove that $\mathrm{in}_\prec(\mathrm{in}_v(I)) \subseteq \mathrm{in}_\prec(I)$. Notice that $\mathrm{in}_\prec(\mathrm{in}_v(I))$ is generated by initial terms of elements $f \in \mathrm{in}_v(I)\backslash\{0\}$ with respect to $\prec$. Suppose $f \in \mathrm{in}_v(I)\backslash\{0\}$. It suffices to show that $\mathrm{in}_\prec(f) \in \mathrm{in}_\prec(I)$. Using Lemma 2.12 we may write $f = \sum_{i=1}^s \mathrm{in}_v(c_i)$ where $c_1, \ldots, c_s \in I$ and $\mathrm{in}_v(c_1), \ldots, \mathrm{in}_v(c_s)$ are v-homogeneous each with distinct degree, so that no cancellations occur. Consequently $\mathrm{in}_\prec(f)$ equals $\mathrm{in}_\prec(\mathrm{in}_v(c_j))$ for some j. We wish to prove that $\mathrm{in}_\prec(\mathrm{in}_v(c_j)) \in \mathrm{in}_\prec(I)$. 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 $m_1, \ldots, m_r$ are monomials and $g_{i_1}, \ldots, g_{i_r}$ 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 2.11. 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

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

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

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

Corollary 2.14 *Let $\prec$ be a term order. If $v \in C_\prec(I)$ then*

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

Proof. By Corollary 2.11, $\langle \mathrm{in}_\prec(\mathrm{in}_v(g)) \rangle_{g \in \mathcal{G}_\prec(I)} = \langle \mathrm{in}_\prec(g) \rangle_{g \in \mathcal{G}_\prec(I)} = \mathrm{in}_\prec(I)$. By Lemma 2.13, $\mathrm{in}_\prec(I)$ equals $\mathrm{in}_\prec(\mathrm{in}_v(I))$. Thus $\mathrm{in}_\prec(\mathrm{in}_v(I)) = \langle \mathrm{in}_\prec(\mathrm{in}_v(g)) \rangle_{g \in \mathcal{G}_\prec(I)}$. This proves that $\{\mathrm{in}_v(g)\}_{g \in \mathcal{G}_\prec(I)}$ is a Gröbner basis of $\mathrm{in}_v(I)$ 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 2.6 which claimed that given $v \in C_\prec(I)$ and $u \in \mathbb{R}^n$, $\mathrm{in}_u(I) = \mathrm{in}_v(I) \iff \forall g \in \mathcal{G}_\prec(I)$, $\mathrm{in}_u(g) = \mathrm{in}_v(g)$.

Proof. $\Leftarrow$: Since $\mathrm{in}_u(g) = \mathrm{in}_v(g)$ for all $g \in \mathcal{G}_\prec(I)$, we get that $\mathrm{in}_\prec(\mathrm{in}_u(g)) = \mathrm{in}_\prec(\mathrm{in}_v(g))$ for all $g \in \mathcal{G}_\prec(I)$. Since $v \in C_\prec(I)$, by Corollary 2.11, $\mathrm{in}_\prec(g) = \mathrm{in}_\prec(\mathrm{in}_v(g))$ for all $g \in \mathcal{G}_\prec(I)$ and hence $\mathrm{in}_\prec(g) = \mathrm{in}_\prec(\mathrm{in}_u(g))$ for all $g \in \mathcal{G}_\prec(I)$ and $u \in C_\prec(I)$ by Corollary 2.11. The Gröbner basis $\mathcal{G}_\prec(\mathrm{in}_u(I))$ is then $\{\mathrm{in}_u(g)\}_{g \in \mathcal{G}_\prec(I)}$ by Corollary 2.14. We get the same Gröbner basis for $\mathrm{in}_v(I)$. Hence, $\mathrm{in}_u(I) = \mathrm{in}_v(I)$.
$\Rightarrow$: Let $g \in \mathcal{G}_\prec(I)$. We need to show that $\mathrm{in}_u(g) = \mathrm{in}_v(g)$. Since the basis is reduced, only one term of g, namely $\mathrm{in}_\prec(g)$, is in $\mathrm{in}_\prec(I)$. We start by proving that the term $\mathrm{in}_\prec(g)$ is a term in $\mathrm{in}_v(g)$ and a term in $\mathrm{in}_u(g)$. For $\mathrm{in}_v(g)$ we apply Corollary 2.11 which says $\mathrm{in}_\prec(g) = \mathrm{in}_\prec(\mathrm{in}_v(g))$. For $\mathrm{in}_u(g)$ we apply Lemma 2.13 and get $\mathrm{in}_\prec(\mathrm{in}_u(g)) \in \mathrm{in}_\prec(\mathrm{in}_u(I)) = \mathrm{in}_\prec(\mathrm{in}_v(I)) = \mathrm{in}_\prec(I)$. Only one term of g is in $\mathrm{in}_\prec(I)$, so $\mathrm{in}_\prec(\mathrm{in}_u(g)) = \mathrm{in}_\prec(g)$. If the difference $\mathrm{in}_u(g) - \mathrm{in}_v(g)$, belonging to $\mathrm{in}_u(I) = \mathrm{in}_v(I)$, is non-zero we immediately reach a contradiction since the difference contains no terms from $\mathrm{in}_\prec(I) = \mathrm{in}_\prec(\mathrm{in}_v(I))$. $\square$

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 the argument following Proposition 2.6 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 2.15 *Let $\prec$ be a term order. If $v \in \mathbb{R}^n_{\geq 0}$ then $v \in C_{\prec_v}(I)$.*

Proof. This follows from Corollary 2.11 since $\mathrm{in}_{\prec_v}(\mathrm{in}_v(g)) = \mathrm{in}_{\prec_v}(g)$ for all $g \in \mathcal{G}_{\prec_v}(I)$. $\square$

Proposition 2.16 *The relative interior of a cone in the Gröbner fan is an equivalence class (with respect to $u \sim u' \Leftrightarrow \mathrm{in}_u(I) = \mathrm{in}_{u'}(I)$).*

Proof. By definition every cone in the fan is the face of the closure of an equivalence class for a positive vector $v \in \mathbb{R}^n_{>0}$. Let $\prec'$ be an arbitrary term order and define $\prec$ as $\prec'_v$. According to Lemma 2.15 the vector v belongs to $C_\prec(I)$. Notice that by (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 $\mathrm{in}_\prec(\mathrm{in}_u(g)) = \mathrm{in}_\prec(\mathrm{in}_v(\mathrm{in}_u(g))) = \mathrm{in}_\prec(\mathrm{in}_v(g)) = \mathrm{in}_\prec(g)$ of Corollary 2.11 is satisfied. By (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 2.6 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. □

It remains to be shown that the intersection of two cones in the Gröbner fan is a face of both cones (Proposition 2.18). We need a few observations.

Corollary 2.17 *Let C be a cone in the Gröbner fan. If $v \in C$ then for $u \in \mathbb{R}^n$,*

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

Proof. The vector v is in the relative interior of some face of C. This face is also in the Gröbner fan. By Proposition 2.16 u is in the relative interior of the same face and, consequently, also in C. □

By Remark 2.4 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 2.6 and Proposition 2.16.

Proposition 2.18 *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 2.17, C_1 and C_2 are unions of equivalence classes. Further, if $v \in C_1 \cap C_2$, then again by Corollary 2.17, 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 2.16 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 $\omega \in C_1 \cap C_2 \backslash \overline{E}$. Then $\text{conv}(\overline{E}, \omega) \backslash \overline{E}$ is contained in $C_1 \cap C_2$ and has dimension at least the dimension of E. This is a contradiction since $\text{conv}(\overline{E}, \omega) \backslash \overline{E}$ cannot be covered by finitely many lower dimensional equivalence classes. This completes the proof. □

Theorem 2.19 *The Gröbner fan is a polyhedral complex of cones and hence a fan.*

Proof. We already argued using Proposition 2.6 and Lemma 2.15 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 2.18. □

3 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* [1] can be used for traversing the graph of a Gröbner fan. This follows from the main theorem, Theorem 3.6, which says that the graph of a Gröbner fan can be oriented easily without cycles and with a unique sink. In Definition 3.4 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 proposition which we shall not prove.

Proposition 3.1 *Let $G = (V, E)$ be an oriented graph without cycles and with a unique sink s. If for every vertex $v \in V \backslash \{s\}$ some outgoing* search edge *$e_v = (v, \cdot)$ is chosen then the set of chosen edges is a spanning tree for G.*

The spanning tree in Proposition 3.1 is referred to as the *search tree*. The proposition implies that the graph is connected.

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 *Let $G = (V, E)$ be the oriented graph of Proposition 3.1 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*.

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

Example 3.3 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 $\mathbf{x}^p \prec \mathbf{x}^q$.

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 2.6 and 2.16 this is no restriction for Gröbner fans.

Definition 3.4 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 $C_1 \backslash F$ having negative inner product with $q - p$ and all points in $C_2 \backslash F$ having positive inner product with $q - p$. We orient the edge in direction from C_1 to C_2 if and only if $\mathbf{x}^p \prec \mathbf{x}^q$.

Note that the orientation of an edge in Definition 3.4 does not depend on the particular choice of p and q. Note also that for normal fans of polytopes this orientation agrees with the orientation of the edge graphs of the polytopes in Example 3.3. Not every fan has the reverse search property:

Example 3.5 Figure 3 shows a fan with support $\mathbb{R}^3_{\geq 0}$ 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 3.3 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 [18] for a proof. (We should mention that in [15] it was proven that this is only true if we use the state polytope definition in [18] and not true with the original definition in [2].) 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 $\mathbb{R}^n_{\geq 0}$ and being the normal fan of a polyhedron. Since the restricted Gröbner fan of any 0-dimensional or principal ideal satisfies these conditions it is clear that these fans have the reverse search property.

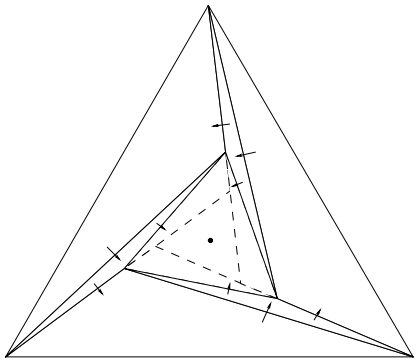

Figure 3: A fan not having the reverse search property, see Example 3.5.

It is shown in [13] that this line of reasoning cannot be applied to Gröbner fans in general. In particular, 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 3.6 *The Gröbner fan of any ideal $I \subseteq R$ has the reverse search property.*

The proof is given in the next section. In Section 4 we will argue that all parts of Algorithm 3.2 (finding adjacent edges, finding adjacent vertices and finding search edges) can be implemented efficiently for Gröbner fans.

3.1 Proof: The Gröbner fan has the reverse search property

In this section we prove Theorem 3.6. We start by recalling how the polynomial ring can be graded by semigroups. This leads to a more general notion of homogeneous ideals.

Definition 3.7 By a *grading* on $R = k[x_1, \ldots, x_n]$ we mean a pair $(A, \mathcal{A})$ consisting of an abelian semigroup A and a semigroup homomorphism:

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

such that $\mathcal{A}^{-1}(a)$ is finite for all $a \in A$. The $\mathcal{A}$-*degree* of a term $c\mathbf{x}^b$ 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 2.12 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*:

$$H_{I,\mathcal{A}} : A \to \mathbb{N} \tag{4}$$
$$a \mapsto \dim_k(R_a/I_a) \tag{5}$$

Remark 3.8 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 3.9 *Let I be an $\mathcal{A}$-homogeneous ideal and $\prec$ a term order then*

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

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

$$R_a/I_a \ \to \ R_a/\mathrm{in}_\prec(I)_a.$$

□

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" $\mathrm{in}_v(I)$ 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 $\mathrm{span}_{\mathbb{R}}(C_v(I))$. Recall that $C_v(I)$ is the closure of the equivalence class of v. Pick a basis $u_1, \ldots, u_{n-1} \in \mathbb{N}^n$ for $\mathrm{span}_{\mathbb{R}}(C_v(I))$ with u_1 being a positive vector. The vectors induce a grading $\mathcal{A}_v : \mathbb{N}^n \to \mathbb{N}^{n-1}$ on R by

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

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

Lemma 3.10 *Let $\prec$ be a term order, I an ideal, (C_1, C_2) a directed edge with respect to the orientation in Definition 3.4 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* $\mathrm{in}_\prec(\mathrm{in}_v(I)) = M_2$.

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

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

The first one holds by the choice of ω_2. The second one is clear since $\mathrm{in}_{\omega_2}(I)$ is a monomial ideal. The third one holds by Lemma 2.13 and Lemma 2.15. By Lemma 2.13 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 $\mathcal{G}_\prec(\mathrm{in}_v(I))$. If we can show that $\mathrm{in}_{\prec_{\omega_2}}(g) = \mathrm{in}_\prec(g)$ for all elements $g \in \mathcal{G}_\prec(\mathrm{in}_v(I))$ then we know that $\mathcal{G}_\prec(\mathrm{in}_v(I))$ is also a Gröbner basis with respect to $\prec_{\omega_2}$ and the generators for the initial ideal $\mathrm{in}_{\prec_{\omega_2}}(\mathrm{in}_v(I))$ are exactly the same as those for $\mathrm{in}_\prec(\mathrm{in}_v(I))$. This would complete the proof.

The reduced Gröbner basis $\mathcal{G}_\prec(\mathrm{in}_v(I))$ 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 3.4 there exists a normal $q-p$ of the facet with $\mathbf{x}^p \prec \mathbf{x}^q$ and $\langle \omega_2, q-p\rangle > 0$. Since $\prec$ and $\prec_{\omega_2}$ agree on one normal vector they must agree on all exponent differences of elements in $\mathcal{G}_\prec(\mathrm{in}_v(I))$. □

Notice that by Proposition 3.9 any initial ideal $\mathrm{in}_\prec(\mathrm{in}_v(I))$ of $\mathrm{in}_v(I)$ has the same $\mathcal{A}_v$-graded Hilbert function as $\mathrm{in}_v(I)$.

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 $\mathrm{in}_\prec(\mathrm{in}_v(I))$ 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 3.11 Let $\prec, M_1, M_2, u_1, \ldots, u_{n-1}$ and v be as above with $\mathrm{in}_\prec(\mathrm{in}_v(I)) = M_2$. Let N_1 and N_2 be the monomials in M_1 and M_2 respectively. We define the bijection $\phi_{\prec M_1 M_2} : N_1 \to N_2$ in the following way: For a monomial $\mathbf{x}^b \in N_1$ look at the monomials $B_1 \subseteq N_1$ and $B_2 \subseteq N_2$ with the same $\mathcal{A}$-degree as $\mathbf{x}^b$. 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 [14, Lemma 4.1]:

Lemma 3.12 *Let $\leq_1$ and $\leq_2$ be two term orders. If $f_1^1, \ldots, f_s^1$ is a vector space basis for I_a such that $\mathrm{in}_{\leq_1}(f_1^1), \ldots, \mathrm{in}_{\leq_1}(f_s^1)$ is a basis for $\mathrm{in}_{\leq_1}(I)_a$, then there exists a basis $f_1^2, \ldots, f_s^2$ for I_a such that $\mathrm{in}_{\leq_2}(f_1^2), \ldots, \mathrm{in}_{\leq_2}(f_s^2)$ is a basis for $\mathrm{in}_{\leq_2}(I)_a$ and*

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

Corollary 3.13 *Let the setting be as in Definition 3.11. If $\mathbf{x}^b \in M_1$ then $\phi_{\prec M_1 M_2}(\mathbf{x}^b) \not\prec \mathbf{x}^b$.*

Proof. Let a be the $\mathcal{A}$-degree of $\mathbf{x}^b$. We apply Lemma 3.12 with I in the lemma being $\mathrm{in}_v(I)$. Let $\leq_1$ be $\prec$ and $\leq_2$ be the refinement of the preorder induced by u_1 with the reversed order of $\prec$. By the orientation of the graph $M_1 = \mathrm{in}_{\leq_2}(\mathrm{in}_v(I))$ and $M_2 = \mathrm{in}_{\leq_1}(\mathrm{in}_v(I))$. By multiplying elements of $\mathcal{G}_\prec(\mathrm{in}_v(I))$ by monomials we can

construct a k-basis $f_1^1, \ldots, f_s^1$ of $\mathrm{in}_v(I)_a$ with $\mathrm{in}_{\leq_1}(f_1^1), \ldots, \mathrm{in}_{\leq_1}(f_s^1)$ being a basis of $(M_2)_a$. By the lemma there is a basis $\mathrm{in}_{\leq_2}(f_1^2), \ldots, \mathrm{in}_{\leq_2}(f_s^2)$ of $(M_1)_a$. Sort the list of inequalities in the lemma with $\mathrm{in}_{\leq_2}(f_i^2)$ 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 ...

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

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

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

Proof of Theorem 3.6. Suppose $C_1, C_2, \ldots, C_m$ was a path in the oriented graph with $C_1 = C_m$. Let $M_1, \ldots, M_m$ denote the initial ideals and $N_1, \ldots, N_m$ their monomials. We will prove that the bijection $\phi := \phi_{\prec M_{m-1} M_m} \circ \cdots \circ \phi_{\prec M_1 M_2}$ is the identity on M_1. Suppose it is not the identity and let $\mathbf{x}^b$ be the smallest element in M_1 with respect to $\prec$ that is not fixed by ϕ. By Corollary 3.13, $\mathbf{x}^b$ is the image of a smaller element in M_1 with respect to $\prec$. But this element is fixed by the minimality of $\mathbf{x}^b$ — a contradiction. The composition being the identity implies by Corollary 3.13 that $\phi_{\prec M_i M_{i+1}}$ is the identity for all i. Hence $M_i = M_{i+1}$, contradicting that $M_1, M_2, \ldots, M_m$ 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 2.13 $\mathrm{in}_\prec(\mathrm{in}_v(I)) = \mathrm{in}_\prec(I)$. By Lemma 3.10 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 [16] $\prec$ has a matrix representation $(\tau_0, \ldots, \tau_{n-1}) \in \mathbb{R}^{n \times n}$ such that $\tau_\varepsilon := \tau_0 + \varepsilon\tau_1 + \cdots + \varepsilon^{n-1}\tau_{n-1} \in \mathrm{int}\, C_\prec(I)$ for $\varepsilon > 0$ sufficiently small. Furthermore, for any $f \in R$, $\mathrm{in}_{\tau_\varepsilon}(f) = \mathrm{in}_\prec(f)$ for $\varepsilon > 0$ sufficiently small. If $C_{\prec'}(I)$ is a sink then according to Definition 3.4 there exists a complete list of inner normals $q_1 - p_1, \ldots, q_r - p_r$ of $C_{\prec'}(I) \cap \mathbb{R}^n_{\geq 0}$ such that $\mathrm{in}_\prec(\mathbf{x}^{q_i} - \mathbf{x}^{p_i}) = \mathbf{x}^{q_i}$. 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, $\tau_\varepsilon \in \mathrm{int}\, C_{\prec'}(I)$ for $\varepsilon > 0$ sufficiently small. We conclude that $C_{\prec'}(I) = C_\prec(I)$. □

4 Implementation issues

We can find a single Gröbner cone by applying Buchberger's algorithm and Corollary 2.11 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 known theorem which we shall not prove:

Theorem 4.1 *Let $I \subseteq R = k[x_1, \ldots, x_n]$ be an ideal. The marked reduced Gröbner bases of I, the monomial initial ideals of I (w.r.t. a positive vector) and the full-dimensional Gröbner cones are in bijection.*

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. With the right definition of search edges the search tree will only consist of flippable facets.

At the end of the section we will see how the search edge computation in the reverse search algorithm can be implemented and 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 2.6 (or Corollary 2.11) 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 $i = 1, \ldots, n$ 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 [11] 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 2.14 tells us how to compute $\mathcal{G}_\prec(\mathrm{in}_v(I))$. This marked reduced Gröbner basis can be computed from $\mathcal{G}_\prec(I)$ as $\{\mathrm{in}_v(g)\}_{g \in \mathcal{G}_\prec(I)}$ 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 $\{\mathrm{in}_v(g)\}_{g \in \mathcal{G}_\prec(I)}$ 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 $\mathrm{in}_\prec(I)$. 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 $\mathrm{flip}(\mathcal{G}_\prec(I), F)$ 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 $\mathrm{flip}(\mathcal{G}_\prec(I), F)$ 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 $\mathcal{G}$.

Algorithm 4.2
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: $\mathcal{G} = \mathrm{flip}(\mathcal{G}_\prec(I), F)$.
{
Let v be a positive vector in the relative interior of F;
Compute $\mathcal{G}_\prec(\mathrm{in}_v(I)) = \{\mathrm{in}_v(g) : g \in \mathcal{G}_\prec(I)\}$;
Compute the marked basis $\mathcal{G}_{\prec_{-\alpha}}(\mathrm{in}_v(I))$ from $\mathcal{G}_\prec(\mathrm{in}_v(I))
using Buchberger's algorithm;
$\mathcal{G} := \{g - g^{\mathcal{G}_\prec(I)} : g \in \mathcal{G}_{\prec_{-\alpha}}(\mathrm{in}_v(I))\}$;
Mark the term $\mathrm{in}_{\prec_{-\alpha}}(g)$ in each element $g - g^{\mathcal{G}_\prec(I)}$ in $\mathcal{G}$;
Turn $\mathcal{G}$ into a reduced basis;
}

The algorithm is a special case of the local change procedure for a single step in the Gröbner walk [5]. See [8, Proposition 3.2] for a new treatment and a proof. Here we will just 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, $\mathrm{in}_\omega(I) = \mathrm{in}_v(I)$ is homogeneous with respect to the ω-grading. Since F is $(n-1)$-dimensional, $\mathrm{in}_v(I)$ is homogeneous with respect to all vectors inside $\mathrm{span}_{\mathbb{R}}(\alpha)^\perp$. All Gröbner bases of $\mathrm{in}_v(I)$ are homogeneous in the same way. Consequently, each of them must consist of polynomials of the form $\sum_{s=0}^{t} c_s \mathbf{x}^{(a+sb)}$ where $a \in \mathbb{N}^n$ and $b \in \mathbb{Z}^n$ is parallel to α. 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 $\mathbf{x}^\gamma \prec \mathbf{x}^{\gamma+\alpha}$ for $\gamma \in \mathbb{N}^n$ or in the situation where $\mathbf{x}^{\gamma+\alpha} \prec \mathbf{x}^\gamma$ for $\gamma \in \mathbb{N}^n$. Thus specifying α or $-\alpha$ as a term order suffices — no tie-breaker is needed. The initial ideal $\mathrm{in}_v(I)$ can have at most two reduced Gröbner bases. Both term orders are legal since $\mathrm{in}_v(I)$ is homogeneous with respect to the strictly positive vector v.

The Gröbner basis $\mathcal{G}_\prec(\mathrm{in}_v(I))$ 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 2.14. Taking the initial form $\mathrm{in}_v(g)$ of a polynomial $g \in \mathcal{G}_\prec(I)$ without computing v is done as follows. By Corollary 2.11, $\mathrm{in}_\prec(\mathrm{in}_v(g)) = \mathrm{in}_\prec(g)$ and thus we already know one term of $\mathrm{in}_v(g)$ since $\mathrm{in}_\prec(g)$ is the marked term of g in $\mathcal{G}_\prec(I)$. Since every ω in the relative interior of F will have $\mathcal{G}_\prec(\mathrm{in}_\omega(I)) = \mathcal{G}_\prec(\mathrm{in}_v(I))$ the remaining terms of $\mathrm{in}_v(g)$ are exactly the terms in g with the same ω-degree as $\mathrm{in}_\prec(g)$ for all ω in the relative interior of F and consequently for all ω in $\mathrm{span}_{\mathbb{R}}(\alpha)^\perp$. In other words a

term of g is in $\mathrm{in}_v(g)$ if and only if its exponent vector minus the exponent of $\mathrm{in}_\prec(g)$ 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 $\mathcal{G}_{\prec_{-\alpha}}(\mathrm{in}_v(I))$ or in any other subsequent step. The vector v also remains unknown in the entire process.

4.3 Computing the search edge

Let $\prec$ be the term order used for orienting the graph of the Gröbner fan. In Algorithm 3.2 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 3.1 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 $e_{\mathcal{C}_{\prec'}(I)}$ 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.

A better strategy is to draw a straight line from a point in the cone $C_{\prec'}(I)$ to the cone of the sink and choose the first facet intersecting this line as $e_{C_{\prec'}(I)}$. A point in the cone $C_{\prec'}(I)$ can be computed deterministically by linear programming. Two problems arise. The straight line might not intersect a unique facet and we may not know a point in the cone of the sink. Both problems can be solved using formal perturbation of the end points of the line. This was worked out in detail in [8]. Here we explain how it works for lexicographic term orders and with one end point perturbed.

Lemma 4.3 *Let $I \subseteq R$ be an ideal and $\prec$ the lexicographic term order with $x_1 \succ x_2 \succ \cdots \succ x_n$. Define $\tau_\varepsilon = (\varepsilon^0, \varepsilon^1, \ldots, \varepsilon^{n-1})$. There exists a $\delta > 0$ such that $\mathrm{in}_{\tau_\varepsilon}(I) = \mathrm{in}_\prec(I)$ for all $\varepsilon \in (0, \delta)$.*

Proof. This follows from Lemma 2.10 since $\prec$ and τ_ε agree on a finite set of polynomials for small ε. □

Let σ be a deterministically computed interior point of the cone of $C_{\prec'}(I)$ and 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 $C_{\prec'}(I)$ is the sink.

Let $\{\alpha_1, \ldots, \alpha_m\}$ be the set of potential inner facet normals 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).

$$t_i < t_j \qquad \Longleftrightarrow \tag{6}$$

$$\frac{\langle \sigma, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle - \langle \tau_\varepsilon, \alpha_i \rangle} < \frac{\langle \sigma, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle - \langle \tau_\varepsilon, \alpha_j \rangle} \qquad \Longleftrightarrow \tag{7}$$

$$\frac{\langle \sigma, \alpha_i \rangle - \langle \tau_\varepsilon, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle} > \frac{\langle \sigma, \alpha_j \rangle - \langle \tau_\varepsilon, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle} \qquad \Longleftrightarrow \tag{8}$$

$$\frac{\langle \tau_\varepsilon, \alpha_i \rangle}{\langle \sigma, \alpha_i \rangle} < \frac{\langle \tau_\varepsilon, \alpha_j \rangle}{\langle \sigma, \alpha_j \rangle} \qquad \Longleftrightarrow \tag{9}$$

$$\langle \tau_\varepsilon, \langle \sigma, \alpha_j \rangle \alpha_i \rangle < \langle \tau_\varepsilon, \langle \sigma, \alpha_i \rangle \alpha_j \rangle \qquad \Longleftrightarrow \tag{10}$$

$$\mathbf{x}^{\langle \sigma, \alpha_j \rangle \alpha_i} \prec \mathbf{x}^{\langle \sigma, \alpha_i \rangle \alpha_j} \tag{11}$$

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 defining and computing $e_{C_{\prec'}(I)}$. We simply choose the facet defined by a_i where t_i is smallest among $\{t_1, \ldots, t_m\}$ (for small $\varepsilon > 0$).

4.4 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 have done in the usual reverse search or breadth-first enumeration without symmetry. It is not clear how to combine symmetry-exploiting with reverse search.

5 Complexity

In this section we will discuss the complexity of enumerating the maximal cones of the Gröbner fan of an ideal I by reverse search. We will assume that I is homogeneous with respect to a positive vector. This guarantees that any facet of a full-dimensional Gröbner cone is flippable.

We identify the following important sub-algorithms:

- Computation of the facet normals of the Gröbner cone of a marked reduced Gröbner basis $\mathcal{G}$. We denote the time for this operation by $T_{\mathbf{facets}}(\mathcal{G})$.
- Computation of a search edge given a marked reduced Gröbner basis $\mathcal{G}$ as described in Subsection 4.3. We denote the time for this operation by $T_{\mathbf{shoot}}(\mathcal{G})$.
- Conversion of a marked reduced Gröbner basis $\mathcal{G}_1$ into a marked reduced Gröbner basis $\mathcal{G}_2$ where the Gröbner cones of $\mathcal{G}_1$ and $\mathcal{G}_2$ are assumed to share a facet. We denote the time for this operation by $T_{\mathbf{flip}}(\mathcal{G}_1,\mathcal{G}_2)$.

For simplicity we will assume that the time used for solving any of the linear programs in our algorithms only depends on the dimensions of its matrix form. We let $T_{\mathbf{lp}}(n,r)$ be the time needed to solve a linear programming problem with n variables and r constraints. Then $T_{\mathbf{facets}}(\mathcal{G})$ and $T_{\mathbf{shoot}}(\mathcal{G})$ can be expressed in terms of $T_{\mathbf{lp}}(n,r)$.

The time $T_{\mathbf{facets}}(\mathcal{G}) \in O(T_{\mathbf{lp}}(n,r)r)$ where r is the number of non-leading terms in $\mathcal{G}$. The reason is that each non-leading term in $\mathcal{G}$ gives an inequality in the description of the Gröbner cone. Checking if the inequality defines a facet takes one linear program. In addition duplicates should be removed from the set of facet normals and further vectors should be eliminated until no parallel vectors exist. The time for this step is dominated by the time for solving LPs.

The time $T_{\mathbf{shoot}}(\mathcal{G}) \in O(rn^2 + T_{\mathbf{lp}}(n,r))$ where r is the number of non-leading terms in $\mathcal{G}$. The first step in the algorithm is to deterministically find a relative interior point of the Gröbner cone. This is done in time $T_{\mathbf{lp}}(n,r)$. After this the smallest vector among the r defining vectors for the cone with respect to the ordering in Subsection 4.3 needs to be found. Comparing two vectors takes $O(n^2)$ operations in the worst case. These are operations in $\mathbb{Q}$. In the above estimate we assume that each operation takes constant time.

We have no good bound for the complexity of flipping. Now we count the number of times each of the three sub-algorithms are applied when enumerating the graph of the Gröbner fan of I using reverse search.

- The facets of each Gröbner cone are computed exactly once in Algorithm 3.2 (right after the Gröbner basis has been output). We remark that since we are only interested in facets with the correct orientation the number of LPs that really need to be solved is lower than the r in the discussion above. We will not take this into account in our analysis.
- Checking if an edge is a search edge is done once for every edge. Every time we need to recompute a search edge and compare it to the edge. Hence the total

number of times we need to compute a search edge is equal to the number of edges in the graph of the Gröbner fan of I.

- When a vertex $\mathcal{G}$ is processed by Algorithm 3.2 we must test for every ingoing edge if the edge is a search edge. To test this we first compute $\text{flip}(\mathcal{G}, F)$ where F is the facet of the cone corresponding to the edge in question. If the edge is equal to the search edge of $\text{flip}(\mathcal{G}, F)$ we do an enumeration of the subtree with root $\text{flip}(\mathcal{G}, F)$. If not, $\text{flip}(\mathcal{G}, F)$ is forgotten. Since all vertices are processed once and every edge is ingoing for exactly one vertex the number of times $\text{flip}(\mathcal{G}, F)$ needs to be computed is equal to the number of edges in the graph. We remark that the variant of the reverse search where the search path for the current vertex is not stored on the recursion stack would require twice as many computations of this kind.

Let E be the edges and V be the vertices of the graph. The total time complexity of the enumeration of (V, E) is:

$$O\Big(\sum_{\mathcal{G}\in V} T_{\textbf{facets}}(\mathcal{G}) + \sum_{(\mathcal{G}_1,\mathcal{G}_2)\in E} T_{\textbf{shoot}}(\mathcal{G}_1) + \sum_{(\mathcal{G}_1,\mathcal{G}_2)\in E} T_{\textbf{flip}}(\mathcal{G}_1, \mathcal{G}_2)\Big)$$

Substituting with the time needed for solving the LPs we get the following theorem:

Theorem 5.1 *Let (V, E) be the graph of the Gröbner fan of I. The time complexity for computing this graph given a marked reduced Gröbner basis is in the class of functions*

$$O\Big(\sum_{\mathcal{G}\in V} T_{\textbf{lp}}(n, r(\mathcal{G}))r(\mathcal{G}) + \sum_{(\mathcal{G}_1,\mathcal{G}_2)\in E} r(\mathcal{G}_1)n^2 + T_{\textbf{lp}}(n, \mathcal{G}_1) + \sum_{(\mathcal{G}_2,\mathcal{G}_1)\in E} T_{\textbf{flip}}(\mathcal{G}_1, \mathcal{G}_2)\Big)$$

where $r(\mathcal{G})$ is the number non-leading terms in the marked reduced Gröbner basis $\mathcal{G}$. In particular, the first two terms are bounded by a polynomial in the size of the output.

Corollary 5.2 *If for a given class of ideals the time $T_{\textbf{flip}}(\mathcal{G}_1, \mathcal{G}_2)$ is bounded by a polynomial in the size of the binary encoding of $\mathcal{G}_1$ and $\mathcal{G}_2$ then the enumeration of the reduced Gröbner bases for an ideal in the class by reverse search is a polynomial time algorithm in the size of the output.*

6 Computational results and examples

The algorithms presented in this paper were implemented in the software package Gfan [12]. In this section we present examples of Gröbner fans computed using this package. The first example comes with a picture and gives an idea of the kind of geometric shape a Gröbner fan might have.

Example 6.1 [18, 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 $\mathbb{R}^3_{\geq 0}$. The intersection of the fan with the standard simplex in $\mathbb{R}^3$ is shown in Figure 4.

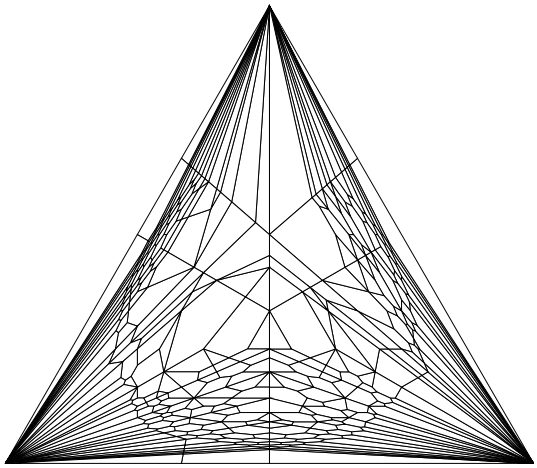

Figure 4: The Gröbner fan of the ideal in Example 6.1 intersected with the standard 2-simplex. The a-axis is on the right, the b-axis on the left and the c-axis at the top.

We now list some families of ideals used in our computations. The Gröbner fans of these ideals have been computed for the parameters listed in the table of Figure 5. The ambient field is always $\mathbb{Q}$. The columns of the table 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 lowest dimensional Gröbner cone $C_0(I)$. Note that h is the dimension of the *homogeneity space* of the ideal which is the common subspace contained in every Gröbner cone of the ideal. The quantity "d" is the lowest total degree of any reduced Gröbner basis of the ideal and "D" is the highest. The *f-vector* of the Gröbner fan is an ordered list of the number of h-dimensional cones, $h+1$-dimensional cones etc., up to the number of n-dimensional cones.

Example	n	h	d	D	f-vector
$\text{Det}_{3,3,4}$	12	6	3	3	(1,12,66,204,342,288,96)
$\text{Det}_{3,3,5}$	15	7	3	3	(1,45,585,3390,10710,19890,21750,12960,3240)
$\text{Det}_{3,4,4}$	16	7	3	5	(1,?,?,?,?,?,?,?,?,163032)
$\text{Detsym}_{3,4}$	10	4	3	8	(1,518,5412,20505,36024,29808,9395)
$\text{Grass}_{2,5}$	10	5	2	3	(1,20,120,300,330,132)
Cyclic_5	5	0	8	15	(1,?,?,?,?,55320)
J_4	4	1	3	8	(1,200,516,318)

Figure 5: Statistics for the Gröbner fans computed using Gfan.

Example 6.2 Let $\mathrm{Det}_{t,m,n}$ 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 6.3 Let $\mathrm{Grass}_{d,n}$ 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 6.4 Let $\mathrm{Detsym}_{t,n}$ 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, $\mathrm{Detsym}_{3,4}$ 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 6.5 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 [3]. These polynomial systems have become benchmarks for computer algebra packages and their lexicographic Gröbner bases are notoriously hard to compute.

Example 6.6 Let K_n denote the complete graph on n vertices and I_n be the *Stanley-Reisner* ideal of this graph. The Stanley-Reisner ideal of a simplicial complex Δ is the ideal generated by all monomials $x_{i_1}x_{i_2}\cdots x_{i_n}$ such that $\{i_1,\ldots,i_n\}$ is not a face of Δ. Apply a generic linear change of coordinates to I_n to obtain the ideal J_n. The generators of J_n typically have very complicated coefficients. For example, the first generator in our J_4 was

$$\begin{array}{llll} a^3 & +\frac{4980248985}{343338664}a^2c & +\frac{2079196217}{257503998}abc & +\frac{86858380}{128751999}b^2c \\ & -\frac{359584197}{171669332}bc^2 & -\frac{84523033581}{1373354656}c^3 & -\frac{11737327991}{51500799600}a^2d \\ & +\frac{1194144014}{9656399925}b^2d & +\frac{394500908221}{25750399800}acd & -\frac{47953955497}{25750399800}bcd \\ & -\frac{4583330213}{3862559970}ad^2 & +\frac{181743499}{364392450}bd^2 & -\frac{429736138279}{25750399800}cd^2 \end{array} \quad \begin{array}{l} -\frac{2205648949}{42917333}ac^2 \\ -\frac{16299027451}{38625599700}abd \\ +\frac{195491595943}{2985553600}c^2d \\ +\frac{8566043731}{12875199900}d^3. \end{array}$$

The initial ideals of J_n are known as the *generic initial ideals* of I_n. The reverse lexicographic generic initial ideals of an ideal have played an important role in commutative algebra and algebraic geometry while other generic initial ideals have not been explored too much. We computed the Gröbner fan of J_4.

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 these examples. In example $\mathrm{Det}_{3,3,4}$ 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 Pentium processor. Using reverse search without symmetry the same computation would take approximately 14 hours. The f-vector extraction routine in Gfan only works for complete fans. This is why the f-vector for the Cyclic_5 example is not shown.

References

[1] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65:21–46, 1996.

[2] David Bayer and Ian Morrison. Standard bases and geometric invariant theory I. Initial ideals and state polytopes. *J. Symb. Comput.*, 6(2/3):209–217, 1988.

[3] 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.

[4] Tristram Bogart, Anders Jensen, Rekha Thomas, David Speyer, and Bernd Sturmfels. Computing tropical varieties. 2005, math.AG/0507563.

[5] Stéphane Collart, Michael Kalkbrener, and Daniel Mall. Converting bases with the Gröbner walk. *J. Symb. Comput.*, 24(3/4):465–469, 1997.

[6] David Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms.* Springer, 1996.

[7] 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`.

[8] Komei Fukuda, Anders Jensen, Niels Lauritzen, and Rekha Thomas. The generic Gröbner walk. 2005, math.AC/0501345.

[9] William Fulton. *Introduction to Toric Varieties.* Princeton University Press, 1993.

[10] Torbjörn Granlund et al. GNU multiple precision arithmetic library 4.1.2, December 2002. `http://swox.com/gmp/`.

[11] Birkett Huber and Rekha R. Thomas. Computing Gröbner fans of toric ideals. *Experimental Mathematics*, 9(3/4):321–331, 2000.

[12] Anders N. Jensen. Gfan, a software system for Gröbner fans. Available at `http://home.imf.au.dk/ajensen/software/gfan/gfan.html`.

[13] Anders N. Jensen. A non-regular Gröbner fan. 2005, math.CO/0501352.

[14] Niels Lauritzen. Truncated Gröbner fans and lattice ideals. 2005, math.AG/0509247.

[15] Daniel Mall. Gröbner fans and projective schemes. *Progress in Computer Science and Applied Logic*, 15:181–191, 1998.

[16] Teo Mora and Lorenzo Robbiano. The Gröbner fan of an ideal. *J. Symb. Comput.*, 6(2/3):183–208, 1988.

[17] Jörg Rambau. TOPCOM: Triangulations of point configurations and oriented matroids. *ZIB report*, 02-17, 2002.

[18] Bernd Sturmfels. *Gröbner bases and Convex Polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, 1996.

Komei Fukuda, Institute for Operations Research and Institute of Theoretical Computer Science ETH Zentrum, CH-8092 Zurich, Switzerland and Mathematics Institute / ROSO EPFL, CH-1015 Lausanne, Switzerland `fukuda@ifor.math.ethz.ch`.

Anders N. Jensen, Institut for Matematiske Fag, Aarhus Universitet, DK-8000 Århus, Denmark, `ajensen@imf.au.dk`.

Rekha R. Thomas, Department of Mathematics, University of Washington, Seattle, WA 98195-4350, USA, `thomas@math.washington.edu`.