COMPUTING GRÖBNER FANS OF TORIC IDEALS

BIRKETT HUBER AND REKHA R. THOMAS

Abstract. The monomial initial ideals of a graded polynomial ideal are in bijection with the vertices of a convex polytope known as the state polytope of the ideal. The Gröbner fan of the ideal is the normal fan of its state polytope. In this paper we present a software system called TiGERS for computing the Gröbner fan of a toric ideal by enumerating the edge graph of its state polytope. The key contributions are an inexpensive algorithm for local change of Gröbner bases in toric ideals and the identification of a reverse search tree on the vertices of the state polytope. Using these ideas we obtain a combinatorial Gröbner walk procedure for toric ideals. TiGERS has been used to compute state polytopes with over 200,000 vertices.

1. Introduction

Consider the polynomial ring $k[\mathbf{x}] := k[x_1, \ldots, x_n]$ where k is a field and an ideal $I \subset k[\mathbf{x}]$ that is homogeneous with respect to a positive grading $degree(x_i) = \omega_i \in \mathbf{N}\backslash\{0\}$. We use $\mathbf{N}$ to denote the set of non-negative integers. The *initial ideal* of I with respect to a *term order* $\succ$ on $k[\mathbf{x}]$ is the *monomial ideal* $in_\succ(I) := \langle in_\succ(f) \,:\, f \in I\rangle$ where $in_\succ(f)$ is the *initial term* of $f \in I$ with respect to $\succ$. The *reduced Gröbner basis* of I with respect to $\succ$ is the unique finite set of monic polynomials $\mathcal{G}_\succ(I) = \{g_1, \ldots, g_t\} \subset I$ such that (i) $in_\succ(I) = \langle in_\succ(g_1), \ldots, in_\succ(g_t)\rangle$ and (ii) for $i \neq j$, no term of g_i is divisible by $in_\succ(g_j)$. Reduced Gröbner bases of polynomial ideals can be computed using *Buchberger's algorithm.* See [AL] or [CLO] for further details.

Given an arbitrary *weight vector* $c \in \mathbf{R}^n$, and a polynomial $f = \sum k_\alpha x^\alpha \in k[\mathbf{x}]$, the *initial term* of f with respect to c is defined to be the sum of all terms $k_\alpha x^\alpha$ in f such that the inner product $c \cdot \alpha$ is maximal. The *initial ideal* of I with respect to c is then $in_c(I) := \langle in_c(f) \,:\, f \in I\rangle$. If $in_c(I)$ is a monomial ideal then c is said to be *generic* for I. It is known that for a term order $\succ$ on $k[\mathbf{x}]$, there exists a weight vector $c \in \mathbf{N}^n$ such that $in_c(I) = in_\succ(I)$. In this case we say that c represents $\succ$ and it can be shown that $in_c(I) = in_\succ(I)$ if and only if $in_c(g) = in_\succ(g)$ for each g in $\mathcal{G}_\succ(I)$.

Two weight vectors c_1 and c_2 in $\mathbf{R}^n$ are said to be *equivalent modulo* I whenever $in_{c_1}(I) = in_{c_2}(I)$. The set of all weight vectors that are equivalent to $c \in \mathbf{R}^n$ form a relatively open polyhedral cone in $\mathbf{R}^n$, the closure of which is called the *Gröbner cone* of c. The Gröbner cone of c is n-dimensional if and only if c is generic for I. The collection of all Gröbner cones of I form a *polyhedral fan* in $\mathbf{R}^n$ called the *Gröbner fan* of I [MR]. Since I is homogeneous with respect to a positive grading, this fan is, in fact, *complete* (i.e., covers $\mathbf{R}^n$) and further, each Gröbner cone of I contains a strictly positive vector of $\mathbf{R}^n$.

The Gröbner fan of I is the *normal fan* of a polytope in $\mathbf{R}^n$ called the *state polytope* of I [BM], denoted as $St(I)$. Therefore, I has only finitely many distinct reduced Gröbner bases as c varies over all weight vectors in $\mathbf{R}^n$. (See [Mac] for a new proof.) The faces of $St(I)$ are in bijection with the distinct initial ideals of I, with the vertices of $St(I)$ corresponding bijectively to the distinct monomial initial ideals of I. The distinct monomial initial ideals

Research at MSRI is supported in part by NSF grant DMS-9701755.

of I, in turn, are in bijection with the distinct reduced Gröbner bases of I obtained from term orders. Hence, computing all such initial ideals amounts to searching the *edge graph* of $St(I)$. Two monomial initial ideals $in_{c_1}(I)$ and $in_{c_2}(I)$ are said to be *adjacent* if the corresponding vertices of $St(I)$ are adjacent or equivalently, if the Gröbner cones of the generic weight vectors c_1 and c_2 share a common *facet.* See Chapters 1-3 in [Stu] for proofs of the results quoted above and a full discussion of Gröbner fans and state polytopes of graded polynomial ideals. Algorithms for their construction are also included.

Given a matrix $A = [a_1, \ldots, a_n] \in \mathbf{Z}^{d\times n}$ of rank d, the *toric ideal* of A, denoted as I_A, is the kernel of the homomorphism $k[x_1, \ldots, x_n] \to k[t_1^{\pm 1}, \ldots, t_d^{\pm 1}]$, such that $x_j \mapsto t^{a_j}$ (see Chapter 4 in [Stu]). The ideal I_A is a d-dimensional prime ideal that is generated by the *binomials* $x^{u^+} - x^{u^-}$ where $u = u^+ - u^-$ lies in the $(n-d)$-dimensional saturated lattice $ker_{\mathbf{Z}}(A) := \{u \in \mathbf{Z}^n : Au = 0\}$. Here $u^- = (-u)^+$ and u^+ is defined as $u_i^+ = u_i$ if $u_i > 0$ and $u_i^+ = 0$ otherwise. Hence $u^+, u^- \in \mathbf{N}^n$. The mechanics of the Buchberger algorithm ensure that every reduced Gröbner basis of I_A again consists of finitely many binomials of the above type. We will assume that $ker_{\mathbf{Z}}(A) \cap \mathbf{N}^n = \{0\}$ which guarantees a positive integral vector ω in the row space of A. Then I_A is homogeneous with respect to the grading $degree(x_i) = \omega_i$.

Let $\mathcal{G}_c = \{x^{\alpha_i} - x^{\beta_i} : i = 1, \ldots, t\}$ be the reduced Gröbner basis of I_A with respect to a generic weight vector $c \in \mathbf{R}^n$. (The positive term of a binomial in $\mathcal{G}_c$ is always assumed to be the initial term with respect to c.) The Gröbner cone of c is then the n-dimensional polyhedral cone $\mathcal{K}_c := \{u \in \mathbf{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u,\ i = 1, \ldots, t\}$ whose lineality space $(\mathcal{K}_c \cap -\mathcal{K}_c)$ is the row space of A. We may assume that c is a strictly positive integral vector since $\mathcal{K}_c$ is a rational cone and $\omega \in \mathcal{K}_c$. Further, c' is equivalent to c if and only if $in_{c'}(x^{\alpha_i} - x^{\beta_i}) = x^{\alpha_i}$ for each binomial $x^{\alpha_i} - x^{\beta_i} \in \mathcal{G}_c$. State polytopes and Gröbner fans of toric ideals were studied in [ST]. This paper gives several custom-tailored construction methods for these entities when the ideal is toric.

The Gröbner fan of I_A has various applications: In [ST] it was used as a model for *sensitivity analysis* for the family of *integer programs* $min\{c \cdot x : Ax = b,\ x \in \mathbf{N}^n\}$ as b varies in $\mathbf{Z}^d$ and c in $\mathbf{R}^n$. The *secondary fan* of A [BFS], [GKZ] is a coarsening of the Gröbner fan of I_A (see Chapter 8 in [Stu]). This fan has important applications to problems in discrete geometry. In between the secondary fan and the Gröbner fan of A lives the *hypergeometric fan* of A which has been used for studying *A-hypergeometric* differential equations using *Gröbner deformations* [SST]. Both these coarser fans can be obtained from the Gröbner fan of I_A. Finally, the *Gröbner walk* procedure introduced in [CKM] uses the Gröbner fan for Gröbner basis conversions.

Example 1.0.1. *Let* $A := \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 1 \end{pmatrix}$. *For ease of exposition we associate the variables* a, b, c, d, e *to the five columns of* A *and consider the toric ideal* $I_A \subset k[a, b, c, d, e]$. *This ideal yields eight distinct monomial initial ideals. The state polytope* $St(I_A)$ *is an octagon in* $\mathbf{R}^5$ *and the Gröbner fan of* I_A *can be drawn in* $\mathbf{R}^2$ *after moding out the row space of* A *from each Gröbner cone. Figure 1 shows a schematic representation of the resulting* pointed *Gröbner fan of* I_A. *Each maximal cone is labeled by the reduced Gröbner basis induced by the weight vectors in the interior of that cone. The binomials in a reduced Gröbner basis that contribute the facet inequalities of its Gröbner cone are marked with dots. Notice that for two adjacent reduced Gröbner bases, their common facet binomial*

$x^\alpha - x^\beta$ appears in both reduced Gröbner bases with x^α as initial term in one basis and x^β as initial term in the other.

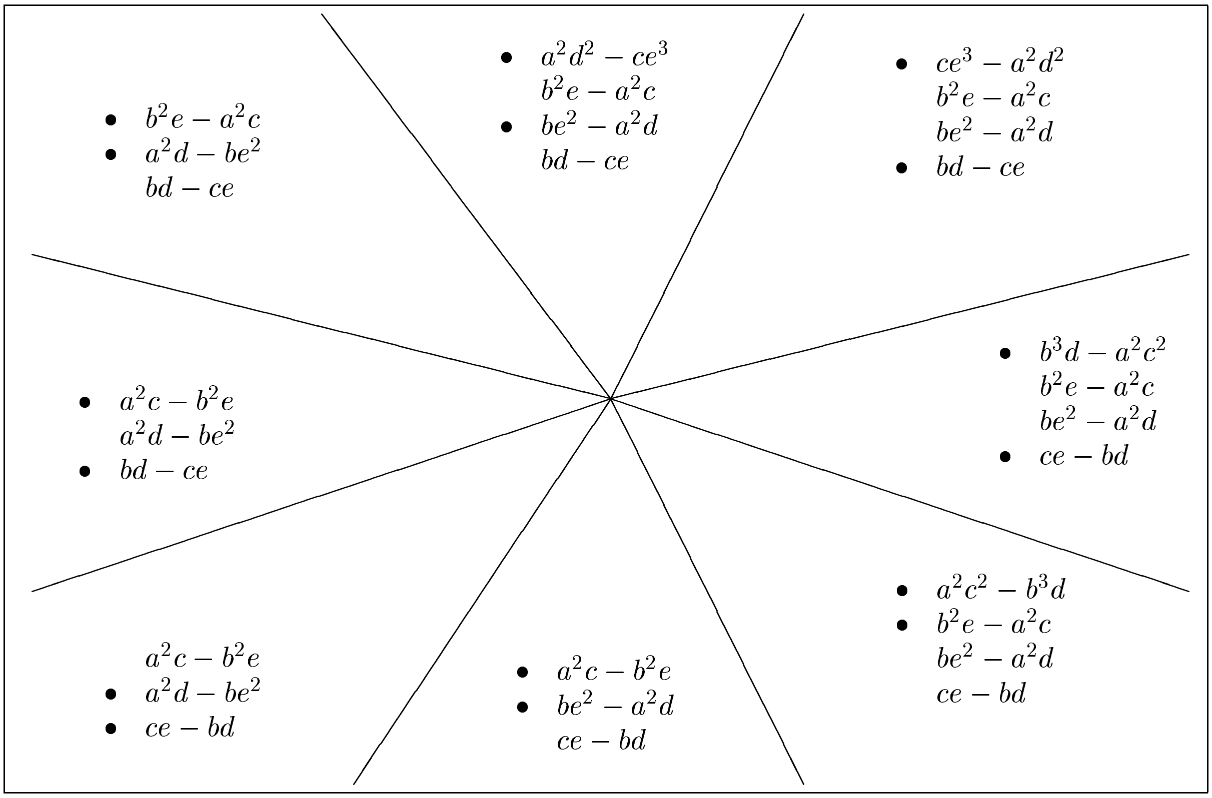

FIGURE 1. A schematic of the Gröbner fan in Example 1.0.1

In this paper we introduce a software system called TiGERS for computing the Gröbner fan of a toric ideal. The program searches the edge graph of the state polytope $St(I_A)$ to find all the distinct monomial initial ideals (reduced Gröbner bases) of I_A. This search can be done in two ways: For moderate sized examples, the graph is searched by a breadth-first search strategy on the entire edge-graph of $St(I_A)$. This approach needs to maintain and search the list of all Gröbner bases found. When $St(I_A)$ is too large this approach will bog down and we use a *reverse search technique* instead. This involves a depth-first search on a subgraph of the edge graph of $St(I_A)$ called the reverse search tree, and can be implemented so that no more than one Gröbner basis ever needs to be stored. In Section 2 we give an overview of the main algorithm in TiGERS and describe the reverse search procedure. These ideas allow *combinatorial Gröbner walks* in toric ideals. For each reduced Gröbner basis found, the algorithm does two main local computations. The first is to determine all the facets of that Gröbner cone and the second is to determine an adjacent Gröbner basis to the current one. In Section 3 we give an algorithm for local change of Gröbner bases in toric ideals based on [CKM]. Unlike usual local change algorithms, our procedure does not require any weight vectors to be computed. We then discuss some special tricks to find the facets of a fixed Gröbner cone in the case of toric ideals. Section 4 reports computational experience with TiGERS.

2. The main algorithm in TiGERS

Let $\mathcal{G}_c = \{x^{\alpha_i} - x^{\beta_i}, i = 1, \ldots, t\}$ be the reduced Gröbner basis of I_A with respect to the generic weight vector $c \in \mathbf{R}^n$. An inequality $\alpha_k \cdot u \geq \beta_k \cdot u$, $k \in \{1, \ldots, t\}$ is irredundant for the Gröbner cone $\mathcal{K}_c = \{u \in \mathbf{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u, i = 1, \ldots, t\}$ if the relaxed cone $\{u \in \mathbf{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u, i = 1, \ldots, t, i \neq k\}$ properly contains $\mathcal{K}_c$. If $\alpha_k \cdot u \geq \beta_k \cdot u$ is irredundant for $\mathcal{K}_c$ then $\mathcal{K}_c \cap \{u \in \mathbf{R}^n : \alpha_k \cdot u = \beta_k \cdot u\}$ is called a *facet* of $\mathcal{K}_c$ and $x^{\alpha_k} - x^{\beta_k}$ is a *facet binomial* of $\mathcal{G}_c$.

Lemma 2.0.2. [ST] *The binomial $x^{\alpha_k} - x^{\beta_k} \in \mathcal{G}_c$ is a facet binomial of $\mathcal{G}_c$ if and only if the linear system $\{u \in \mathbf{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u\, , i = 1, \ldots, t, \, i \neq k\} \cup \{u \in \mathbf{R}^n : \beta_k \cdot u \geq \alpha_k \cdot u\}$ is feasible.*

Therefore, all facets of the Gröbner cone $\mathcal{K}_c$ can be found by checking the feasibility of t systems of linear inequalities, which in turn can be done by linear programming. In practice this can be computationally expensive when t is large and we discuss certain speed ups in Section 3.2.

Suppose $x^\alpha - x^\beta$ is a facet binomial of $\mathcal{G}_c$ such that $in_c(x^\alpha - x^\beta) = x^\alpha$ and $\mathcal{G}_{c'}$ be adjacent to $\mathcal{G}_c$ with $in_{c'}(x^\alpha - x^\beta) = x^\beta$. In order to compute the edge graph of $St(I_A)$, we require a subroutine to make a local change of reduced Gröbner bases from $\mathcal{G}_c$ to $\mathcal{G}_{c'}$ through the facet given by $x^\alpha - x^\beta$. We use $flip(\mathcal{G}_c, x^\alpha - x^\beta)$ to denote both the subroutine that yields $\mathcal{G}_{c'}$ from $\mathcal{G}_c$ as well as the reduced Gröbner basis $\mathcal{G}_{c'}$ that results from the "flip". In Section 3.1 we describe the precise local change algorithm in TiGERS.

Remark 2.0.3. If $x^\alpha - x^\beta$ lies in the reduced Gröbner basis of a toric ideal, then the supports of α and β are disjoint. This guarantees that each facet in a toric Gröbner cone corresponds to precisely one binomial in the corresponding Gröbner basis. However, this is not true for general homogeneous binomial ideals. Consider $J = \langle b^3d - b^2ce, a2c - b2e, bcd - c2e\rangle$. Under the reverse lexicographic order $a \succ c \succ d \succ e \succ b$, its reduced Gröbner basis is $G = \{ceb^2 - db^3, c^2e - cdb, a^2db^3 - e^2b^4, a^2c - eb^2\}$. The inequality $u_3 + u_5 \geq u_2 + u_4$ is a facet inequality of the Gröbner cone of G. For ω in the relative interior of this facet, we get both $in_\omega(ceb^2 - db^3) = ceb^2 - db^3$ and $in_\omega(c^2e - cdb) = c^2e - cdb$.

Algorithm 2.0.4. *Enumerating the edge graph of $St(I_A)$ via breadth-first search.*

Input: *Any reduced Gröbner basis $\mathcal{G}_0$ of I_A.*
Output: *All reduced Gröbner bases of I_A (all vertices of $St(I_A)$).*

```
Todo := [G_0];
Verts := [ ];
While(Todo not empty) do
  G := first-element-in(Todo);
  Remove G from Todo;
  add G to Verts;
  determine list L of facet binomials of G
  for each x^α − x^β in L do
    Compute G' = flip(G, x^α − x^β)
    If G' ∉ Todo ∪ Verts then add G' to Todo;
  End
End
output Verts
```

This algorithm works well in practice but does have the drawback that all vertices must be stored, and that every time a vertex is visited it must be checked against all other vertices seen thus far in order to determine if it is indeed a new vertex. The storage and search costs involved in this procedure can become prohibitive as the size of $St(I_A)$ increases. As an example consider $A := \begin{pmatrix} 3 & 2 & 2 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 2 & 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 1 & 0 & 2 & 0 & 1 & 2 & 3 \end{pmatrix}$. The ideal I_A involves only nine variables and most reduced Gröbner bases of this ideal have fewer than 36 elements each of degree no greater than seven. Yet, $St(I_A)$ has $54,828$ vertices and our breadth first search algorithm was exhausting a personal computer with 64 kilobytes of memory before getting through about $13,000$ vertices. In order to push this calculation through we resorted to the reverse enumeration paradigm of Avis and Fukuda [AF]. The basic idea there is to do a depth first search on a directed subgraph of $St(I_A)$.

We say that a polynomial $f \in k[\mathbf{x}]$ is *marked* by a term order $\succ$ on $k[\mathbf{x}]$, if the initial term of f with respect to $\succ$ has been distinguished from among all terms in f. A polynomial f that has been marked with respect to $\succ$ is said to be *mismarked* with respect to $\succ'$, if $in_{\succ}(f) \neq in_{\succ'}(f)$.

Lemma 2.0.5. *Let $\mathcal{G}_{\succ}$ be the reduced Gröbner basis of I_A with respect to the term order $\succ$. Then for a term order $\succ' \neq \succ$, the reduced Gröbner basis $\mathcal{G}_{\succ'}$ equals $\mathcal{G}_{\succ}$ if and only if no facet binomial of $\mathcal{G}_{\succ}$ is mismarked with respect to $\succ'$.*

Proof: Suppose no facet binomial of $\mathcal{G}_{\succ}$ is mismarked with respect to $\succ'$ and let c' be a weight vector from the interior of the Gröbner cone $\mathcal{K}_{\succ'}$. Then for each facet binomial $x^\alpha - x^\beta$ in $\mathcal{G}_{\succ}$, we have $c' \cdot \alpha > c' \cdot \beta$ which implies that c' lies in the interior of the Gröbner cone $\mathcal{K}_{\succ}$. Hence $\mathcal{K}_{\succ'} = \mathcal{K}_{\succ}$ which implies that $\mathcal{G}_{\succ} = \mathcal{G}_{\succ'}$. Conversely, if $\mathcal{G}_{\succ} = \mathcal{G}_{\succ'}$, then no binomial in $\mathcal{G}_{\succ}$ is mismarked with respect to $\succ'$.

Definition 2.0.6. *For a given term order $\succ$ we define the* reverse search tree $T_{\succ}(I_A)$ *as follows:*
The vertices of $T_{\succ}(I_A)$ are the vertices of $St(I_A)$ (i.e. the various reduced Gröbner bases of I_A arising from term orders). For two reduced Gröbner bases $\mathcal{G}_i$ and $\mathcal{G}_j$, $[\mathcal{G}_i, \mathcal{G}_j]$ directed from $\mathcal{G}_i$ to $\mathcal{G}_j$ is an edge of $T_{\succ}(I_A)$ if $\mathcal{G}_j$ is obtained from $\mathcal{G}_i$ by the subroutine $flip(\mathcal{G}_i, x^\alpha - x^\beta)$ where $x^\alpha - x^\beta$ is the unique facet binomial of $\mathcal{G}_i$ whose leading term is lexicographically maximal among all facet binomials of $\mathcal{G}_i$ that are mismarked with respect to $\succ$.

Theorem 2.0.7. *The reverse search tree $T_{\succ}(I_A)$ is an acyclic directed graph with a unique sink.*

Proof: By Lemma 2.0.5 each reduced Gröbner basis $\mathcal{G}$ of I_A (vertex of $St(I_A)$) except $\mathcal{G}_{\succ}$ has at least one mismarked facet binomial with respect to $\succ$. By the definition of $T_{\succ}(I_A)$, each such $\mathcal{G}$ has a unique adjacent reduced Gröbner basis given by $flip(\mathcal{G}, x^\alpha - x^\beta)$ where $x^\alpha - x^\beta$ is the unique facet binomial of $\mathcal{G}$ whose leading term is lexicographically maximal among all mismarked facet binomials of $\mathcal{G}$. Therefore $T_{\succ}(I_A)$ is a directed graph such that the out-degree of $\mathcal{G}_{\succ}$ in $T_{\succ}(I_A)$ is zero and of all other reduced Gröbner bases is one.

Suppose there was a cycle $C = (v_1, \ldots, v_l)$ of length l in $T_{\succ}(I_A)$ where vertex v_i corresponds to the reduced Gröbner basis $\mathcal{G}_i$ of I_A and $v_{l+1} = v_1$ and $\mathcal{G}_{l+1} = \mathcal{G}_1$. If $x^{\alpha_i} - x^{\beta_i}$ is the common facet binomial of $\mathcal{G}_i$ and $\mathcal{G}_{i+1}$ with x^{α_i} the leading term with respect to $\mathcal{G}_i$ and x^{β_i} the leading term with respect to $\mathcal{G}_{i+1}$ for $i = 1, \ldots, l$, then $v_{i+1} - v_i = \beta_i - \alpha_i$ for $i = 1, \ldots, l$. Since C is a cycle in $T_{\succ}(I_A)$ we get $(\beta_1 - \alpha_1) + (\beta_2 - \alpha_2) + \cdots + (\beta_l - \alpha_l) = 0$. However, each binomial $x^{\alpha_i} - x^{\beta_i}$ for $i = 1, \ldots, l$ is mismarked with respect to $\succ$ which implies that for a weight vector c representing $\succ$ we have, $c \cdot (\beta_i - \alpha_i) > 0$. This leads to the contradiction $0 = c \cdot \sum_{i=1}^{l} (\beta_i - \alpha_i) = \sum_{i=1}^{l} c \cdot (\beta_i - \alpha_i) > 0$.

Corollary 2.0.8. *From any reduced Gröbner basis $\mathcal{G}_c$ of I_A there is a unique path in the reverse search tree $T_{\succ}(I_A)$ to the sink $\mathcal{G}_{\succ}$.*

Corollary 2.0.8 proves that toric ideals admit combinatorial Gröbner walks that can be used for converting one reduced Gröbner basis of I_A into another. Given any reduced Gröbner basis $\mathcal{G}$ of I_A and any term order $\succ$, we can move from $\mathcal{G}$ to $\mathcal{G}_{\succ}$ by tracing the unique path in the reverse search tree $T_{\succ}(I_A)$ from $\mathcal{G}$ to $\mathcal{G}_{\succ}$. Unlike in the usual Gröbner walk procedure [CKM],[AGK] to convert one Gröbner basis into another there are no explicit cost vectors involved in these walks. These combinatorial walks also have the advantage that there is never the danger of walking through a lower dimensional face of a Gröbner cone thus eliminating several numerical considerations that otherwise have to be dealt with. The trade off is that at every vertex of the state polytope, all facets of the current Gröbner cone have to be computed which can be highly non-trivial for general ideals, but is relatively easy for toric ideals.

Algorithm 2.0.9. *Enumerating the edge graph of $St(I_A)$ via reverse search.*
Input: *A reduced Gröbner basis $\mathcal{R}_{\succ}$ of I_A and its term order $\succ$.*
Output: *All reduced Gröbner bases of I_A (all vertices of $St(I_A)$).*

```
set G := R_≻; j:=0 ; L := list of facet binomials of G which are marked by ≻
output G
repeat
  while j < |L| do
    j := j + 1.
    G' := flip(G, L[j]);
    if [G',G] in T_≻(I_A) then
      G:= G'; j := 0; L := list of facets of G marked by ≻
      output G
    endif
  endwhile
  If G ≠ R_≻ then
    G' := unique element such that [G,G'] in T_≻(I_A); j := 0; L := list of facets of G' marked by ≻
    repeat j := j + 1 until the common facet of G and G' is the jth facet of L.
    G := G'
  endif
until G = R_≻ and j = |L|.
```

3. Local Computations

Having given an overview of the algorithm in TiGERS for computing the Gröbner fan of a toric ideal, we now focus on the computations that have to be made at a fixed reduced Gröbner basis encountered while searching the edge graph of $St(I_A)$. The two main issues here are (i) how to compute a Gröbner basis that shares a prescribed facet binomial with the current basis and (ii) how to compute all the facets of a toric Gröbner cone.

3.1. Local change of Gröbner bases in toric ideals.

Algorithm 3.1.1. *Local change of reduced Gröbner bases in I_A.*
(a specialization of [CKM] *and Subroutine 3.7 in* [Stu]*).*

Input: (i) A reduced Gröbner basis $\mathcal{G} = \{\underline{x^{a_k}} - x^{b_k} : k = 1, .., t\}$ of I_A.
(The weight vector inducing $\mathcal{G}$ is generic and the underlined terms are the leading terms.)
(ii) A prescribed facet binomial $\underline{x^{a_i}} - x^{b_i}$ of $\mathcal{G}$.
Output: The reduced Gröbner basis adjacent to $\mathcal{G}$ in which $\underline{x^{b_i}} - x^{a_i}$ is a facet binomial.

a. Let $Old := \{\underline{x^{a_i}} - x^{b_i}\} \cup \{\underline{x^{a_j}} : \underline{x^{a_j}} - x^{b_j} \in \mathcal{G}, j \neq i\}$.
b. Let $Temp$ be got from Old by switching the marking on **the** binomial in Old.
   i.e., $Temp := \{\underline{x^{b_i}} - x^{a_i}\} \cup \{\underline{x^{a_j}} : x^{a_j} \in Old\}$.
c. Compute the reduced Gröbner basis of $Temp$ with respect to the new marking.
   Store the (marked) result in New.
d. Set $\mathcal{G}' := \emptyset$
e. For each h in New DO
   (i) Reduce h to zero modulo Old keeping track of the coefficients during reduction. This gives the representation

$$h = \sum_{g \in Old} p_g \cdot g$$

   (ii) Compute the polynomial

$$\sum_{full(g) \in \mathcal{G}} p_g \cdot full(g)$$

   and add it to the set $\mathcal{G}'$. ($full(g) \in \mathcal{G}$ is the completion of $g \in Old$.)
f. Auto-reduce $\mathcal{G}'$ with respect to the markings in New to get $\mathcal{G}_{new}$.

Then $\mathcal{G}_{new}$ is the reduced Gröbner basis adjacent to $\mathcal{G}$ such that $\underline{x^{a_i}} - x^{b_i}$ is a facet binomial in $\mathcal{G}$ and $\underline{x^{b_i}} - x^{a_i}$ is a facet binomial in $\mathcal{G}_{new}$.

Proof of correctness: Let $\mathcal{K}$ and $\mathcal{K}_{new}$ be the Gröbner cones of $\mathcal{G}$ and $\mathcal{G}_{new}$ respectively. The linear span of the common facet of $\mathcal{K}$ and $\mathcal{K}_{new}$ is $\{u \in \mathbf{R}^n : a_i \cdot u = b_i \cdot u\}$. Let $c_1 \in interior(\mathcal{K})$, $c_2 \in interior(\mathcal{K}_{new})$ and c be a vector in the relative interior of $K \cap K_{new}$. By definition, $Old = in_c(\mathcal{G}) := \{in_c(f) : f \in \mathcal{G}\}$ and with respect to the markings specified in Step a, Old is the reduced Gröbner basis of $in_c(I_A)$ with respect to c_1. Since $Temp$ is got from Old by simply reversing the marking on the facet binomial $x^{a_i} - x^{b_i}$, $Temp$ is a generating set for $in_c(I_A)$.

We first show that the set New computed in Step c, using $Temp$ as input, is the reduced Gröbner basis of $in_c(I_A)$ with respect to c_2. Clearly the marked monomials in $Temp$ are their own initial terms with respect to c_2. The non-trivial S-pair computations in Step c are those between a monomial x^{a_j}, $j \neq i$, and the binomial $\underline{x^{b_i}} - x^{a_i}$. This results in a monomial which either reduces to zero with respect to the current partial Gröbner basis or reduces to a monomial that gets added to the current partial Gröbner basis. There is no point during this process at which an unmarked binomial is produced that is required to be marked. No binomials are produced during such an S-pair reduction either. Hence all subsequent S-pair computations are also of the above nature and so in fact, the set New is as claimed above and consists of the binomial $\underline{x^{b_i}} - x^{a_i}$ along with monomials some of which were possibly produced during the Buchberger process. Steps d and e simply lift New to a minimal Gröbner basis $\mathcal{G}'$ of I_A with respect to c_2. Finally, Step f auto-reduces this minimal Gröbner basis to the reduced Gröbner basis $\mathcal{G}_{new}$ of I_A with respect to c_2.

The most important computational advantage of the above local change algorithm is that it does not require the computation of a weight vector in the interior of $\mathcal{K}_{new}$ in order to compute the reduced Gröbner basis $\mathcal{G}_{new}$. This is possible due to the binomial/monomial nature of all intermediate sets produced during the algorithm. The weight vectors are carried implicitly in the markings of these elements. This observation leads to considerable speed up of the procedure.

Remark 3.1.2. *As in the proof of Algorithm 3.1.1, let $c_1 \in interior(\mathcal{K})$ and $c_2 \in interior(\mathcal{K}_{new})$. Then notice that $in_{c_1}(I_A)$ is the initial ideal of $W_{a_i - b_i} := \langle x^{a_i} - x^{b_i}\rangle + \langle x^{a_j} : i \neq j, x^{a_j} \in in_{c_1}(I_A)\rangle$ with respect to the marking $x^{a_i} > x^{b_i}$ while $in_{c_2}(I_A)$ is the initial ideal of $W_{a_i-b_i}$ with respect to $x^{b_i} > x^{a_i}$. See* [MT] *for a generalization of this observation to the theory of* A-graded ideals *(see Chapter 10 in* [Stu]*) and the* toric Hilbert scheme.

Example 3.1.3. For $A := \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 1 \end{pmatrix}$, consider the adjacent reduced Gröbner bases $\mathcal{G}_1$ and $\mathcal{G}_2$ that share the facet binomial $a^2d - be^2$. The basis $\mathcal{G}_1 = \{bd - ce, a^2d - be^2, b^2e - a^2c\}$ has facet binomials $a^2d - be^2$ and $b^2e - a^2c$ and $\mathcal{G}_2 = \{bd - ce, be^2 - a^2d, a^2d^2 - ce^3, b^2e - a^2c\}$ has facet binomials $a^2d^2 - ce^3$ and $be^2 - a^2d$. We use Algorithm 3.1.1 to compute $\mathcal{G}_2$ from $\mathcal{G}_1$.

Step a. $Old := \{a^2d - be^2, bd, b^2e\}$.

Step b. $Temp := \{be^2 - a^2d, bd, b^2e\}$.

Step c. (i) $S\text{-pair}(be^2 - a^2d, bd) = a^2d^2$.
$Temp = \{be^2 - a^2d, bd, b^2e, a^2d^2\}$.
(ii) $S\text{-pair}(be^2 - a^2d, b^2e) = a^2bd \to 0$.
(iii) $S\text{-pair}(be^2 - a^2d, a^2d^2) \to 0$.
Therefore $New = \{be^2 - a^2d, bd, b^2e, a^2d^2\}$.

Step d. $\mathcal{G}' := \emptyset$.

Step e. $be^2 - a^2d = -1(a^2d - be^2)$. Therefore it lifts to $-1(full(a^2d - be^2)) = -1(a^2d - be^2) = be^2 - a^2d$.
$bd = 1(bd)$. Therefore it lifts to $full(bd) = bd - ce$.
$b^2e = 1(b^2e)$. Therefore it lifts to $full(b^2e) = b^2e - a^2c$.
$a^2d^2 = d(a^2d - be^2) + e^2(bd)$. Therefore it lifts to $d(full(a^2d - be^2)) + e^2(full(bd)) = d(a^2d - be^2) + e^2(bd - ce) = a^2d^2 - bde^2 + bde^2 - ce^3 = a^2d^2 - ce^3$.
Hence $\mathcal{G}' = \{be^2 - a^2d, bd - ce, b^2e - a^2c, a^2d^2 - ce^3\}$.

Step f. $\mathcal{G}_{new} = \{be^2 - a^2d, bd - ce, b^2e - a^2c, a^2d^2 - ce^3\}$.

3.2. **Finding the facets of a toric Gröbner cone.** For a general ideal, the Gröbner cone of one of its reduced Gröbner bases is described by a large set of inequalities many of which are redundant. Even for a toric ideal, empirical evidence shows that the number of facet binomials of a reduced Gröbner basis may be much smaller than the cardinality of the Gröbner basis. In fact, it was conjectured in [ST] that there is a function $\phi : \mathbf{N} \to \mathbf{N}$ such that the number of facet binomials of a reduced Gröbner basis of I_A of codimension k is bounded above by $\phi(k)$. As a special case, it was conjectured in [ST] that $\phi(3) = 4$ based on empirical evidence with existing codes at the time. Recently, Serkan Hosten and Diane Maclagan have found counterexamples to this second conjecture using TiGERS. Lower bounds for ϕ are given in [ST], although no good upper bound is known for the number of facets of a toric Gröbner cone. Hence, identifying the facet binomials in a reduced Gröbner basis can become a computationally expensive subroutine during the computation of the Gröbner fan. In this section we discuss several ways to find the facet binomials of a Gröbner cone in the case of a toric ideal.

A first algorithm to compute the facets of a toric Gröbner cone follows from Lemma 2.0.2. In practice this can be an expensive procedure since we need to solve as many linear programs as the cardinality of $\mathcal{G}$ and most binomials in $\mathcal{G}$ are not facet binomials.

Algorithm 3.2.1. *How to find the facet binomials of a reduced Gröbner basis of I_A.*

Input: A reduced Gröbner basis $\mathcal{G} = \{\underline{x}^{a_i} - x^{b_i} : i = 1, .., t\}$ of I_A.
Output: The facet binomials of $\mathcal{G}$.
Set Facets $:= \emptyset$.
For each binomial $\underline{x}^{a_i} - x^{b_i}$ in $\mathcal{G}$ do
If $a_i - b_i$ is not in the cone generated by the vectors $\{a_j - b_j : \underline{x}^{a_j} - x^{b_j} \in \mathcal{G}, i \neq j\}$
then Facets $:=$ Facets $\cup \{x^{a_i} - x^{b_i}\}$.
Output Facets.

Algorithm 3.2.1 is dual to the algorithm suggested by Lemma 2.0.2 since $a_i \cdot u \geq b_i \cdot u$ is a facet inequality of the Gröbner cone $\mathcal{K}$ of $\mathcal{G}$ if and only if $a_i - b_i$ is an extreme ray (essential generator) of $\mathcal{K}^* := \{v \in \mathbf{R}^n : v \cdot u \geq 0, \forall u \in \mathcal{K}\}$ the polar cone of $\mathcal{K}$. The vector $a_i - b_i$ is an extreme ray of $\mathcal{K}^*$ if and only if $a_i - b_i$ cannot be expressed as a non-negative linear combination of the vectors $a_j - b_j, i \neq j$ where $\underline{x}^{a_j} - x^{b_j} \in \mathcal{G}$. Algorithm 3.2.1 can be implemented either by solving one linear program per binomial in $\mathcal{G}$ or finding a generator representation of the cone $\mathcal{K}^*$ using a convex hull package. We also obtain an easy sufficient condition for a binomial in $\mathcal{G}$ to be a facet binomial.

Lemma 3.2.2. *Let $\underline{x}^{a_i} - x^{b_i}$ be an element of a reduced Gröbner basis $\mathcal{G}$ of $I_A \subset k[x_1, \ldots, x_n]$ and x_k be a variable in $k[x_1, \ldots, x_n]$ such that x_k divides the leading term x^{a_i} (respectively, the trailing term x^{b_i}) of $\underline{x}^{a_i} - x^{b_i}$ but does not divide the leading terms (respectively, trailing terms) of any other binomial in $\mathcal{G}$. Then $\underline{x}^{a_i} - x^{b_i}$ is a facet binomial of $\mathcal{G}$.*

Proof: If k is in the support of a_i (respectively b_i) but not in the support of a_j (respectively b_j) for $j \neq i$, $j = 1, \ldots, t$, then $a_i - b_i$ cannot be a non-negative linear combination of $a_j - b_j$ for $j \neq i$, $j = 1, \ldots, t$. Hence $a_i - b_i$ is an extreme ray of the cone polar to the Gröbner cone of $\mathcal{G}$.

We now describe an algorithm to find a superset of the facet binomials of a fixed reduced Gröbner basis of I_A that does not require linear programming. Our idea comes from results in [MT] (c.f. Remark 3.1.2).

Theorem 3.2.3. [MT] *Let $\mathcal{G}_c$ be the reduced Gröbner basis of I_A with respect to the generic weight vector $c \in \mathbf{R}^n$. Then $x^a - x^b \in \mathcal{G}_c$ is a facet binomial of $\mathcal{G}_c$ only if $in_c(I_A)$ is the initial ideal of $W_{a-b} := \langle x^a - x^b \rangle + \langle x^c : x^c \in in_c(I_A), x^c \neq x^a \rangle$ with respect to $x^a > x^b$.*

The exact result in [MT] is that if $x^a - x^b \in \mathcal{G}_c$ and $in_c(I_A)$ is the initial ideal of W_{a-b} with respect to $x^a > x^b$, then the initial ideal M of W_{a-b} with respect to $x^b > x^a$ has the same A-graded Hilbert function as $in_c(I_A)$. The latter is a necessary condition for M to be an initial ideal of I_A. For M to be an adjacent initial ideal to $in_c(I_A)$, you need the additional geometric requirement that M and $in_c(I_A)$ share the facet given by $x^a - x^b$.

Hence the binomials in $\mathcal{G}_c$ that satisfy the condition in Theorem 3.2.3 form a superset of the facet binomials of $\mathcal{G}_c$. Once this superset has been found, we use linear programming as before to identify the true facet binomials.

Algorithm 3.2.4. *How to find a superset of the facet binomials of a reduced Gröbner basis of I_A.*

Input: A reduced Gröbner basis $\mathcal{G}_c$ of I_A.
Output: A superset SS of the facet binomials of $\mathcal{G}_c$.
Set $SS := \emptyset$
For $x^a - x^b \in \mathcal{G}_c$ do
Let $W_{a-b} := \langle x^a - x^b \rangle + \langle x^c \,:\, x^c \in in_c(I_A),\, x^c \neq x^a \rangle$.
If $in_c(I_A)$ is the initial ideal of W_{a-b} with respect to $x^a > x^b$ then $SS := SS$ union $\{x^a - x^b\}$.
Output SS.

We remark that computing the reduced Gröbner basis and hence the initial ideal of W_{a-b} with respect to $x^a > x^b$ proceeds exactly as in Algorithm 3.1.1 and is possible because of the specific monomial/binomial structure of W_{a-b}. Surprisingly it turns out that using Algorithm 3.2.4 can often result in a 50% speed up over using linear programming alone.

4. Computational Experience

The algorithms described in this paper have been implemented in `C` in a program called TiGERS which is available from the authors. In this section we will describe some implementation issues, optimizations and timings. All timings in Table 1 were obtained by running TiGERS on a dual processor Pentium 450 with $1GB$ of RAM. With each example problem, we list the following information about its state polytope: d its dimension, f_0 the number of vertices, f_1 the number of edges, and td the tree depth – the longest chain in the reverse search tree. We also list mn, the cardinality of the largest Gröbner basis computed, mf, the largest number of facets in any Gröbner basis, and md, the highest degree of a binomial appearing in any of the Gröbner bases. Timings are then given both for exhaustive search and reverse search. Those in parentheses are timings using Algorithm 3.2.4 to cut down on the number of linear programs needed to find facets.

We first give a brief description of the examples in Table 1. The first four examples are unrelated: In `pent` the matrix $A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 1 \end{pmatrix}$ the convex hull of whose columns is a pentagon in the plane. The matrix for `V23` is $A = \begin{pmatrix} 2 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 2 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 2 \end{pmatrix}$ whose toric variety is the second Veronese embedding of $\mathbf{P}^2$, while $PV33$ refers to the *pinched Veronese* surface for which $A = \begin{pmatrix} 3 & 2 & 2 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 2 & 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 1 & 0 & 2 & 0 & 1 & 2 & 3 \end{pmatrix}$. The symbol `gti` stands for the generic toric ideal from Example 4.5 in [PS] for which $A = \begin{pmatrix} 20 & 24 & 25 & 31 \end{pmatrix}$. Examples $K5$ and $K6$ are specific instances of Kn, the complete graph on n vertices. The matrix associated to Kn is the node-edge incidence matrix of the graph (see Chapter 9 in [Stu]). The matrix $An := (1\ 2\ 3\ \ldots n)$ and the Gröbner basis elements of the toric ideal

I_{An} correspond to *primitive partition identities* with largest part n (see Chapter 6 in [Stu] for details). The name $Dr \times s$ refers to the $(r+s) \times rs$ node-edge incidence matrix of an undirected bipartite graph with r nodes in one vertex class and s nodes in the other.

The exhaustive search approach required a large amount of memory - in the A9 example for instance, it was using about 600M by the end. In addition to this, the amount of time needed to determine if a vertex has been seen increases as the number of vertices found increases. Even when memory size is not an issue we found that the reverse search approach could end up by being faster than the exhaustive search (see the A9 example).

On the other hand, the reverse search implementation requires us to traverse every edge and then check if it belonged to the reverse search tree by finding the down edge for the new vertex. If the edge used belongs to the reverse search tree, we keep the new vertex, otherwise we discard it knowing that it has been seen before or will be seen again. Thus instead of a large list search we need only find the first mismarked facet binomial to decide whether a vertex should be output. Furthermore, at each node we keep, we must recompute its facet list every time we pass through it. By comparison the exhaustive search algorithm required us to find facets only once per vertex. One trick that we used to mitigate this re-computation of facets was to save vertices and facet information every time we went up in the tree, so as to avoid recomputing the facets already visited. While this approach means that we are no longer storing just one Gröbner basis, the amount of storage required is still quite small, being bounded by the tree depth. In all the examples listed, the reverse search (with caching) ran in under 750K.

We conclude by drawing attention to the last entry in Table 1 denoted as `HM` for which $A = (247\ 248\ 345\ 15)$. This was the original counter example found by Hosten and Maclagan (c.f. page 10), using TiGERS, to the conjecture that the maximum valency of a vertex in the state polytope of a corank three matrix is four [ST].

Name	d	f_0	f_1	td	mn	mf	md	RS	ES
pent	2	8	8	4	4	2	4	0.00	0.00
V23	3	29	45	6	7	4	3	0.02	0.02(0.01)
gti	3	288	467	30	18	4	31	2.30(1.27)	1.50(0.76)
PV33	6	54828	190253	48	36	12	7	(4343.31)	(3731.24)
K5	5	102	255	14	11	5	3	0.36(0.32)	0.28(0.24)
K6	9	195720		56	37	14	4	110111.68(50662.29)	
A4	3	20	31	6	8	4	4	0.01(0.01)	0.01(0.01)
A5	4	114	249	11	14	8	5	0.35(0.28)	0.25(0.16)
A6	5	488	1394	18	20	12	6	6.78(4.39)	4.33(2.37)
A7	6	4073	14800	28	29	18	7	239.80(139.40)	159.96(82.01)
A8	7	25334	111558	41	38	24	8	5010.37(2732.71)	3624(1867.85)
A9	8	206444	1080981	58	49	32	9	127978.46(67565.22)	(71404.29)
A10	9	>578435	-	-	-	-	-	-	-
D2x2	3	108	222	9	10	6	3	0.24(0.20)	0.23(0.16)
D2x3	5	4488	14184	19	20	8	3	171.74(97.48)	124.06(71.02)
D3x3	8	>257057							
HM	3	904	1546	52	40	5	345	96.72(35.80)	76.08(21.53)

TABLE 1. TiGERS Timings

References

[AGK] B. Amrhein, O. Gloor and W. Küchlin, On the walk, *Theoretical Computer Science* **187** (1997) 179–202.

[AF] D. Avis and K. Fukuda, A basis enumeration algorithm for convex hulls and vertex enumeration of arrangements and polyhedra, *Discrete Computational Geometry* **8** (1992) 295–313.

[AL] W. W. Adams and P. Loustaunau, *An Introduction to Gröbner Bases*, American Mathematical Society, 1994.

[BM] D. Bayer and I. Morrison, Standard bases and geometric invariant theory I, *Journal of Symbolic Computation* **6** (1988) 209–217.

[BFS] L.J. Billera, P. Filliman and B. Sturmfels, Constructions and complexity of secondary polytopes, *Advances in Mathematics* **83** (1990) 155–179.

[CKM] S.Collart, M.Kalkbrener and D.Mall, Converting bases with the Gröbner walk, *Journal of Symbolic Computation* **24** (1997) 465-469.

[GKZ] I.M. Gel'fand, M. Kapranov and A. Zelevinsky, *Multidimensional Determinants, Discriminants and Resultants*, Birkhäuser, Boston, 1994.

[CLO] D. Cox, J. Little and D. O'Shea, *Ideals, Varieties, and Algorithms*, Springer-Verlag, New York, Second Edition, 1996.

[Mac] D. Maclagan, Antichains of monomial ideals are finite, *Manuscript*, 1998.

[MT] D. Maclagan and R. Thomas, On the combinatorics of the toric Hilbert scheme, in preparation.

[MR] T. Mora and L. Robbiano, The Gröbner fan of an ideal, *Journal of Symbolic Computation* **6** (1988) 183–208.

[PS] I. Peeva and B. Sturmfels, Generic lattice ideals, Journal of the American Mathematical Soceity **11** (1998) 363–373.

[Stu] B. Sturmfels, Gröbner bases and Convex Polytopes, *American Mathematical Society*, Providence, RI, 1995.

[SST] M. Saito, B. Sturmfels and N. Takayama, *Gröbner Deformations of Hypergeometric Differential Equations*, in preparation.

[ST] B. Sturmfels and R. Thomas, Variation of cost functions in integer programming, *Mathematical Programming* **77** (1997) 357-387.

MSRI, 1000 Centennial Drive, Berkeley, CA 94720
E-mail address: `birk@msri.org`

Dept. of Mathematics, Texas A & M University, College Station, TX 77843
E-mail address: `rekha@math.tamu.edu`