Computing Gröbner Fans of Toric Ideals

Birkett Huber and Rekha R. Thomas

CONTENTS

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 (Toric Gröbner bases Enumeration by Reverse Search) 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 $k[\boldsymbol{x}] := k[x_1, \dots, x_n]$ where k is a field and an ideal $I \subset k[\boldsymbol{x}]$ that is homogeneous with respect to a positive grading

$$\text{degree}(x_i) = \omega_i \in \mathbb{N}\backslash\{0\}.$$

We use $\mathbb{N}$ to denote the set of nonnegative integers. The *initial ideal* of I with respect to a *term order* $\succ$ on $k[\boldsymbol{x}]$ is the *monomial ideal*

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

where $\text{in}_\succ(f)$ 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 $\mathcal{G}_\succ(I) = \{g_1, \dots, g_t\} \subset I$ such that

(i) $\text{in}_\succ(I) = \langle \text{in}_\succ(g_1), \dots, \text{in}_\succ(g_t) \rangle$, and
(ii) for $i \neq j$, no term of g_i is divisible by $\text{in}_\succ(g_j)$.

Reduced Gröbner bases of polynomial ideals can be computed using *Buchberger's algorithm*. See [Adams and Loustaunau 1994] or [Cox et al. 1997] for further details.

Given an arbitrary *weight vector* $c \in \mathbb{R}^n$, and a polynomial $f = \sum k_\alpha x^\alpha \in k[\boldsymbol{x}]$, the *initial term* of f with respect to c is defined to be the sum of all

1058-6458/2000 \$0.50 per page

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 $\mathrm{in}_c(I) := \langle \mathrm{in}_c(f) : f \in I \rangle$. If $\mathrm{in}_c(I)$ is a monomial ideal then c is said to be *generic* for I. It is known that for a term order $\succ$ on $k[\boldsymbol{x}]$, there exists a weight vector $c \in \mathbb{N}^n$ such that $\mathrm{in}_c(I) = \mathrm{in}_\succ(I)$. In this case we say that c represents $\succ$ and it can be shown that $\mathrm{in}_c(I) = \mathrm{in}_\succ(I)$ if and only if $\mathrm{in}_c(g) = \mathrm{in}_\succ(g)$ for each g in $\mathcal{G}_\succ(I)$.

Two weight vectors c_1 and c_2 in $\mathbb{R}^n$ are said to be *equivalent modulo* I whenever $\mathrm{in}_{c_1}(I) = \mathrm{in}_{c_2}(I)$. The set of all weight vectors that are equivalent to $c \in \mathbb{R}^n$ form a relatively open polyhedral cone in $\mathbb{R}^n$, 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 set of all equivalence classes of cost vectors fit together to form a *polyhedral fan* in $\mathbb{R}^n$ called the *Gröbner fan* of I [Mora and Robbiano 1988]. Since I is homogeneous with respect to a positive grading, this fan is, in fact, *complete* (i.e., covers $\mathbb{R}^n$) and further, each Gröbner cone of I contains a strictly positive vector of $\mathbb{R}^n$ in its relative interior.

The Gröbner fan of I is the *normal fan* of a polytope in $\mathbb{R}^n$ called the *state polytope* of I [Bayer and Morrison 1988], denoted as $\mathrm{St}(I)$. Therefore, I has only finitely many distinct reduced Gröbner bases as c varies over all weight vectors in $\mathbb{R}^n$. (See [Maclagan 1998] for a new proof.) The faces of $\mathrm{St}(I)$ are in bijection with the distinct initial ideals of I, with the vertices of $\mathrm{St}(I)$ corresponding bijectively to the distinct monomial initial ideals of I. The distinct monomial initial ideals of I, in turn, are in bijection with the distinct reduced Gröbner bases of I obtained from term orders. Hence, computing all monomial initial ideals (or reduced Gröbner bases from term orders) of I amounts to searching the *edge graph* of $\mathrm{St}(I)$. Two monomial initial ideals $\mathrm{in}_{c_1}(I)$ and $\mathrm{in}_{c_2}(I)$ are said to be *adjacent* if the corresponding vertices of $\mathrm{St}(I)$ are adjacent or equivalently, if the Gröbner cones of the generic weight vectors c_1 and c_2 share a common *facet*. See [Sturmfels 1996, Chapters 1–3] 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 $A = [a_1 \cdots a_n] \in \mathbb{Z}^{d \times n}$ of rank d, the *toric ideal* of A, denoted as I_A, is the kernel of the homomorphism $k[x_1, \dots, x_n] \to k[t_1^{\pm 1}, \dots, t_d^{\pm 1}]$, such that $x_j \mapsto t^{a_j}$ [Sturmfels 1996, Chapter 4]. 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

$$\ker_{\mathbb{Z}}(A) := \{u \in \mathbb{Z}^n : Au = 0\}.$$

Here $u^- = (-u)^+$ and u^+ is defined as $u_i^+ = u_i$ if $u_i > 0$ and $u_i^+ = 0$ otherwise. Hence $u^+, u^- \in \mathbb{N}^n$. 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 $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^n = \{0\}$, which guarantees a positive integral vector ω in the row space of A. Then I_A is homogeneous with respect to the grading $\mathrm{degree}(x_i) = \omega_i$ for $i = 1, \dots, n$.

Let $\mathcal{G}_c = \{x^{\alpha_i} - x^{\beta_i} : i = 1, \dots, t\}$ be the reduced Gröbner basis of I_A with respect to a generic weight vector $c \in \mathbb{R}^n$. (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

$$\mathcal{K}_c := \{u \in \mathbb{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u,\ i = 1, \dots, t\},$$

whose lineality space $\mathcal{K}_c \cap -\mathcal{K}_c$ is the row space of A. The equivalence class of c is the interior $\mathring{\mathcal{K}}_c$ of $\mathcal{K}_c$; in particular, c lies in $\mathring{\mathcal{K}}_c$. We may assume that c is a strictly positive integral vector since $\mathcal{K}_c$ is a rational cone and $\omega \in \mathcal{K}_c$. The weight vector c' is equivalent to c if and only if $\mathrm{in}_{c'}(x^{\alpha_i} - x^{\beta_i}) = x^{\alpha_i}$ for each binomial $x^{\alpha_i} - x^{\beta_i} \in \mathcal{G}_c$. State polytopes and Gröbner fans of toric ideals were studied in [Sturmfels and Thomas 1997]. That paper gives several custom-tailored construction methods for these entities.

The Gröbner fan of I_A has various applications: In [Sturmfels and Thomas 1997] it was used as a model for *sensitivity analysis* for the family of *integer programs* $min\{c \cdot x : Ax = b,\ x \in \mathbb{N}^n\}$ as b varies in $\mathbb{Z}^d$ and c in $\mathbb{R}^n$. The *secondary fan* of A [Billera et al. 1990; Gel'fand et al. 1994] is a coarsening of the Gröbner fan of I_A [Sturmfels 1996, Chapter 8]. This fan has important applications 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* [Saito et al. 2000]. Both these coarser fans can be obtained from the Gröbner fan of I_A. Finally, the *Gröbner walk* procedure introduced in [Collart et al. 1997] uses the Gröbner fan for Gröbner basis conversions.

Example 1.1. 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 $\mathrm{St}(I_A)$ is an octagon in $\mathbb{R}^5$ and the Gröbner fan of I_A can be drawn in $\mathbb{R}^2$ 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.

In this paper we introduce a software system called TiGERS (Toric Gröbner bases Enumeration by Reverse Search) for computing the Gröbner fan of a toric ideal. The program searches the edge graph of the state polytope $\mathrm{St}(I_A)$ to find all the distinct monomial initial ideals (reduced Gröbner bases) of I_A. (For now on, when we refer to a reduced Gröbner basis of I_A we assume that it has been computed using a generic weight vector and hence is indexed by a vertex of $\mathrm{St}(I_A)$.) 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 $\mathrm{St}(I_A)$. This approach needs to maintain and search the list of all Gröbner bases found. When $\mathrm{St}(I_A)$ is too large this approach will bog down and we use a *reverse search technique* instead. This involves a depth-first search on a directed subgraph of the edge graph of $\mathrm{St}(I_A)$ 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

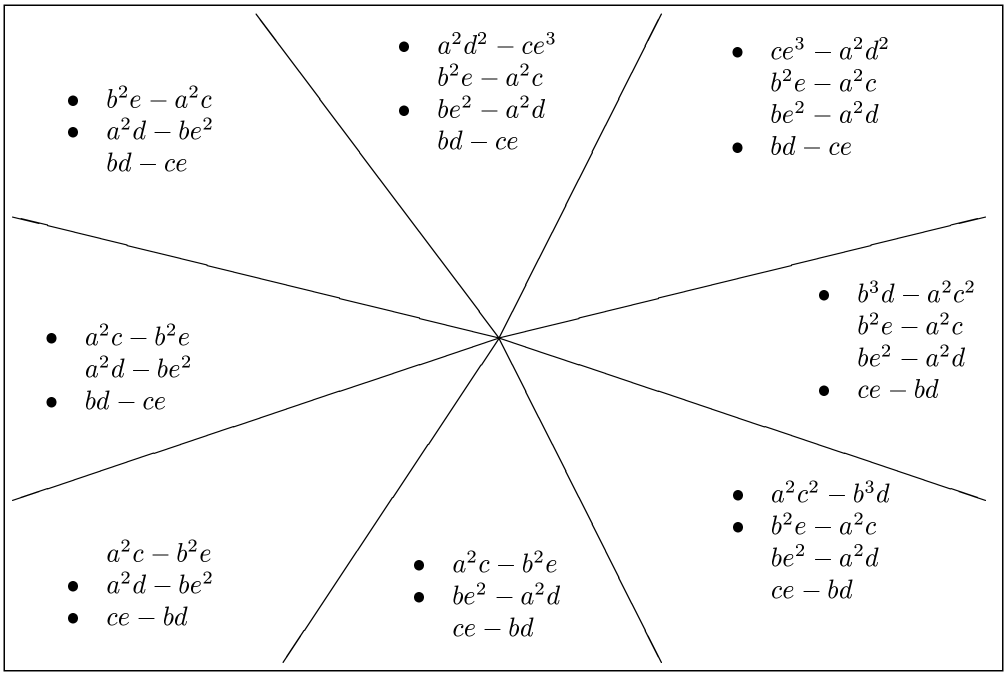

FIGURE 1. A schematic of the Gröbner fan in Example 1.1.

on [Collart et al. 1997]. 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 $c \in \mathbb{R}^n$. An inequality $\alpha_k \cdot u \geq \beta_k \cdot u$, with $k \in \{1, \ldots, t\}$, is irredundant for the Gröbner cone

$$\mathcal{K}_c = \{u \in \mathbb{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u \text{ for } i = 1, \ldots, t\}$$

if the relaxed cone $\{u \in \mathbb{R}^n : \alpha_i \cdot u \geq \beta_i \cdot u$ for $i = 1, \ldots, t,\ i \neq k\}$ properly contains $\mathcal{K}_c$. If $\alpha_k \cdot u \geq \beta_k \cdot u$ is irredundant for $\mathcal{K}_c$ then

$$\mathcal{K}_c \cap \{u \in \mathbb{R}^n : \alpha_k \cdot u = \beta_k \cdot u\}$$

is called a *facet* of $\mathcal{K}_c$ and $x^{\alpha_k} - x^{\beta_k}$ is a *facet binomial* of $\mathcal{G}_c$.

Lemma 2.1 [Sturmfels and Thomas 1997]. *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 $\{\alpha_i \cdot u > \beta_i \cdot u : i = 1, \ldots, t,\ i \neq k,\ u \in \mathbb{R}^n\} \cap \{\beta_k \cdot u > \alpha_k \cdot u : u \in \mathbb{R}^n\}$ 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 $\mathrm{in}_c(x^\alpha - x^\beta) = x^\alpha$ and $\mathcal{G}_{c'}$ be adjacent to $\mathcal{G}_c$ with $\mathrm{in}_{c'}(x^\alpha - x^\beta) = x^\beta$. In order to compute the edge graph of $\mathrm{St}(I_A)$, 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 $\mathrm{flip}(\mathcal{G}_c, x^\alpha - x^\beta)$ to denote both the subroutine that yields $\mathcal{G}_{c'}$ from $\mathcal{G}_c$ and 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 2.2. 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 $J = \langle b^3d - b^2ce, a^2c - b^2e, bcd - c^2e \rangle$. 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 $\mathrm{in}_\omega(ceb^2 - db^3) = ceb^2 - db^3$ and $\mathrm{in}_\omega(c^2e - cdb) = c^2e - cdb$.

Algorithm 2.3 (Enumerating the edge graph of $\mathrm{St}(I_A)$ 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 $\mathrm{St}(I_A)$).

```
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 $\mathrm{St}(I_A)$ 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, $\mathrm{St}(I_A)$ has 54,828 vertices and our breadth first search algorithm was exhausting a personal computer with 64 megabytes of memory before getting through about 13,000 vertices.

To push this calculation through we resorted to the reverse enumeration paradigm of Avis and Fukuda [1992], a memoryless algorithm that runs in linear time in the size of the output. It was originally formulated for enumerating the vertices of the edge graphs of polytopes. Suppose G is the edge graph of a polytope P and c a generic cost vector such that no two vertices of P have the same cost value $c \cdot x$. Let v_0 be the unique vertex of P at which $c \cdot x$ is maximized. We fix a *pivot rule* on the vertices of P that assigns a unique neighbor to a vertex v from among all neighboring vertices of v with higher cost value than v. Together, these create a directed graph on the vertices of P with a unique sink at v_0 and out degree one for every other vertex. The basic idea in the Avis-Fukuda algorithm is to start at v_0 and do a depth first reverse search on this directed subgraph. No intermediate storage of vertices is required during the search since a vertex seen once will not be encountered again as one traverses down a branch in the directed tree. See Algorithm 2.8 for details. This algorithm has been used with great success in several applications such as enumerating all vertices of a polytope, all regular triangulations of a point configuration [Masada et al. 1996], all cells in a hyperplane arrangement and all lattice points in a polytope [Sturmfels 1996, Algorithm 5.7].

We say that a polynomial $f \in k[\boldsymbol{x}]$ is *marked* by a term order $\succ$ on $k[\boldsymbol{x}]$, 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 $\mathrm{in}_{\succ}(f) \neq \mathrm{in}_{\succ'}(f)$.

Lemma 2.4. *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 $\mathcal{G}_\succ = \mathcal{G}_{\succ'}$, no binomial in $\mathcal{G}_\succ$ is mismarked with respect to $\succ'$. □

Definition 2.5. 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 $\mathrm{St}(I_A)$ (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 $\mathrm{flip}(\mathcal{G}_i, x^\alpha - x^\beta)$ 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 2.6. *The reverse search tree $T_\succ(I_A)$ is an acyclic directed graph with a unique sink.*

Proof. By Lemma 2.4 each reduced Gröbner basis $\mathcal{G}$ of I_A (vertex of $\mathrm{St}(I_A)$) 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 $\mathrm{flip}(\mathcal{G},\ x^\alpha - x^\beta)$ 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 $C = (v_1, \dots, v_l)$ 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 in $\mathcal{G}_i$ and x^{β_i} the leading term in $\mathcal{G}_{i+1}$ for $i = 1, \dots, l$, then $v_{i+1} - v_i = \beta_i - \alpha_i$ for $i = 1, \dots, l$. 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 $i = 1, \dots, l$ 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 2.7. *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 $\mathcal{G}_\succ$.*

This implies that toric ideals admit combinatorial Gröbner walks that can be used for converting one reduced Gröbner basis of I_A into another. Given any

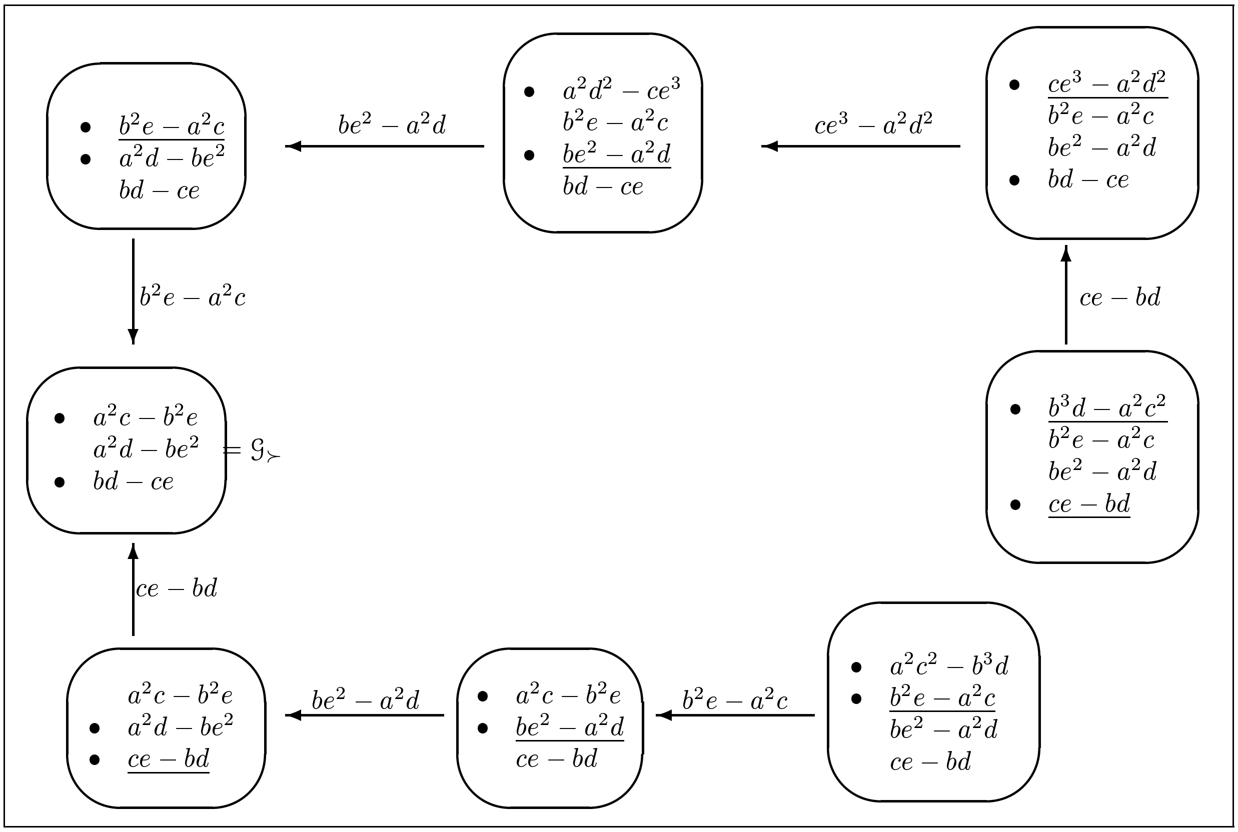

FIGURE 2. $T_\succ(I_A)$ with A as in Example 1.1 using pure lex for $\succ$. Mismarked facet binomials are underlined and arrows are labeled by the binomials which are being flipped.

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 [Collart et al. 1997; Amrhein et al. 1997], 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 no 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 tradeoff is that at every vertex of the state polytope, all facets of the current Gröbner cone have to be computed, which can be highly nontrivial for general ideals, but is relatively easy for toric ideals.

Algorithm 2.8 (Enumerating the edge graph of $\mathrm{St}(I_A)$ 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 $\mathrm{St}(I_A)$).

$\mathcal{G} := \mathcal{R}_\succ$; $j := 0$;
$L :=$ list of facet binomials of $\mathcal{G}$ marked by $\succ$
output $\mathcal{G}$
repeat
  while $j < |L|$ do
    $j := j + 1$
    $\mathcal{G}' := \mathrm{flip}(\mathcal{G}, L[j])$;
    if $[\mathcal{G}', \mathcal{G}]$ in $T_\succ(I_A)$ then
      $\mathcal{G} := \mathcal{G}'$; $j := 0$;
      $L :=$ list of facets of $\mathcal{G}$ marked by $\succ$
      output $\mathcal{G}$
    endif
  endwhile
  If $\mathcal{G} \neq \mathcal{R}_\succ$ then
    $\mathcal{G}' :=$ unique element such that $[\mathcal{G}, \mathcal{G}']$ in $T_\succ(I_A)$;
    $j := 0$; $L :=$ list of facets of $\mathcal{G}'$ marked by $\succ$
    repeat $j := j + 1$ until the common facet of $\mathcal{G}$ and $\mathcal{G}'$ is the j-th facet of L.
    $\mathcal{G} := \mathcal{G}'$
  endif
until $\mathcal{G} = \mathcal{R}_\succ$ and $j = |L|$.

3. LOCAL COMPUTATIONS

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

3A. Local Change of Gröbner Bases in Toric Ideals

The following algorithm is a specialization of [Collart et al. 1997] and [Sturmfels 1996, Subroutine 3.7].

Algorithm 3.1 (Local change of reduced Gröbner bases in I_A).

Input: (i) A reduced Gröbner basis $\mathcal{G} = \{\underline{x^{a_k}} - x^{b_k} : k = 1, \dots, t\}$ 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. 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 obtained from Old by switching the marking on *the* binomial in Old, that is, Temp $:= \{\underline{x^{b_i}} - x^{a_i}\} \cup \{\underline{x^{a_j}} : x^{a_j} \in \text{Old}\}$.

c. Compute the reduced Gröbner basis of Temp with respect to the new marking. Store the (marked) result in New.

d. $\mathcal{G}' := \{\underline{x^{b_i}} - x^{a_i}\}$.

e. For each monomial h in New, do:
Reduce h with $\mathcal{G}$ to obtain the monomial h'.
Add $h - h'$ to $\mathcal{G}'$ with h marked as the leading term.

f. Auto-reduce $\mathcal{G}'$ to get $\mathcal{G}_{\text{new}}$.
Then $\mathcal{G}_{\text{new}}$ 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 $\mathcal{G}_{\text{new}}$.

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

We first show that the set New computed in Step c, using Temp as input, is the reduced Gröbner basis of $\mathrm{in}_c(I_A)$ with respect to c_2. The marked monomials in Temp are the leading terms with respect to c_2 of each polynomial in Temp. The nontrivial 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. Step e lifts New to the set of marked binomials $\mathcal{G}'$, the leading terms of whose elements are precisely the minimal generators of $\mathrm{in}_{c_2}(\mathrm{in}_c(I_A)) = \mathrm{in}_{c+\varepsilon c_2}(I_A)$ for some small $\varepsilon > 0$. However, $c + \varepsilon c_2$ lies in the interior of $\mathcal{K}_{\text{new}}$ and hence, $\mathrm{in}_{c_2}(\mathrm{in}_c(I_A)) = \mathrm{in}_{c_2}(I_A)$. Hence $\mathcal{G}'$ is a minimal Gröbner basis of I_A with respect to c_2. Step f then auto-reduces this minimal Gröbner basis to the reduced Gröbner basis $\mathcal{G}_{\text{new}}$ 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 $\mathcal{K}_{\text{new}}$ in order to compute the reduced Gröbner basis $\mathcal{G}_{\text{new}}$. This is possible due to the binomial/monomial nature of all intermediate polynomials 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 3.2. As in the proof of Algorithm 3.1, let $c_1 \in \mathring{\mathcal{K}}$ and $c_2 \in \mathring{\mathcal{K}}_{\text{new}}$. Then notice that $\mathrm{in}_{c_1}(I_A)$ is the initial ideal of $W_{a_i - b_i} := \langle x^{a_i} - x^{b_i} \rangle + \langle x^{a_j} : i \neq j,\ x^{a_j}$ minimal generator of $\mathrm{in}_{c_1}(I_A) \rangle$ with respect to the marking $x^{a_i} > x^{b_i}$ while $\mathrm{in}_{c_2}(I_A)$ is the initial ideal of $W_{a_i-b_i}$ with respect to $x^{b_i} > x^{a_i}$. See [Maclagan and Thomas 1999] for a generalization of this observation to the theory of *A-graded ideals* [Sturmfels 1996, Chapter 10] and the *toric Hilbert scheme*.

Example 3.3. 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 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-pair$(be^2 - a^2d, bd) = a^2d^2$. Temp $= \{be^2 - a^2d, bd, b^2e, a^2d^2\}$. (ii) S-pair$(be^2 - a^2d, b^2e) = a^2bd \to 0$. (iii) S-pair$(be^2 - a^2d, a^2d^2) \to 0$. Therefore New $= \{be^2 - a^2d, bd, b^2e, a^2d^2\}$.
Step d. $\mathcal{G}' := \{\underline{be^2} - a^2d\}$.
Step e. bd reduces modulo $\mathcal{G}_1$ to ce. Therefore we add $\underline{bd} - ce$ to $\mathcal{G}'$.
b^2e reduces to a^2c modulo $\mathcal{G}_1$. Therefore we add $\underline{b^2e} - a^2c$ to $\mathcal{G}'$.
$a^2d^2 = d(a^2d - be^2) + e^2(bd - ce) + ce^3$. Therefore we add $\underline{a^2d^2} - ce^3$ to $\mathcal{G}'$.
Hence $\mathcal{G}' = \{be^2 - a^2d, bd - ce, b^2e - a^2c, a^2d^2 - ce^3\}$, with the positive terms as the leading terms.
Step f. $\mathcal{G}_{\text{new}} = \{be^2 - a^2d, bd - ce, b^2e - a^2c, a^2d^2 - ce^3\}$.

3B. 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 [Sturmfels and Thomas 1997] that there is a function $\varphi : \mathbb{N} \to \mathbb{N}$ such that the number of facet binomials of a reduced Gröbner basis of I_A of codimension k is bounded above by $\varphi(k)$. As a special case, it was conjectured in [Sturmfels and Thomas 1997] that $\varphi(3) = 4$ 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 [Sturmfels and Thomas 1997], 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 2.1. 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 3.4 (Finding the facet binomials of a reduced Gröbner basis of I_A).
Input: A reduced Gröbner basis $\mathcal{G} = \{\underline{x}^{a_i} - x^{b_i} : i = 1, \dots, t\}$ of I_A.
Output: The facet binomials of $\mathcal{G}$.

Facets $:= \varnothing$.
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 $\{a_j - b_j : \underline{x}^{a_j} - x^{b_j} \in \mathcal{G}, i \neq j\}$, set Facets := Facets $\cup \{x^{a_i} - x^{b_i}\}$.
Output Facets.

Algorithm 3.4 is dual to the algorithm suggested by Lemma 2.1 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 $\mathcal{K}^* := \{v \in \mathbb{R}^n : v \cdot u \geq 0, \forall u \in \mathcal{K}\}$ 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 nonnegative linear combination of the vectors $a_j - b_j$, $i \neq j$ where $\underline{x}^{a_j} - x^{b_j} \in \mathcal{G}$. Algorithm 3.4 can be implemented by solving one linear program per binomial in $\mathcal{G}$ or by finding the generators 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 3.5. *Let $\underline{x}^{a_i} - x^{b_i}$ be an element of a reduced Gröbner basis $\mathcal{G}$ of $I_A \subset k[x_1, \dots, x_n]$ and x_k be a variable in $k[x_1, \dots, x_n]$ 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 (re-*

spectively, 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 nonnegative 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 [Maclagan and Thomas 1999] (compare Remark 3.2).

Theorem 3.6 [Maclagan and Thomas 1999]. *Let* $\mathcal{G}_c$ *be the reduced Gröbner basis of* I_A *with respect to the generic weight vector* $c \in \mathbb{R}^n$. *Then* $x^a - x^b \in \mathcal{G}_c$ *is a facet binomial of* $\mathcal{G}_c$ *only if* $\mathrm{in}_c(I_A)$ *is the initial ideal of* $W_{a-b} := \langle x^a - x^b \rangle + \langle x^c : x^c$ *is a minimal generator of* $\mathrm{in}_c(I_A)$, $x^c \neq x^a \rangle$ *with respect to* $x^a > x^b$.

The exact result in [Maclagan and Thomas 1999] is that if $x^a - x^b \in \mathcal{G}_c$ and $\mathrm{in}_c(I_A)$ 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 $\mathrm{in}_c(I_A)$. 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 $\mathrm{in}_c(I_A)$, you need the additional geometric requirement that M and $\mathrm{in}_c(I_A)$ share the facet given by $x^a - x^b$. Hence the binomials in $\mathcal{G}_c$ that satisfy the condition in Theorem 3.6 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 3.7 (Finding 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 $\mathcal{G}_c$.
$SS := \varnothing$
For $x^a - x^b \in \mathcal{G}_c$, do:
  Set $W_{a-b} := \langle x^a - x^b \rangle + \langle x^c : x^c$ minimal generator of $\mathrm{in}_c(I_A)$, $x^c \neq x^a \rangle$.
  If $\mathrm{in}_c(I_A)$ is the initial ideal of W_{a-b} with respect to $x^a > x^b$, set $SS := SS \cup \{x^a - x^b\}$.
Output SS.

The computation of the reduced Gröbner basis and hence of the initial ideal of W_{a-b} with respect to $x^a > x^b$ proceeds exactly as in Algorithm 3.1 and is possible because of the specific monomial/binomial structure of W_{a-b}. Surprisingly, the use of Algorithm 3.7 can often result in a 50% speed-up over using linear programming alone.

4. COMPUTATIONAL EXPERIENCE

TiGERS, the program that implements the algorithms described in this paper, is written in C and is available from http://www.math.washington.edu/~thomas/programs.html. In this section we describe some implementation issues, optimizations and timings. All timings in Table 1 were obtained by running TiGERS on a dual processor Pentium 450 with 1 GB 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 m_n, the cardinality of the largest Gröbner basis computed, m_f, 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 reverse search (RS) and exhaustive search (ES). Timings computed using Algorithm 3.7 to cut down on the use of linear programming are given in parentheses.

We first give a brief description of the examples in Table 1. The first four examples are unrelated: In Pent the matrix is

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 1 & 0 \\ 0 & 0 & 1 & 2 & 1 \end{pmatrix},$$

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 $\mathbb{P}^2$. The name 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 name gti stands for the generic toric ideal from [Peeva and Sturmfels 1998, Example 4.5], for which $A = (20\ 24\ 25\ 31)$. 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 [Sturmfels 1996, Chapter 9]. The matrix A$n := (1\ 2\ 3\ \dots n)$ and the Gröbner basis elements of the toric ideal I_{An} correspond to *primitive partition identities* with largest part n (see [Sturmfels 1996, Chapter 6] for details). The name D$r \times s$ 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 last example, HM, is explained at the end of this section.

The exhaustive search approach required a large amount of memory — in the A9 example, for instance, it was using about 600MB 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 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 750KB.

The last line on the table, HM, is Hosten and Maclagan's original counterexample, found using TiGERS, to the conjecture that the maximum valency of a vertex in the state polytope of a corank three matrix is four [Sturmfels and Thomas 1997]. It corresponds to $A = (247\ 248\ 345\ 15)$.

Example	d	f_0	f_1	td	m_n	m_f	m_d	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							
D2×2	3	108	222	9	10	6	3	0.24 (0.20)	0.23 (0.16)
D2×3	5	4488	14184	19	20	8	3	171.74 (97.48)	124.06 (71.02)
D3×3	8	> 257057							
HM	3	904	1546	52	40	5	345	96.72 (35.80)	76.08 (21.53)

TABLE 1. TiGERS performance. The middle columns give the characteristics of the state polytope (see beginning of Section 4) and the two rightmost columns give timings, in seconds, for the reverse search and exhaustive search.

ELECTRONIC AVAILABILITY

TiGERS, the program that implements the algorithms described in this paper, is written in C and is available from http://www.math.washington.edu/~thomas/programs.html.

REFERENCES

[Adams and Loustaunau 1994] W. W. Adams and P. Loustaunau, *An introduction to Gröbner bases*, Graduate Studies in Math. **3**, Amer. Math. Soc., Providence, 1994.

[Amrhein et al. 1997] B. Amrhein, O. Gloor, and W. Küchlin, "On the walk", *Theoret. Comput. Sci.* **187**:1-2 (1997), 179–202.

[Avis and Fukuda 1992] D. Avis and K. Fukuda, "A basis enumeration algorithm for convex hulls and vertex enumeration of arrangements and polyhedra", *Disc. Comput. Geom.* **8** (1992), 295–313.

[Bayer and Morrison 1988] D. Bayer and I. Morrison, "Standard bases and geometric invariant theory, I: Initial ideals and state polytopes", *J. Symbolic Comput.* **6**:2-3 (1988), 209–217.

[Billera et al. 1990] L. J. Billera, P. Filliman, and B. Sturmfels, "Constructions and complexity of secondary polytopes", *Adv. Math.* **83**:2 (1990), 155–179.

[Collart et al. 1997] S. Collart, M. Kalkbrener, and D. Mall, "Converting bases with the Gröbner walk", *J. Symbolic Comput.* **24**:3-4 (1997), 465–469.

[Cox et al. 1997] D. Cox, J. Little, and D. O'Shea, *Ideals, varieties, and algorithms*, 2nd ed., Springer, New York, 1997.

[Gel'fand et al. 1994] I. M. Gel'fand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants, resultants, and multidimensional determinants*, Birkhäuser, Boston, 1994.

[Maclagan 1998] D. Maclagan, "Antichains of monomial ideals are finite", preprint, 1998. See http://arxiv.org/abs/math.CO/9909168/.

[Maclagan and Thomas 1999] D. Maclagan and R. Thomas, "Combinatorics of the toric Hilbert scheme", preprint, 1999. See http://arxiv.org/abs/math.AG/9912014/.

[Masada et al. 1996] T. Masada, H. Imai, and K. Imai, "Enumeration of regular triangulations", pp. 224–233 in *Proc. 12th annual Symposium on Computational Geometry* (Philadelphia, PA), ACM Press, New York, 1996.

[Mora and Robbiano 1988] T. Mora and L. Robbiano, "The Gröbner fan of an ideal", *J. Symbolic Comput.* **6**:2-3 (1988), 183–208.

[Peeva and Sturmfels 1998] I. Peeva and B. Sturmfels, "Generic lattice ideals", *J. Amer. Math. Soc.* **11**:2 (1998), 363–373.

[Saito et al. 2000] M. Saito, B. Sturmfels, and N. Takayama, *Gröbner deformations of hypergeometric differential equations*, Algorithms and Computation in Mathematics **6**, Springer, Berlin, 2000.

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

[Sturmfels and Thomas 1997] B. Sturmfels and R. R. Thomas, "Variation of cost functions in integer programming", *Math. Programming* **77**:3, Ser. A (1997), 357–387.

Birkett Huber, Institute for Defense Analysis, United States (birk@member.ams.org)

Rekha R. Thomas, Rekha R. Thomas, Department of Mathematics, University of Washington, Box 354350, Seattle, WA 98195, United States (thomas@math.washington.edu)

Received June 8, 1999; accepted in revised form September 27, 1999