Gröbner bases and integer programming

Serkan Hoşten and Rekha Thomas

1 Introduction

This article is a brief survey of recent work on Gröbner bases (Buchberger 1965) of toric ideals and their role in integer programming. Toric varieties and ideals are crucial players in the interaction between combinatorics, discrete geometry, commutative algebra and algebraic geometry. For a detailed treatment of this topic see (Sturmfels 1995). Our survey focuses on the application of toric ideals to integer programming, a specific branch of discrete optimization, and for the sake of brevity we leave details to the references that are included.

We study a family of integer programs associated with a fixed matrix A and the corresponding toric ideal I_A. In Section 2, we show that reduced Gröbner bases of I_A are *test sets* for these integer programs. In Section 3, we define the *universal Gröbner basis* of I_A and we identify it as a subset of the *Graver basis* of A. Section 4 deals with the effect of varying the cost function while keeping the matrix A fixed; there we give a self contained construction of the *state polytope* and *Gröbner fan* of I_A. Moreover, we show that the edge directions of the state polytope are precisely the elements in the universal Gröbner basis of I_A. We conclude in Section 5 with a discussion of practical issues that arise while computing Gröbner bases of toric ideals. In particular we discuss algorithms for finding generating sets for I_A.

2 Gröbner bases as test sets in integer programming

Let $IP_{A,c}(b)$ denote the integer program

$$\text{minimize } \ cx : \ Ax = b, \ \ x \in \mathbf{N}^n$$

where the *coefficient matrix* $A \in \mathbf{Z}^{d\times n}$ has rank d, the *right hand side vector* $b \in \mathbf{Z}^d$ and the *cost vector* $c \in \mathbf{R}^n$. The program $IP_{A,c}(b)$ has a solution if and only if b lies in the additive monoid $\mathbf{N}(A) := \{Ax \ : \ x \in \mathbf{N}^n\}$. We study $IP_{A,c}$, the family of integer programs of the form $IP_{A,c}(b)$ obtained by varying b in $\mathbf{N}(A)$, while keeping A and c fixed. We call the polyhedron

$P_b^I = \mathit{convex\,hull}\{x \in \mathbf{N}^n : Ax = b\}$, the b-fiber of A and the minimum value of $c \cdot x$ over all $x \in P_b^I$, the *optimal value* of $IP_{A,c}(b)$. For simplicity, we assume that the *recession cone* of A, $\{x \in \mathbf{R}_+^n : Ax = 0\} = \{0\}$, which guarantees that P_b^I is a polytope for each $b \in \mathbf{N}(A)$. Hence every cost vector $c \in \mathbf{R}^n$ attains a bounded optimal value on each program in $IP_{A,c}$.

The *toric ideal* of A, denoted I_A, is the prime ideal in the polynomial ring $S = k[x_1, \ldots, x_n]$ generated by monomial differences (binomials) as follows:

$$I_A = \langle \mathbf{x}^\alpha - \mathbf{x}^\beta : \alpha, \beta \in \mathbf{N}^n, A(\alpha - \beta) = 0 \rangle.$$

Proposition 2.1 *Any reduced Gröbner basis of the toric ideal I_A consists of binomials.*

Proof: By Hilbert's basis theorem we can find a finite binomial generating set for I_A. If the input to the Buchberger algorithm is a set of binomials, then all intermediate polynomials created, as well as the final output, consists of binomials. □

The cost vector c typically induces a partial order on $\mathbf{N}^n$ and hence on the monomials in S via the inner product $c \cdot x$. In this situation, we pick a term order $\succ$ on $\mathbf{N}^n$ like say the lexicographic order and we replace c in $IP_{A,c}$ by the total order $\succ_c$ defined as:

$$\mathbf{x}^\alpha \succ_c \mathbf{x}^\beta \text{ if } c \cdot \alpha > c \cdot \beta, \quad \text{or if } \ c \cdot \alpha = c \cdot \beta, \ \text{ and } \ \mathbf{x}^\alpha \succ \mathbf{x}^\beta. \tag{2.1}$$

In general $\succ_c$ is not a term order on $\mathbf{N}^n$ but since all P_b^Is are bounded with respect to c, $\succ_c$ behaves like a term order. Our primary object of study is the reduced Gröbner basis $\mathcal{G}_{\succ_c}$ of the toric ideal I_A with respect to the total order $\succ_c$.

Definition 2.2 *A set $\mathcal{T} \subseteq \{t \in \mathbf{Z}^n : At = 0, t \succ_c 0\}$ is a* **test set** *for $IP_{A,\succ_c}$ if*

1. *given any non-optimal solution α to any program in $IP_{A,\succ_c}$, there exists $t \in \mathcal{T}$ such that $\alpha - t$ is a feasible solution to the same program, and*
2. *for the optimal solution β to a program in $IP_{A,\succ_c}$, $\beta - t$ is infeasible for every t in $\mathcal{T}$.*

Various test sets can be found in the integer programming literature (see Scarf (1986), Graver (1975), and Cook et al. (1986)). Whenever a test set $\mathcal{T}$ is known for $IP_{A,\succ_c}$, there is a natural optimization algorithm which finds the unique optimal solution to any program in $IP_{A,\succ_c}$: given a feasible solution α for $IP_{A,\succ_c}(b)$, we *test* whether $\alpha - t$ is non-negative for some $t \in \mathcal{T}$. If such a t exists we replace α with $\alpha - t$ which has a better cost value since $t \succ_c 0$ for all

$t \in \mathcal{T}$. We iterate until no $t \in \mathcal{T}$ can improve the current feasible solution β. Then, since $\mathcal{T}$ is a test set for $IP_{A,\succ_c}$ we may conclude that β is the optimal solution of $IP_{A,\succ_c}(b)$. The following theorem connects reduced Gröbner bases of I_A and test sets for $IP_{A,\succ_c}$.

Theorem 2.3 *The reduced Gröbner basis $\mathcal{G}_{\succ_c}$ of I_A is a minimal test set for $IP_{A,\succ_c}$.*

Proof: Let $u \in \mathbf{N}^n$ be a non-optimal solution for the integer program $IP_{A,\succ_c}(Au)$ and let $v \in \mathbf{N}^n$ be the optimal solution for the same program. If we let $u - v = (u-v)_+ - (u-v)_-$, the binomial $\mathbf{x}^{(u-v)_+} - \mathbf{x}^{(u-v)_-}$ is in I_A and $\mathbf{x}^{(u-v)_+}$ is the initial term of this binomial with respect to $\succ_c$. This means there is a Gröbner basis element $\mathbf{x}^\alpha - \mathbf{x}^\beta$ such that the initial term $\mathbf{x}^\alpha$ divides $\mathbf{x}^{(u-v)_+}$. Therefore $u - (\alpha - \beta)$ is an improved solution for $IP_{A,\succ_c}(Au)$. By the same reasoning the optimal solution v cannot be improved by any reduced Gröbner basis element. This shows that $\mathcal{G}_{\succ_c}$ is a test set. To prove minimality, we have to show that no Gröbner basis element can be dropped from this test set. For any reduced Gröbner basis element $\mathbf{x}^\alpha - \mathbf{x}^\beta$ the non-optimal solution α for the integer program $IP_{A,\succ_c}(A\alpha)$ cannot be improved by any other element of $\mathcal{G}_{\succ_c}$ since $\mathcal{G}_{\succ_c}$ is a reduced Gröbner basis. □

The Gröbner basis algorithm for $IP_{A,\succ_c}$ was introduced by Conti and Traverso (1991) and consists of the following steps:

1. Compute the reduced Gröbner basis $\mathcal{G}_{\succ_c}$ for I_A.
2. For a solution α to the program $IP_{A,\succ_c}(b)$, compute the *normal form* $\mathbf{x}^\beta$ of the monomial $\mathbf{x}^\alpha$. Then β is the unique optimal solution of $IP_{A,\succ_c}(b)$.

Two important issues that are bypassed in the above version of the Conti-Traverso algorithm are those of 1) finding a generating set for the toric ideal I_A and 2) finding an initial solution α for $IP_{A,\succ_c}(b)$. The original algorithm of Conti and Traverso (1991) meets both of these requirements by computing the reduced Gröbner basis of a toric ideal in a larger polynomial ring with respect to an elimination order that depends on $\succ_c$. We use the above short version to highlight the main features of the procedure. The issue of finding a generating set for I_A will be discussed in Section 5.

Example 2.4

$$A = \begin{bmatrix} 2 & 4 & 3 & 1 \\ 7 & 2 & 4 & 0 \end{bmatrix}, \quad c = \begin{bmatrix} 10 & 3 & 1 & 8 \end{bmatrix}, \quad b = \begin{bmatrix} 72 \\ 79 \end{bmatrix}$$

Suppose we refine the cost vector c by the graded reverse lexicographic order to obtain $\succ_c$. Then $\mathcal{G}_{\succ_c} = \{\underline{cd^5} - b^2, \underline{a^2d^9} - bc^3, \underline{a^2bd^4} - c^4, \underline{a^2b^3} - c^5d\}$ where $S = k[a,b,c,d]$. The vector $\alpha = (3,7,11,5)$ is a feasible solution for the above program, and the normal form of $a^3b^7c^{11}d^5$ with respect to $\mathcal{G}_{\succ_c}$ is $ab^6c^{15}d$. So the optimal solution to this program is $(1,6,15,1)$. □

Corollary 2.5 *Let $\mathcal{G}_{\succ_c}$ be the reduced Gröbner basis of I_A.*

1. *The initial monomial ideal of I_A with respect to $\succ_c$, denoted $in_{\succ_c}(I_A)$, is generated by the monomials $\mathbf{x}^\alpha$ where α is a non-optimal solution for some program in $IP_{A,\succ_c}$. The minimal generators of $in_{\succ_c}(I_A)$ are the initial terms of the binomials in $\mathcal{G}_{\succ_c}$.*

2. *A vector $\beta \in \mathbf{N}^n$ is the optimal solution to the program $IP_{A,\succ_c}(A\beta)$ if and only if $\mathbf{x}^\beta$ is a standard monomial of $in_{\succ_c}(I_A)$.*

We saw earlier that an element $\mathbf{x}^\alpha - \mathbf{x}^\beta \in \mathcal{G}_{\succ_c}$ can be thought of as the vector $\alpha - \beta \in ker_{\mathbf{Z}}(A)$. Alternately we may think of $\mathbf{x}^\alpha - \mathbf{x}^\beta \in \mathcal{G}_{\succ_c}$ as the line segment $[\alpha, \beta]$ in the $A\alpha$-fiber of A directed from α to β. This interpretation of the elements of $\mathcal{G}_{\succ_c}$ says that there exists a monotone directed path in P_b^I, from *any* solution α of the program $IP_{A,\succ_c}(b)$ to the optimum, whose edges are elements of $\mathcal{G}_{\succ_c}$. This observation can be used for enumerating lattice points in polyhedra (see Chapter 5, Sturmfels 1995). It was also used in Tayur et. al (1995) to solve a special family of stochastic integer programs that arose in a manufacturing setting. In fact, the entire Buchberger algorithm in the toric situation can be reduced to a geometric process on vectors in $ker_{\mathbf{Z}}(A)$ (Thomas 1995).

3 Universal Gröbner bases of toric ideals

A subset $\mathcal{U}$ of a polynomial ideal $I \subseteq S$ is called a *universal Gröbner basis* of I (Weispfenning 1987) if it contains a Gröbner basis for I with respect to every cost vector. If we let IP_A denote all integer programs with fixed coefficient matrix A, then we call $\mathcal{U}_A \subset ker_{\mathbf{Z}}(A)$ a *universal test set* for IP_A if $\mathcal{U}_A$ contains a test set for every family $IP_{A,c}$ as c varies. Therefore a universal Gröbner basis for I_A would be a universal test set for IP_A.

The *Graver basis* of A (Graver 1975) is a universal test set for IP_A, constructed as follows. For each $\sigma \in \{+,-\}^n$, consider the semigroup $S_\sigma = ker_{\mathbf{Z}}(A) \cap \mathbf{R}^n_\sigma$, where $\mathbf{R}^n_\sigma$ is the orthant with sign pattern σ. Then $pos(S_\sigma) = \{\sum_{i=1}^p \lambda_i x_i : x_i \in S_\sigma, \lambda_i \geq 0\}$ is a pointed polyhedral $(n-d)$-dimensional cone in $\mathbf{R}^n$. The *Hilbert basis* of a polyhedral cone K in $\mathbf{R}^n$ is a minimal subset of $K \cap \mathbf{Z}^n$ such that every integral vector in K can be written as a non-negative integral combination of the elements in the basis. Further, pointed cones have unique Hilbert bases (see Chapter 16, Schrijver 1986). Let H_σ denote the unique Hilbert basis of S_σ . This is finite since $A \in \mathbf{Z}^{d\times n}$. The Graver basis of A is the set $Gr_A := \cup_\sigma H_\sigma \backslash \{0\}$ which was shown by Graver (1975) to be a universal test set for IP_A. Blair and Jeroslow (1982) and Cook et al. (1986) gave equivalent test sets (see also Section 17.4, Schrijver 1986).

Let UGB_A denote the union of all reduced Gröbner bases $\mathcal{G}_c$ as c varies over all *generic* cost vectors in $\mathbf{R}^n$. A cost vector c is generic for A if the optimal solution with respect to c in every fiber P_b^I is a unique vertex. Clearly, UGB_A is a uniquely defined universal test set for IP_A and we call it the *universal Gröbner basis* of A. The sets Gr_A and UGB_A are related as follows.

Theorem 3.1 *(cf. Thomas 1995) The Graver basis Gr_A contains the universal Gröbner basis UGB_A.*

Corollary 3.2 *There exists only finitely many distinct reduced Gröbner bases for I_A as the cost function is varied. In particular, UGB_A is finite.*

Corollary 3.2 is a well known result for polynomial ideals (Mora and Robbiano 1988, Bayer and Morrison 1988) and what we provide above is an independent proof in the case of toric ideals. Both Gr_A and UGB_A can be computed using Gröbner basis methods as described by Sturmfels and Thomas (1997). We outline below the main ingredients of these algorithms. In order to compute Gr_A we consider the *Lawrence lifting* of A which is the enlarged matrix $\Lambda(A) = \begin{bmatrix} A & 0 \\ I & I \end{bmatrix} \in \mathbf{Z}^{(n+d)\times 2n}$ where 0 is a $d \times n$ matrix of all zeros and I is the $n \times n$ identity matrix. The matrices A and $\Lambda(A)$ have isomorphic kernels: $ker_{\mathbf{Z}}(\Lambda(A)) = \{(u, -u) : u \in ker_{\mathbf{Z}}(A)\}$. The toric ideal $I_{\Lambda(A)}$ is the homogeneous prime ideal

$$I_{\Lambda(A)} = \langle \mathbf{x}^\alpha \mathbf{y}^\beta - \mathbf{x}^\beta \mathbf{y}^\alpha \ : \ \alpha, \beta \in \mathbf{N}^n, A\alpha = A\beta \rangle$$

in the polynomial ring $k[x_1, \ldots, x_n, y_1, \ldots, y_n]$.

Theorem 3.3 *For the matrix $\Lambda(A)$, the following sets coincide:*
1. the Graver basis of $\Lambda(A)$,
2. the universal Gröbner basis of $\Lambda(A)$,
3. any reduced Gröbner basis of $I_{\Lambda(A)}$, and
4. any minimal generating set of $I_{\Lambda(A)}$ (up to scalar multiples).

Algorithm 3.4 How to compute the Graver basis of A.

1. Compute the reduced Gröbner basis $\mathcal{G}$ of $I_{\Lambda(A)}$ with respect to any term order.
2. The Graver basis Gr_A consists of all elements $\alpha - \beta$ such that $\mathbf{x}^\alpha \mathbf{y}^\beta - \mathbf{x}^\beta \mathbf{y}^\alpha$ appears in $\mathcal{G}$.

Proof: By Theorem 3.3 any reduced Gröbner basis of $I_{\Lambda(A)}$ is also the Graver basis of $\Lambda(A)$. The bijection between the kernels of A and $\Lambda(A)$ implies that a reduced Gröbner basis of $I_{\Lambda(A)}$ with the variables y_j set to one, is the Graver basis of A. $\square$

Since $UGB_A \subseteq Gr_A$, all we need now is a characterization of the elements of UGB_A so that they can be identified from among the elements of Gr_A. In the next section we will see a geometric characterization of the elements in UGB_A that can be used for this purpose. A second test can be found in Sturmfels and Thomas (1997).

Example 2.4 continued. For the matrix A in Example 2.4, $Gr_A = UGB_A = \{b^2 - cd^5,\ a^2d^9 - bc^3,\ a^2d^{14} - b^3c^2,\ a^2d^{19} - b^5c,\ a^2d^{24} - b^7,\ a^2bd^4 - c^4,\ a^2b^3 - c^5d,\ a^4d^{13} - c^7,\ a^4b^4d^3 - c^9,\ a^6b^7d^2 - c^{14},\ a^8b^{10}d - c^{19},\ a^{10}b^{13} - c^{24}\}$. $\square$

A natural question now is to ask whether it is possible to bound the *degree* of elements in UGB_A and/or Gr_A in terms of the matrix A. By the degree of a binomial $\mathbf{x}^{v^+} - \mathbf{x}^{v^-}$ we mean $\sum_{i=1}^{n} |v_i|$, the 1-norm of the vector v. Let $D(A) = max\{|det(a_{i_1}, \ldots, a_{i_d})| : 1 \leq i_1 < i_2 < \cdots < i_d \leq n\}$. (Recall that A is assumed to have rank d.) Then the following theorem of Sturmfels (1991) gives a bound on the degree of elements in Gr_A.

Theorem 3.5 *The degree of a binomial in the Graver basis of A is at most $(d+1)(n-d)D(A)$.*

Improving degree bounds of elements in Gr_A is an important open problem in this area. It has been conjectured that the true degree bound for elements in Gr_A should not involve n. It was shown by Hoşten (1997) that if the cone $pos(S_\sigma)$ is simplicial then the maximum degree of a Graver basis element from this cone is at most $(d+1)D(A)$. A bound that is independent of n but depends exponentially on d is also given by Hoşten (1997). In the special case of $2 \times n$ matrices, this new bound can be improved further and the proof also yields an algorithm to compute Gr_A.

4 Variations of cost functions in integer programming

In this section we collect certain polyhedral results for integer programming that arise from the Gröbner basis approach to the theory. These results first appeared in (Sturmfels and Thomas 1997) and we refer the reader to that paper for proofs and details. For the basics of polyhedral theory see (Ziegler 1995).

A polytope P is a bounded polyhedron that can either be written as $P = \{x \in \mathbf{R}^n : Mx \leq m\}$ for some $M \in \mathbf{R}^{p \times n}$ and $m \in \mathbf{R}^p$, or as $P = \{\lambda_1 v_1 + \cdots + \lambda_t v_t : \sum \lambda_i = 1, \lambda_i \geq 0, \forall i\}$ where the points $v_1, \ldots, v_t$ are the vertices of P. The dimension of P, denoted $dim(P)$, is the dimension of its affine span. For a polytope $P \subset \mathbf{R}^n$ and a vector $c \in \mathbf{R}^n$, we define $face_c(P)$ to be the set of all points in P at which the linear functional $c \cdot x$ is minimized. Hence $face_c(P) = \{x \in P : c \cdot x \leq c \cdot y, \ \forall y \in P\}$. If F is any face of P, then $\mathcal{N}(F; P)$ denotes the cone of (inner) normals, called the *inner normal cone* of P at F. In symbols, $\mathcal{N}(F; P) = \{c \in \mathbf{R}^n : c \cdot x \leq c \cdot y \text{ for all } x \in F, y \in P\}$. A *polyhedral fan* in $\mathbf{R}^n$ is a collection of polyhedral cones (also called *cells* of the fan) with the property that if C is in the fan then so is every face of C, and if C_1 and C_2 are in the fan then their intersection is a common face of each. A polyhedral fan in $\mathbf{R}^n$ is said to be *complete* if the union of all elements in the fan is $\mathbf{R}^n$. It may be noted that the collection of cones $\mathcal{N}(F; P)$ as F ranges over all faces of P, denoted $\mathcal{N}(P)$, is a complete polyhedral fan in $\mathbf{R}^n$ and we call it the (inner) *normal fan* of P. We say that two polytopes are *normally equivalent* if they have the same normal fan.

Given two polytopes P and Q in $\mathbf{R}^n$, their *Minkowski sum* is the polytope $P + Q = \{p + q : p \in P, \, q \in Q\} \subset \mathbf{R}^n$, and P and Q are called *Minkowski summands* of $P + Q$. As in the usual extension of addition to integration, the operation of taking Minkowski sums of finitely many polytopes extends naturally to the operation of taking *Minkowski integrals* of infinitely many polytopes (see Billera and Sturmfels 1992). The common *refinement* of two fans $\mathcal{F}$ and $\mathcal{G}$ in $\mathbf{R}^n$, denoted $\mathcal{F} \cap \mathcal{G}$, is the fan of all intersections of cones from $\mathcal{F}$ and $\mathcal{G}$. We say that $\mathcal{F} \cap \mathcal{G}$ is a refinement of $\mathcal{F}$ (respectively $\mathcal{G}$). The following are two useful facts in this context: (i) for polytopes P and Q in $\mathbf{R}^n$, the fan $\mathcal{N}(P+Q) = \mathcal{N}(P) \cap \mathcal{N}(Q)$ and (ii) the fan $\mathcal{N}(P)$ is a refinement of $\mathcal{N}(Q)$ if and only if λQ is a Minkowski summand of P for some positive real number λ.

We are interested in studying the effect of varying the cost vector c on integer programs in IP_A. We start with the following definition of when two cost vectors (not just generic ones) are equivalent with respect to A.

Definition 4.1 *Two cost vectors c and c' in $\mathbf{R}^n$ are equivalent (with respect to A) if the integer programs $IP_{A,c}(b)$ and $IP_{A,c'}(b)$ have the same set of optimal solutions for all b in $\mathbf{N}(A)$.*

Our goal in this section is to provide a structure theorem (Theorem 4.8) for these equivalence classes. Using the properties of Gröbner bases of toric ideals, one can characterize generic equivalence classes as follows. The proof follows from properties of the reduced Gröbner bases of I_A.

Proposition 4.2 *Given two generic cost vectors c and c' in $\mathbf{R}^n$, the following are equivalent:*
1. For every $b \in \mathbf{N}(A)$, the programs $IP_{A,c}(b)$ and $IP_{A,c'}(b)$ have the same optimal solution.
2. The cost vectors c and c' support the same optimal vertex in each fiber P_b^I.
3. The reduced Gröbner bases $\mathcal{G}_c$ and $\mathcal{G}_{c'}$ of I_A are equal.

Let $u = u^+ - u^- \in ker_{\mathbf{Z}}(A)$. Both u^+ and u^- lie in the Au^+-fiber of A, and we may think of u as the line segment $[u^+, u^-]$ in this fiber. We shall refer to the polytope $P^I_{Au^+}$ as *the fiber of* u. By a *Gröbner fiber* of A we mean the fiber of an element $u \in UGB_A$. Let $St(A)$ denote the Minkowski sum of all Gröbner fibers. This is a well-defined polytope in $\mathbf{R}^n$ which we call the *state polytope* of A. Since the elements in $\mathcal{G}_c$ span $ker_{\mathbf{Z}}(A)$ over the integers, it follows that $St(A)$ is an $(n-d)$-dimensional polytope. The complete polyhedral fan $\mathcal{N}(St(A))$ is called the *Gröbner fan* of A. The state polytope and Gröbner fan of a graded polynomial ideal were introduced by (Bayer and Morrison, 1988) and (Mora and Robbiano, 1988) respectively. What we provide here is a self contained construction of these entities in the toric situation.

We now state a lemma that plays a crucial role in proving Theorem 4.8.

Lemma 4.3 *Every fiber P^I_b of A is a Minkowski summand of $St(A)$.*

Proof: By the facts stated at the beginning of this section, it suffices to show that $\mathcal{N}(St(A))$ is the common refinement of $\mathcal{N}(P^I_b)$ for all $b \in \mathbf{N}(A)$. Let c be a generic cost vector and $w \neq c$ belong to the interior of the cone $\mathcal{N}(face_c(St(A)); St(A))$. Then w lies in $\mathcal{N}(\beta_i; P^I_{A\beta_i})$ for each element $\alpha_i - \beta_i$ in the reduced Gröbner basis $\mathcal{G}_c$. This implies that $w \cdot \alpha_i > w \cdot \beta_i$ for all i, and therefore, $\mathcal{G}_w = \mathcal{G}_c$.

Now consider an arbitrary $b \in N(A)$. Let u be the unique optimum of $IP_{A,c}(b)$. The equality of test sets $\mathcal{G}_w = \mathcal{G}_c$ implies that u is also the unique optimum of $IP_{A,w}(b)$. Hence w lies in the interior of $\mathcal{N}(u; P^I_b)$. Therefore, $\mathcal{N}(face_c(St(A)); St(A)) \subseteq \mathcal{N}(u; P^I_b)$, as desired. □

Proposition 4.4 *Let db denote any probability measure with support $\mathbf{N}(A)$ such that $\int_b b\,db$ is finite. Then the Minkowski integral $\int_b P^I_b db$ is a polytope normally equivalent to $St(A)$.*

Proof: The hypothesis $\int_b b\,db < \infty$ guarantees that $\int_b P^I_b db$ is bounded. By Lemma 4.3, $\int_b P^I_b db$ is a summand of $St(A)$ and is hence a polytope. However, each Gröbner fiber is a summand of $\int_b P^I_b db$ and hence $\int_b P^I_b db$ is a polytope of dimension $n-d$ in $\mathbf{R}^n$ that has the same normal fan as $St(A)$. □

From now on we shall use the term *state polytope* for any polytope normally equivalent to $\int_b P^I_b db$. We define the *Gröbner cone* associated with $\mathcal{G}_c$ to be the closed convex polyhedral cone

$$\mathcal{K}_c \quad := \quad \{\, x \in \mathbf{R}^n \,:\, g_i \cdot x \geq 0,\ g_i \in \mathcal{G}_c \,\}$$

Observation 4.5 *The Gröbner cone $\mathcal{K}_c$ is n-dimensional. Its lineality space (largest subspace contained in the cone) $\mathcal{K}_c \cap -\mathcal{K}_c$ equals $rowspan(A) \simeq \mathbf{R}^d$.*

Proposition 4.6 *The Gröbner fan of A is the collection of all Gröbner cones $\mathcal{K}_c$ together with their faces, as c varies over all generic cost vectors.*

Proof: The argument in the proof of Lemma 4.3 shows that, for c generic, the Gröbner cone $\mathcal{K}_c$ equals $\mathcal{N}(face_c(St(A)); St(A))$. □

Corollary 4.7 *The equivalence classes of cost vectors with respect to IP_A are precisely the cells of the Gröbner fan.*

Proof: Two cost vectors c and c' are equivalent if and only if they support the same optimal face in each fiber of A. Using Propositions 4.4 and 4.6, it follows that c and c' are equivalent if and only if they lie in the relative interior of the same cell in $\mathcal{N}(St(A))$. □

A cost vector w lies in the interior of a Gröbner cone $\mathcal{K}_c$ if and only if w is generic and equivalent to c. Hence the interiors of the maximal cells in the Gröbner fan are precisely the equivalence classes of generic cost vectors. The following theorem summarizes the above discussion.

Theorem 4.8 *1. There are only finitely many equivalence classes of cost vectors with respect to A.*
2. Each equivalence class is the relative interior of a convex polyhedral cone in $\mathbf{R}^n$.
3. The collection of these cones defines a complete polyhedral fan in $\mathbf{R}^n$ called the Gröbner fan of A.
4. Let db denote any probability measure with support $\mathbf{N}(A)$ such that $\int_b b\,db < \infty$. Then the Minkowski integral $St(A) = \int_b P_b^I db$ is an $(n-d)$-dimensional convex polytope, called the state polytope of A. The normal fan of $St(A)$ equals the Gröbner fan of A.

Example 2.4 continued. The toric ideal I_A of the matrix A in Example 2.4 has fifteen distinct reduced Gröbner bases. We list the distinct initial ideals below along with a representative cost vector from each Gröbner cone. The state polytope is hence a two dimensional 15-gon in $\mathbf{R}^4$. The initial ideals are listed in the order you would encounter them as you walk along the boundary of $St(A)$.

No.	Initial ideal	Cost vector
1.	$\{a^2bd^4, cd^5, a^2d^9, a^2b^3\}$	$[66, 32, 15, 62]$
2.	$\{a^8b^{10}d, a^6b^7d^2, a^4b^4d^3, a^2bd^4, cd^5, a^2d^9, c^5d, a^{10}b^{13}\}$	$[29/2, 0, 6, 0]$
3.	$\{a^2bd^4, a^2d^9, b^2\}$	$[9/4, 1/2, 0, 0]$
4.	$\{a^2bd^4, bc^3, a^4d^{13}, b^2\}$	$[1/4, 3/2, 0, 0]$
5.	$\{a^2bd^4, bc^3, b^2, c^7\}$	$[0, 11/7, 1/7, 0]$
6.	$\{bc^3, b^2, c^4\}$	$[0, 9/7, 4/7, 0]$
7.	$\{cd^5, bc^3, b^3c^2, b^5c, b^7, c^4\}$	$[0, 1/7, 9/7, 0]$
8.	$\{cd^5, bc^3, b^3c^2, b^5c, a^2d^24, c^4\}$	$[1/2, 0, 2, 0]$
9.	$\{cd^5, bc^3, b^3c^2, a^2d^{19}, c^4\}$	$[3/2, 0, 2, 0]$
10.	$\{cd^5, bc^3, a^2d^{14}, c^4\}$	$[5/2, 0, 2, 0]$
11.	$\{cd^5, a^2d^9, c^4\}$	$[11/2, 0, 3, 0]$
12.	$\{a^2bd^4, cd^5, a^2d^9, c^5d, c^9\}$	$[13/2, 0, 3, 0]$
13.	$\{a^4b^4d^3, a^2bd^4, cd^5, a^2d^9, c^5d, c^{14}\}$	$[23/2, 0, 5, 0]$
14.	$\{a^6b^7d^2, a^4b^4d^3, a^2bd^4, cd^5, a^2d^9, c^5d, c^{19}\}$	$[33/2, 0, 7, 0]$
15.	$\{a^8b^{10}d, a^6b^7d^2, a^4b^4d^3, a^2bd^4, cd^5, a^2d^9, c^5d, c^{24}\}$	$[43/2, 0, 9, 0]$

We close this section with a geometric characterization of the elements in the universal Gröbner basis UGB_A.

Theorem 4.9 *A vector $\alpha - \beta \in ker_{\mathbf{Z}}(A)$ lies in the universal Gröbner basis UGB_A if and only if $\alpha - \beta$ is primitive and the line segment $[\alpha, \beta]$ is an edge of the $A\alpha$-fiber of A.*

For a proof of this theorem see Sturmfels and Thomas (1997). In view of Proposition 4.4, this says that the edge directions of the fibers of A are precisely the edge directions of the state polytope. If $[\alpha, \beta]$ is the primitive representative of an edge direction, then $[\alpha, \beta]$ is an edge of the $A\alpha$-fiber of A. Therefore, Theorem 4.9 is equivalent to the following assertion: *the universal Gröbner basis UGB_A consists of the edge directions of the state polytope.*

Theorem 4.9 implies several interesting corollaries.

Corollary 4.10 *For an element $\alpha - \beta$ in UGB_A, there exists two cost vectors c and c' in $\mathbf{R}^n$ such that $\alpha - \beta \in \mathcal{G}_c$ and $\beta - \alpha \in \mathcal{G}_{c'}$.*

Proof: Every element in UGB_A appears as a facet normal of some cell in the Gröbner fan. Take as $\mathcal{G}_c$ and $\mathcal{G}_{c'}$ the Gröbner bases associated with the two Gröbner cones that share this facet. □

Corollary 4.11 *For every generic cost vector $c \in \mathbf{R}^n$, the reduced Gröbner basis of $IP_{A,c}$ consists only of edges of certain fibers P_b^I.*

Theorem 4.9 implies that we can trace a monotone *edge path* from every vertex of P_b^I that is non-optimal with respect to the cost vector c to the optimal vertex, using only elements in UGB_A. Thus reduction with respect to the universal Gröbner basis can be viewed as an integer analogue to the simplex method for linear programming.

Theorem 4.9 gives rise to the following algorithm for computing the universal Gröbner basis.

Algorithm 4.12 How to compute the universal Gröbner basis UGB_A

1. Compute the Graver basis Gr_A (see Section 3).
2. For each element $\mathbf{x}^\alpha - \mathbf{x}^\beta$ of Gr_A decide whether $[\alpha, \beta]$ is an edge of its fiber.

5 Implementing toric Gröbner basis algorithms

In principle, one can compute Gröbner bases of toric ideals using available software like Macaulay2 or CoCoA. However, more careful implementations of Buchberger's algorithm along with new ideas for improving this algorithm are possible due to the nice geometric properties of toric ideals. This in turn considerably speeds up computations. Pioneering attempts in this direction are implementations like BaStaT (Pottier 1996) and GRIN (Hoşten and Sturmfels 1995). More implementations of the same kind are emerging (see Bigatti et al.). First of all, these implementations benefit from the fact that binomials require very simple data structures since they can be identified with pairs of lattice points. Moreover, S-pair computations and reduction operations in the Buchberger algorithm can be all formulated in terms of vector additions and translations (see Thomas 1995, and the discussion at the end of Section 2). Another simple observation which leads to considerable speed-up is that I_A is a prime ideal which does not contain monomials; so whenever the terms of a binomial which is created during the computation has a common monomial factor, we can simply divide by that factor.

Although the above ideas improve algorithms for computing Gröbner bases of toric ideals, it seems that the real bottleneck is computing an initial generating set for I_A. In fact, in most of the applications such a generating set is not part of the data. One is usually presented with the matrix A and a term order only. A first idea to circumvent this problem is an elimination technique as described by Conti and Traverso (1991). For this, let $A = [a_1, \ldots, a_n]$, and for simplicity assume that the columns a_i are nonnegative integral vectors.

Proposition 5.1 *Let $J = \langle x_i - \mathbf{y}^{a_i} : i = 1, \ldots, n\rangle \subset k[x_1, \ldots, n, y_1, \ldots, y_d]$. Then $I_A = J \cap k[x_1, \ldots, x_n]$ and the reduced Gröbner basis $\mathcal{G}_\succ(I_A) = k[x_1, \ldots, x_n] \cap \mathcal{G}_{\succ'}(J)$ where $\succ'$ is a term order with $y_1, \ldots, y_d \succ' x_1, \ldots, x_n$ and monomials in $k[x_1, \ldots, x_n]$ are compared by $\succ$.*

The method that this proposition suggests is in general time-consuming for computing a Gröbner basis when a generating set for I_A is absent, since it involves d extra variables and $\mathcal{G}_{\succ'}(J)$ is typically quite large compared to $\mathcal{G}_{\succ}(I_A)$. Now let $\mathcal{B} = \{u_1, \ldots, u_{n-d}\} \subset \mathbf{Z}^n$ be a lattice basis for $\ker_{\mathbf{Z}}(A)$, and let $J = \langle f_{u_1}, \ldots, f_{u_{n-d}} \rangle$ where $f_{u_i} = \mathbf{x}^{u_{i+}} - \mathbf{x}^{u_{i-}}$. The following proposition appears in (Conti and Traverso 1991), (Fischer and Shapiro 1994) and (Hoşten and Sturmfels 1995).

Proposition 5.2 *$I_A = J : (\prod_{i=1}^{n} x_i)^\infty$ and $I_A = (J + \langle t(\prod_{i=1}^{n} x_i) - 1 \rangle) \cap k[x_1, \ldots, x_n]$.*

Using the second statement of the proposition and an elimination term order where t is the most expensive variable we can compute a Gröbner basis of I_A. But although there is just one extra variable involved in this method we still get a fairly large Gröbner basis which contains the wanted Gröbner basis of I_A as a subset. The nice thing about the above ideal quotient is that it can be computed one variable at a time as follows:

$$J : (\prod_{i=1}^{n} x_i)^\infty = ((\cdots (J : x_1^\infty) : x_2^\infty) \cdots) : x_n^\infty).$$

Suppose that A contains a positive vector $c = (c_1, \ldots, c_n)$ in its row space, i.e. both I_A and J are homogeneous with respect to the grading $\deg(x_i) = c_i$. Now each one of the above ideal quotients can be computed as follows: first we compute a Gröbner basis where we use a c-graded reverse lexicographic order which makes the corresponding variable x_i the cheapest variable; then we divide each element of this Gröbner basis by the highest power of x_i. This allows us to compute a generating set for I_A by doing n Gröbner basis computations. These individual computations are observed to be short in practice. See the discussion in (Hoşten and Sturmfels 1995). A similar algorithm is suggested by DiBiase and Urbanke (1995) which involves at most $\lfloor n/2 \rfloor$ Gröbner basis computations. In fact, one needs to make at most $\lfloor n/2 \rfloor$ Gröbner basis computations instead of n in the above algorithm as well: Proposition 5.2 implies that I_A is an associated prime of J and any other associated prime of J contains pure variables. Now suppose M is the $(n-d) \times n$ matrix whose rows are the elements of $\mathcal{B} = \{u_1, \ldots, u_{n-d}\}$. We will call a vector *mixed* if it contains both positive and negative entries. Otherwise it will be called *unmixed*. Since we assume that $\ker(A)$ does not contain nonnegative vectors, all the rows of M are mixed. The next result is new (Hoşten and Shapiro 1997).

Theorem 5.3 *There exist variables $x_{i_1}, \ldots, x_{i_k}$ such that $I_A = J : (\prod_{j=1}^{k} x_{i_j})^\infty$ where $k \leq \lfloor n/2 \rfloor$.*

Proof: We construct a subset $T = \{x_{i_1}, \ldots, x_{i_k}\}$ of the variables such that for every associated prime P of J we have $P \cap T \neq \emptyset$. We start with $T = \emptyset$.

We can assume that the first k_1 entries of the first row of M are nonzero and the rest of u_1's entries are zero. Furthermore we can also assume that the first s_1 entries of u_1 are positive where $s_1 \leq \lfloor k_1/2 \rfloor$. We set $T = T \cup \{x_1, \ldots, x_{s_1}\}$. Then we delete the first row and the first k_1 columns of M. If the resulting matrix have all unmixed rows we stop, if not, we can assume that the top row is mixed. We can further assume that the first k_2 entries of this row are nonzero and the first $s_2 \leq \lfloor k_2/2 \rfloor$ of them are positive; and we set $T = T \cup \{x_{k_1+1}, \ldots, x_{k_1+s_2}\}$ and we repeat. When we are done, T will have $s_1 + s_2 + \cdots + s_p \leq \lfloor k_1/2 \rfloor + \lfloor k_2/2 \rfloor + \cdots + \lfloor k_p/2 \rfloor \leq \lfloor n/2 \rfloor$ variables. So M will be in the form

$$\left(\begin{array}{c|c} B & 0 \\ \hline C & D \end{array}\right)$$

where B is an $p \times (\sum_{j=1}^{p} k_j)$ matrix of the form

$$\begin{pmatrix} + & \cdots & + & - & \cdots & - & 0 & \cdots & 0 & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ * & \cdots & * & * & \cdots & * & + & \cdots & + & - & \cdots & - & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ & \vdots & & & \vdots & & & \vdots & & & \vdots & & \ddots & & \ddots & & & \ddots & \\ * & \cdots & * & * & \cdots & * & * & \cdots & * & * & \cdots & * & \cdots & + & \cdots & + & - & \cdots & - \end{pmatrix}$$

and D is a matrix whose rows are unmixed. Now suppose x is a variable contained in an associated prime of J. If $x \in T$ then we are done. So suppose that $x = x_i \notin T$ and it corresponds to a column of B. This means that x_i corresponds to an entry in the final block of negative entries in a row u_j of B. Then P must contain yet another variable $x_{i'}$ which corresponds to one of the positive entries of row j. If $x_{i'}$ is in T, we are done. If not, since $i' < i$, we repeat this finitely many times and eventually find a variable in $T \cap P$. If x corresponds to a column of D, then we know that P must contain another variable corresponding to one of the columns of B because each row of D is unmixed. This brings us back to the above case. □

Often times one is interested in solving $IP_{A,c}(b)$ for just one fixed b. In this situation, the full Gröbner basis might be much larger than a sufficient test set for this problem. Given a fixed b, Thomas and Weismantel (1995) show how one can truncate the Buchberger algorithm using b as a *degree bound* to obtain a *truncated* Gröbner basis that will solve the particular $IP_{A,c}(b)$ as well as other integer programs whose right hand side vector are "lower" than b in a specific sense. When one is interested in $0-1$-solutions to integer programs even more can be done: Urbaniak et al. (1997) give a variant of the Buchberger algorithm in this setting.

Gröbner bases of toric ideals can be used to gain more information about the structure of integer programs. The results in Section 4 are good examples of such use. We believe the future will bring more results in this direction. Conversely, we expect that the integer programming literature and in general convex geometry will help our understanding of toric ideals, toric varieties and their Gröbner bases. For an example for this see (Hoşten 1997) and (Hoşten and Thomas 1997).

References

[1] Bayer, D., Morrison, I. (1988): Gröbner bases and geometric invariant theory I. J. Symb. Comp. 6:209–217.

[2] Bigatti, A.M., LaScala, R., Robbiano, L.: Computing toric ideals, in preparation.

[3] Billera, L.J., Sturmfels, S. (1992): Fiber polytopes. Annals of Math. 135:527–549.

[4] Blair, C.E., Jeroslow, R.G. (1982): The value function of an integer program. Mathematical Programming 23:237–273.

[5] Buchberger, B. (1965): On Finding a Vector Space Basis of the Residue Class Ring Modulo a Zero Dimensional Polynomial Ideal (German). PhD Thesis, Univ. of Innsbruck, Austria.

[6] Conti, P., Traverso, C. (1991): Gröbner bases and integer programming. In: Proceedings AAECC-9, New Orleans. Springer-Verlag, LNCS **539**, pp. 130–139.

[7] Cook, W., Gerards, A.M.H., Schrijver, A., Tardos, É. (1986): Sensitivity theorems in integer linear programming. Mathematical Programming, 34:251–264.

[8] DiBiase, F., Urbanke, R. (1995): An algorithm to calculate the kernel of certain polynomial ring homomorphisms. Experimental Mathematics, 4:227–234.

[9] Fischer, K.G., Shapiro, J. (1994): Generating prime ideals in the Minkowski ring of polytopes. In: Fischer, K.G., Loustaunau, P., Shapiro, J., Green, E.L., Farkas, D. (eds.): Computational Algebra, Marcel Dekker Inc., Lecture Notes in Pure and Applied Mathematics **151**, New York, pp. 111–130.

[10] Graver, J.E. (1975): On the foundations of linear and integer programming I. Mathematical Programming, 8:207–226.

[11] Hoşten, S. (1997): Degrees of Gröbner bases of integer programs. Ph.D thesis, Cornell University.

[12] Hoşten, S., Shapiro, J. (1997): Primary decompositions of lattice basis ideals, in preparation.

[13] Hoşten, S., Sturmfels, B. (1995): GRIN: An implementation of Gröbner bases for integer programming. In: Balas, E., Clausen, J. (eds.): Integer Programming and Combinatorial Optimization. Springer Verlag, LNCS **920**, pp. 267–276.

[14] Hoşten, S., Thomas, R.R. (1997): Arithmetic degree of initial ideals of toric ideals, in preparation.

[15] Mora, T., Robbiano, L. (1988): Gröbner fan of an ideal. J. Symb. Comp., 6:183–208.

[16] Tayur, S.R., Thomas, R.R., Natraj, N.R., (1995): An algebraic geometry algorithm for scheduling in presence of setups and correlated demands. Mathematical Programming, 69, Ser. A:396–401.

[17] Pottier, L. (1996): BaStaT (downloadable from `http://www.inria.fr/safir/WHOSWHO/Loic/Bastat/bastatdemo.html`)

[18] Scarf, H.E. (1986): Neighborhood systems for production sets with indivisibilities. Econometrica, 54:507-532.

[19] Schrijver, A. (1986): Theory of Linear and Integer Programming. Series in Discrete Mathematics. Wiley-Interscience, New York.

[20] Sturmfels, B. (1991): Gröbner bases of toric varieties. Tôhoku Math. Journal, 43:249-261.

[21] Sturmfels, B. (1995): Gröbner Bases and Convex Polytopes. AMS University Lecture Series **8**, Providence.

[22] Sturmfels, B., Thomas, R.R. (1997): Variation of cost functions in integer programming. Mathematical Programming, 77:357–387.

[23] Thomas, R.R. (1995): A geometric Buchberger algorithm for integer programming. Mathematics of Operations Research, 20:864–884.

[24] Thomas, R.R., Weismantel, R. (1995): Truncated Gröbner bases for integer programming. Preprint SC 94-29, ZIB Berlin.

[25] Urbaniak, R., Weismantel, R., Ziegler, G. (1997): A variant of Buchberger's algorithm for integer programming. SIAM J. Discrete Math., 10:96–108.

[26] Weispfenning, V. (1987): Constructing universal Gröbner bases. In: Proceedings AAEEC 5, Menorca 1987, Springer Verlag, LNCS **356**, pp. 408-417.

[27] Ziegler, G. (1995): Lectures on Polytopes. Graduate Texts in Math., Springer-Verlag, New York.