

Problems in Computational Algebra and Integer Programming

Tristram Bogart

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2007

Program Authorized to Offer Degree: Mathematics

UMI Number: 3265300

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3265300

Copyright 2007 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Tristram Bogart

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

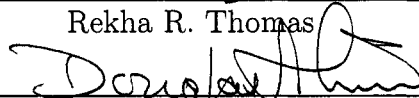


Rekha R. Thomas

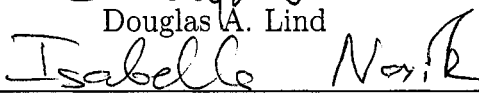
Reading Committee:



Rekha R. Thomas



Douglas A. Lind



Isabella Novik

Date:

7 June 2007

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Proquest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature Justin D. [unclear]

Date 7 June 2007

University of Washington

Abstract

Problems in Computational Algebra and Integer Programming

Tristram Bogart

Chair of the Supervisory Committee:
Associate Professor Rekha R. Thomas
Department of Mathematics

This thesis is a compendium of several projects that span the gap between commutative algebra and geometric combinatorics: one in tropical geometry, one in computational algebra, and two in discrete geometry and integer programming. Key tools used in this thesis include the theory of polyhedra, Gröbner bases, and computational complexity theory.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	v
Chapter 1: Introduction	1
1.1 Overview	1
1.2 Background	3
Chapter 2: Tropical Varieties	15
2.1 Tropical Algebra and Geometry	15
2.2 Computing Tropical Varieties	20
Chapter 3: Circuit Ideals	54
3.1 Introduction	54
3.2 Properties of $I_{\mathcal{A}}$ versus $J_{\mathcal{B}}$	58
3.3 Monomial Initial Ideals of the Circuit Ideal	62
3.4 Associated Primes of the Circuit Ideal	74
3.5 Fans of Toric and Circuit Ideals	83
Chapter 4: State polytopes of toric ideals	91
4.1 Introduction	91
4.2 Definitions and Properties	92
4.3 Toric state polygons	97
4.4 Combinatorial types	101
Chapter 5: Small Chvátal Rank	108
5.1 Introduction	109
5.2 Main Definitions	112

5.3	Contrasting small Chvátal rank with Chvátal rank	116
5.4	Supernormality and small Chvátal rank	120
5.5	Lower Bounds on SCR	124
	Bibliography	132

LIST OF FIGURES

Figure Number	Page
1.1 A triangle with outer normal cones (left) and its outer normal fan (right)	6
1.2 The staircase diagrams of the monomial ideals $M_1 = \langle x^2, y^3, z^4, xy^2, xz^2, xyz \rangle$ and $M_2 = \langle x^2yz^2, xy^3z^2 \rangle$.	9
2.1 The tropical hyperplane in \mathbb{R}^2 of $F(X, Y) = X \oplus (2 \otimes Y) \oplus 3$.	17
2.2 The Newton polygon of Example 2.1 with the rays of its inner normal fan.	17
2.3 Two tropical hyperplanes $\mathcal{T}(x + y + t)$ and $\mathcal{T}(x + y + t^3)$. Their intersection, shown as a thick ray, is not a tropical variety.	20
2.4 A projective drawing of the situation in Algorithm 2.29, with $\mathcal{T}(I)$ on the left and $\mathcal{T}(\text{in}_{\prec_{\mathbf{u}}}(I))$ on the right.	39
2.5 The tropical variety of the ideal generated by the 3×3 minors of the generic 4×4 Hankel matrix.	46
2.6 $p = 1, 3$	50
2.7 $p = 5, 7$	50
3.1 The distribution of values of \mathbf{H}_{J_B} for the circuit ideal of the matrix A in Example 3.12.	63
3.2 The point configuration of Example 3.35.	76
3.3 The Gröbner fan from Example 3.57 and the two radical equivalence classes. The fan is drawn in the standard simplex with $(1, 0, 0)$ at the right bottom, $(0, 1, 0)$ at the left bottom and $(0, 0, 1)$ at the top.	88
3.4 The Gröbner fans in the proof of Proposition 3.61 intersected with the simplex with coordinates $(0, 1, 0, 0)$ (right), $(0, 0, 1, 0)$ (left) and $(0, 0, 0, 1)$ (top). The Gröbner fan of $J_{\mathcal{C}_A}$ is to the left, the Gröbner fan of I_A in the middle, and the hypergeometric fan at the right.	89
4.1 The pointed secondary fan (left) and pointed Gröbner fan (right) in Example 4.13.	100

4.2	Schlegel diagrams of the seven combinatorially distinct 3-polytopes with six vertices. The first is a toric state polytope; the others are all decomposable. The last one has two non-triangular facets but its triangles are numbered to form a ridge path.	107
5.1	The polygon R_j^{j-1} (with $j = 6$) used to prove Theorem 5.19. Each integer point (a, b) in the polygon is labeled with the smallest number k such that $(1, a, b)$ appears in $\mathcal{A}_6^{(k)}$	125

LIST OF TABLES

Table Number	Page
3.1 Comparison of initial ideals of $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ from Example 3.16.	66

ACKNOWLEDGMENTS

My advisor, Rekha Thomas, has been a wonderful mentor throughout my years in Seattle by suggesting lines of research, actively collaborating in much of my work, helping me participate in valuable conferences and workshops, and always encouraging me when progress is slow. Anders Jensen, David Speyer, and Bernd Sturmfels have all been brilliant, friendly, and patient collaborators while introducing me to the world of tropical geometry. Bernd's work lies behind many of my problems and techniques. Thanks also to Anders for his software package `Gfan` on which I have relied extensively. My work has benefitted from discussion with many other people, including Eric Babson, Anton Dochtermann, Jesus De Loera, Diane Maclagan, Ezra Miller, Alex Papazoglou, Edwin O'Shea, and Vic Reiner. Brooke Miller, Mary Sheetz, and the rest of the Washington math department staff have always been helpful. Thanks also to all of my friends in Seattle and to my parents Joanne and Rick, my sister Zoë, and the rest of my family for their support.

Chapter 1

INTRODUCTION

1.1 Overview

The intersection of commutative algebra and algebraic geometry with combinatorics and discrete geometry is a rich and expanding field. Integral polytopes, purely discrete objects, provide enlightening examples to algebraic geometers via the theory of toric varieties; conversely, structural results in the theory of polytopes such as the Upper Bound Theorem have been proved using commutative algebra. This thesis is a compendium of several projects that cross this gap: one in tropical geometry, one in computational algebra, and two in discrete geometry and integer programming.

Chapter 2 concerns *tropical geometry* and derives mainly from a joint paper [8] with Anders Jensen, David Speyer, Bernd Sturmfels, and Rekha Thomas. Tropical geometry [5, 6, 41] is a recently developed tool to study algebraic varieties via discrete and computational methods. To every affine algebraic variety V is associated a polyhedral fan, its *tropical variety*. This object satisfies standard algebro-geometric relations, such as Bezout's theorem, that apply to V itself. In our work, we define *tropical bases*, analogous to generating sets of ordinary ideals. We show that tropical bases always exist, though they may be surprisingly large. We present an algorithm to compute tropical varieties and illustrate its implementation in the software package `Gfan` [34]. To prove correctness of the algorithm, we show that any irreducible tropical variety is connected in codimension one. Finally, we demonstrate the superiority of our algorithm to a naive one via a family of examples.

Chapter 3, taken from a joint paper [9] with Anders Jensen and Rekha Thomas,

introduces and contrasts two types of polynomial ideals defined from a finite vector configuration $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{Z}^d$. A *circuit* of \mathcal{A} is a vector \mathbf{c} representing a minimal integer dependency among the elements of \mathcal{A} . The *circuit ideal* is the ideal generated by the binomials $\mathbf{x}^{\mathbf{c}^+} - \mathbf{x}^{\mathbf{c}^-} \in \mathbb{C}[x_1, \dots, x_n]$ as $\mathbf{c} = \mathbf{c}^+ - \mathbf{c}^- \in \mathbb{Z}^n$ varies over the circuits of \mathcal{A} . The radical of the circuit ideal is the *toric ideal* of \mathcal{A} , defined similarly but where *all* integer dependencies are used, rather than just the minimal ones. Toric ideals are the defining ideals of *toric varieties* [25] and have applications to combinatorics, optimization, algebra, and algebraic geometry [47]. However, they are difficult to compute. Since circuits can be computed using linear algebra and the two ideals often coincide, it is worthwhile to understand when equality occurs. We give a complete characterization of when the circuit and toric ideals are equal in terms of the existence of certain polytopes and their integer points. When the two ideals are not equal, we explain the primary decomposition of the circuit ideal in terms of its initial ideals.

Chapter 4 is a study of *state polytopes* of toric ideals. These polytopes are of interest in optimization because they summarize families of integer programming problems: those that involve optimizing over any integer polytope

$$P_{\mathbf{b}} := \text{conv}\{\mathbf{x} \in \mathbb{N}^n : A\mathbf{x} = \mathbf{b}\}$$

where A is fixed but \mathbf{b} is allowed to vary. We give various conditions for a polytope to arise as a toric state polytope. We construct examples of toric state polytopes that include k -gons for every k and m -dimensional simplices and cubes for every m . Using a construction of Sturmfels and Thomas [49], we show that with the exception of simplices, state polytopes must be *decomposable*: that is, they must be nontrivial Minkowski sums of other polytopes. Among the consequences of this statement are the fact that *simplicial* polytopes (again except for simplices) cannot be toric state polytopes as well as a classification of three-dimensional toric state polytopes with few vertices.

Finally, Chapter 5 is joint work with Rekha Thomas in the theory of integer programming. Given a system of inequalities $Ax \leq \mathbf{b}$ defining a polytope P , the *Chvátal procedure* is an iterative method to obtain the *integer hull* of P ; that is, the convex hull of its integer points. The number of iterations required is known as the *Chvátal rank* of the system and is known to be finite [43, §23]. However, Chvátal rank is not just a function of the dimension of the polytope P and can be exponentially large, even in dimension two. Using a variant of the Chvátal procedure that depends only on the matrix A and not on any particular right-hand-side \mathbf{b} , we define a new parameter, the *small Chvátal rank* (SCR) of the system $Ax \leq \mathbf{b}$ and of A . This number is at most the Chvátal rank but in certain cases it is much smaller. In particular, in dimension two the SCR is always at most one. Another such case is the *clique polytope* of a complete graph K_n . However, we demonstrate that in dimension three, the SCR of a fixed-size matrix, like the Chvátal rank, can be exponentially large in the input size. We also prove a nontrivial lower bound on the SCR of polytopes contained in the n -dimensional unit cube. In the process of studying the computation of SCR, we answer a request of Hoşten, Maclagan, and Sturmfels [30, §2] for an infinite family of *supernormal* point configurations.

1.2 Background

Although each of the succeeding chapters concerns an independent project, many constructions and techniques recur throughout this thesis. Polytopes and polyhedral fans are objects of study in every chapter. Gröbner bases are a main tool in Chapters 3 and 4 and are needed to define and compute tropical varieties in Chapter 2. State polytopes and Gröbner fans summarize the set of *all* Gröbner bases of an ideal; they are the focus of Chapter 4 and also appear in sections of Chapters 3 and 2. Finally, toric ideals are a key element of Chapters 3 and 4.

1.2.1 Discrete geometry and optimization

We begin with some key definitions in polyhedral geometry with few results and no proofs. Two excellent general references on polytopes are the texts of Grünbaum [28] and Ziegler [53]. Schrijver's text [43] is an outstanding source on the theory of optimization (linear and integer programming) over polyhedra.

A *polytope* P is the convex hull of finitely many points v_1, \dots, v_n in a Euclidean space \mathbb{R}^m . A *polyhedron* is the intersection of closed half-spaces H_1^+, \dots, H_m^+ , where H_i^+ denotes one of the two regions bounded by a hyperplane H_i . It is a fundamental result that polytopes coincide with bounded polyhedra. We say that P is a d -polytope (or d -polyhedron) if the affine span of P is d -dimensional. A *face* of a polyhedron P is the intersection of P with any hyperplane H such that P is contained in one of the half-spaces H^+ defined by H . Every polytope P has a *dual* polytope P^* whose facets are in correspondence with the vertices of P , and vice versa.

If P is a polytope, then the faces of P may be ordered by inclusion to define the *face lattice* of P . The 0-dimensional faces are called *vertices*, the 1-dimensional faces *edges*, the $(d - 2)$ -dimensional faces *ridges*, and the $(d - 1)$ -dimensional faces *facets*. The vertices form the unique minimal set that has P as its convex hull. If P is an m -dimensional polytope in \mathbb{R}^m , then the facets are obtained as $H_1 \cap P, \dots, H_m \cap P$, where $P = H_1^+ \cap \dots \cap H_m^+$ is the unique minimal half-space description.

Two polytopes P and Q are *combinatorially isomorphic* if their face lattices are isomorphic. If so, we say that P is a *realization* of Q (or, more accurately, of its face lattice.) A coarser combinatorial invariant of a polytope P than its face lattice is its *f-vector* $(1, f_0, f_1, \dots, f_{d-1})$, where f_i is the number of i -dimensional faces of P . The question of which *f-vectors* can be achieved by various classes of polytopes is a major area of research.

A *d-simplex* is the convex hull of $d + 1$ affinely independent points. A *simplicial polytope* P is a polytope all of whose facets are simplices. The boundary of a simplicial

polytope is a *simplicial complex*, making its combinatorics especially nice: the face lattice is entirely determined by the incidence relation between facets and vertices. A *simple* polytope is one whose dual is simplicial.

A *polyhedral complex* is a finite set \mathcal{C} of polyhedra in \mathbb{R}^m such that

1. the empty polyhedron is in \mathcal{C} ,
2. if $P \in \mathcal{C}$, then every face of P is also in \mathcal{C} , and
3. if $P, Q \in \mathcal{C}$, then $P \cap Q$ is a face of both P and Q .

A (*polyhedral*) *cone* is a polyhedron C such that if \mathbf{x} and \mathbf{y} are in C , then so is $\lambda\mathbf{x} + \mu\mathbf{y}$ for every $\lambda, \mu \geq 0$. A cone is *pointed* if it does not contain both \mathbf{x} and $-\mathbf{x}$ for any $\mathbf{x} \neq 0$. A (*polyhedral*) *fan* is a polyhedral complex that contains $\{0\}$ and consists entirely of cones. A fan is *complete* if the union of its cones is the entire space \mathbb{R}^m . An important construction is the *normal fan* $N(P)$ of a polytope or polyhedron P . This is a fan in \mathbb{R}^m with one cone $C(F)$ for each face F of P ; this cone consists of the vectors whose maximum inner product with any point $p \in P$ is achieved exactly when p is in F . The fan $N(P)$ is complete if and only if P is a polytope. Figure 1.1 shows a triangle and its normal fan in \mathbb{R}^2 . We will study two particular types of fans: *tropical varieties* in Chapter 2 and *Gröbner fans* in various sections.

Optimization is a major application of polyhedral theory. Specifically, *linear programming* is the problem of maximizing a linear function over a polyhedron given by a half-space description. Similarly, *integer programming* consists of maximizing a linear function over the integer points of a polyhedron. Both problems are fundamental in mathematics as well as in applications. However, integer programming is a much more difficult problem: there exist polynomial-time algorithms for linear programming, but integer programming is NP-hard.

It is not surprising that integer programming turns out to be difficult, for many combinatorial optimization problems can easily be interpreted as integer programs.




Figure 1.1: A triangle with outer normal cones (left) and its outer normal fan (right)

For example, given a bipartite graph G on disjoint vertex sets V_1 and V_2 of equal size n , a *matching* is a subset S of the edges of G such that every vertex is incident to at most one edge in S . A matching S is *perfect* if every vertex is incident to exactly one edge in S . Let e_1, \dots, e_m denote the edges of G and define a polytope $P(G)$ in \mathbb{R}^m by

$$\begin{aligned} \sum_{e_i \in E(v)} x_i &\leq 1 && \text{for each vertex } v \\ x_i &\geq 0 && \text{for } 1 \leq i \leq m \end{aligned}$$

where $E(v)$ is the set of edges incident to a vertex v . The integer points of $P(G)$ index matchings and the linear function $x_1 + \dots + x_m$ on \mathbb{R}^m counts the number of edges in a matching. Thus to determine whether G has a perfect matching, it suffices to maximize $x_1 + \dots + x_m$ over the integer points of $P(G)$ and check whether the maximum value is n . This shows that general integer programs are at least as computationally difficult as the perfect matching problem.

In Chapter 3, we explore the extension of an algebraic approach to integer programming, via toric ideals, to a slightly wider class of ideals, including *circuit ideals*. Chapter 5 deals almost entirely with families of integer programs, and more specifically with an extension of the *cutting planes* method of Chvátal and Gomory. A cutting plane is a new inequality used to remove a non-integer vertex from a polyhe-

dron P ; the successive introduction of new cutting planes enables the computation of the integer hull P^I .

1.2.2 Gröbner bases

In this section we briefly introduce the theory of Gröbner bases, a fundamental tool in computational commutative algebra and algebraic geometry. Given the polynomial equations defining an affine or projective algebraic variety V , Gröbner bases can be used to solve the fundamental problems of identifying a point in V and of determining whether a given algebraic hypersurface contains V . The corresponding algebraic problems are *elimination and extension* and *ideal membership*; see [14], Chapters 2 and 3 respectively.

Let k be a field and set $k[\mathbf{x}] := k[x_1, \dots, x_n]$. We denote a monomial $x_1^{u_1} x_2^{u_2} \cdots x_n^{u_n}$ by $\mathbf{x}^{\mathbf{u}}$ where $\mathbf{u} = (u_1, \dots, u_n)$. A *monomial order* is a total ordering \prec on \mathbb{N}^n , the set of exponent vectors of monomials, satisfying the following conditions.

1. If $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathbb{N}^n$ and $\mathbf{u} \prec \mathbf{v}$, then $\mathbf{u} + \mathbf{w} \prec \mathbf{v} + \mathbf{w}$.
2. The ordering \prec is a *well-ordering*; that is, every descending chain eventually stabilizes.

When convenient, we will write $\mathbf{x}^{\mathbf{u}} \prec \mathbf{x}^{\mathbf{v}}$ instead of $\mathbf{u} \prec \mathbf{v}$.

Two common monomial orders are *lexicographic* (lex) and *graded reverse lexicographic* (grevlex). Lex order is given by $\mathbf{x}^{\mathbf{u}} \prec_{\text{lex}} \mathbf{x}^{\mathbf{v}}$ if the first nonzero component of $\mathbf{v} - \mathbf{u}$ is positive. Grevlex order is given by $\mathbf{x}^{\mathbf{u}} \prec_{\text{grevlex}} \mathbf{x}^{\mathbf{v}}$ if either the total degree of $\mathbf{x}^{\mathbf{v}}$ is greater than that of $\mathbf{x}^{\mathbf{u}}$, or if the two total degrees are equal and the last nonzero component of $\mathbf{v} - \mathbf{u}$ is negative. Of course, there are a total of $n!$ possible lex orders and $n!$ grevlex orders obtained by permuting the order of the variables.

Example 1.1. Let $\mathbf{u} = (3, 0, 3)$ and $\mathbf{v} = (1, 3, 2)$. Then $\mathbf{x}^{\mathbf{v}} \prec_{\text{lex}} \mathbf{x}^{\mathbf{u}}$ because $u_1 > v_1$, but $\mathbf{x}^{\mathbf{u}} \prec_{\text{grevlex}} \mathbf{x}^{\mathbf{v}}$ because both monomials have degree six and $u_3 > v_3$.

Note that the ungraded “revlex” order, given by $\mathbf{x}^{\mathbf{u}} \prec_{\text{revlex}} \mathbf{x}^{\mathbf{v}}$ if the last nonzero component of $\mathbf{v} - \mathbf{u}$ is negative, is not a well-ordering: we have $1 \succ_{\text{revlex}} x \succ_{\text{revlex}} x^2 \succ_{\text{revlex}} \dots$. Usually, though, we will work with homogeneous polynomials, so this problem will not concern us and we may simply refer to the grevlex order as reverse lexicographic.

Fix a monomial order \prec . The *initial monomial* of a polynomial $f(\mathbf{x}) := \sum_{\mathbf{u}} c_{\mathbf{u}} \mathbf{x}^{\mathbf{u}}$ is the greatest monomial $\mathbf{x}^{\mathbf{u}}$ with respect to \prec . The *initial form* of f is the term $\text{in}_{\prec}(f) := c_{\mathbf{u}} \mathbf{x}^{\mathbf{u}}$, where $\mathbf{x}^{\mathbf{u}}$ is the initial monomial. If I is an ideal in $k[\mathbf{x}]$, the *monomial initial ideal* of I with respect to \prec is the monomial ideal $\text{in}_{\prec}(I)$ generated by the initial terms of every polynomial in I . The idea is that we can study important properties of I via the much simpler monomial ideal $\text{in}_{\prec}(I)$.

Monomial ideals are indeed very special and easy to study. Dickson’s Lemma [14, Theorem 2.4.5] states that every monomial ideal has a finite monomial generating set. Given that such a set exists, it is easy to see that it is unique. Often it is convenient to work with the *standard monomials* of a monomial ideal M , simply defined to be the monomials that are not in M . Geometrically, a monomial ideal defines a union of coordinate planes. To keep track of the ideal properly, we must associate *multiplicities* to these planes; that is, we must strictly consider non-reduced affine schemes instead of affine varieties. This is a worthwhile trade-off for the advantage of working with such simple geometric objects as coordinate planes!

Example 1.2. Let $M_1 = \langle x^2, y^3, z^4, xy^2, xz^2, xyz \rangle$ and $M_2 = \langle x^2yz^2, xy^3z^2 \rangle$. The variety $V(M_2)$ is the union of the xy -plane, the xz -plane, and the yz -plane. More specifically, the scheme defined by M_2 consists of the xy -plane counted twice, the xz -plane once, the yz -plane once, and the line $z = 0$ (an embedded component) twice. This can be shown via *standard pair decomposition*, discussed in Chapter 3. The monomial ideal M_1 is artinian, so the variety $V(M_1)$ is just the origin with multiplicity 15, the number of standard monomials. The monomial ideals M_1 and M_2




Figure 1.2: The staircase diagrams of the monomial ideals $M_1 = \langle x^2, y^3, z^4, xy^2, xz^2, xyz \rangle$ and $M_2 = \langle x^2yz^2, xy^3z^2 \rangle$.

can be depicted by staircase diagrams (Figure 1.2, drawn by Gfan [34]).

Besides being relatively easy to handle, monomial initial ideals of an ideal I carry important information about I itself. The Krull dimension of $k[\mathbf{x}]/\text{in}_{\prec}(I)$ equals that of $k[\mathbf{x}]/I$. In fact, the two ideals are far more strongly linked than this. The *Hilbert function* of a homogeneous ideal I is given by $\text{HF}_I(s) = \dim(k[\mathbf{x}]_s/I_s)$ for $s \in \mathbb{N}$. In particular, the Hilbert function of a monomial ideal simply enumerates the standard monomials of each degree. There is also a *Hilbert polynomial*: a polynomial in s whose values agree with the Hilbert function for all sufficiently large s , and the degree of this polynomial is the Krull dimension of $k[\mathbf{x}]/I$.

Proposition 1.3. [14, Proposition 9.3.9] *If I is a homogeneous ideal in $k[\mathbf{x}]$ and \prec is a monomial order, then the Hilbert function of $\text{in}_{\prec}(I)$ equals that of I .*

Motivated by the connection between an ideal and its monomial initial ideals, we now define *Gröbner bases*. The idea goes back to the end of the 19th century, but it became widely known only after Buchberger [10] developed his famous algorithm.

Definition 1.4.

1. Given an ideal I and a monomial order \prec , a *Gröbner basis* for I with respect to \prec is a finite set $\mathcal{G} := \{g_1, \dots, g_s\} \subset I$ such that $\langle \text{in}_\prec(g_1), \dots, \text{in}_\prec(g_s) \rangle = \text{in}_\prec(I)$.
2. A Gröbner basis \mathcal{G} is *reduced* if the initial coefficient of every $g \in \mathcal{G}$ is one and if no term of any element of \mathcal{G} is divisible by the \prec -initial term of another element of \mathcal{G} .

Proposition 1.5. [14, Chapter 2] *Let $I \subseteq k[\mathbf{x}]$ be an ideal and \prec be a monomial order.*

1. *There exists a unique reduced Gröbner basis of I with respect to \prec .*
2. *Any Gröbner basis of I generates I as an ideal.*

Example 1.6. Let $f_1 = xz - y^2$, $f_2 = y - z$, and $I = \langle f_1, f_2 \rangle \subseteq k[x, y, z]$. In lex order, the leading monomials of f_1 and f_2 are xz and y , respectively. Since these monomials are coprime, it follows that f_1 and f_2 form a lex Gröbner basis for I (though coprimality is not a necessary condition.) On the other hand, in grevlex order the leading monomials are y^2 and y , and f_1 and f_2 do not form a grevlex Gröbner basis. The grevlex reduced Gröbner basis for I is $\{y - z, xz - z^2\}$.

Buchberger's algorithm to compute Gröbner bases is quite simple and is implemented in computer algebra systems such as Maple and Macaulay 2 [27]. However, its computational complexity is high and it may run quite slowly in practice. Some Gröbner bases for an ideal can be computed much faster than others; lex Gröbner bases tend to be especially slow. In the next section, we introduce Gröbner fans, which allow one Gröbner basis to be converted to another. This is often far more efficient than computing a lex Gröbner basis directly from Buchberger's algorithm.

The main step in Buchberger's algorithm is the computation of *S-polynomials*:

$$S(f, g) = \frac{\mathbf{x}^\gamma}{\text{in}_\prec(f)} \cdot f - \frac{\mathbf{x}^\gamma}{\text{in}_\prec(g)} \cdot g.$$

The algorithm terminates when no new initial terms can be created this way. We will employ *S-polynomials* in key proofs in Chapters 3 and 4.

1.2.3 The Gröbner fan and state polytope

Although there are infinitely many ways to define term orders, a given ideal $I \in k[\mathbf{x}]$ has only finitely many monomial initial ideals. The *Gröbner fan* of I , introduced by Mora and Robbiano in [38], is a full-dimensional fan whose maximal cones index these initial ideals. The lower-dimensional cones index non-monomial initial ideals. The cone containing the origin, always the unique lowest-dimensional cone of a fan, indexes I itself. To define general initial ideals we must change our focus from abstract term orders to those given by *weight vectors*. Our main source in this section and the next is Sturmfels' monograph [47].

Given a vector ω in the positive orthant $\mathbb{R}_{\geq 0}^n$ and a polynomial $f(x_1, \dots, x_n) = \sum c_{\mathbf{u}} \mathbf{x}^{\mathbf{u}}$, the *initial form* of f with respect to ω , denoted $\text{in}_\omega(f)$, is the sum of all terms $c_{\mathbf{u}} \mathbf{x}^{\mathbf{u}}$ such that $\omega \cdot \mathbf{u}$ is maximized. For most but not all ω , $\text{in}_\omega(f)$ will consist of a single term. If I is an ideal in $k[\mathbf{x}]$, define

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

Proposition 1.7. [47, Proposition 1.11] *Every monomial initial ideal $\text{in}_\prec(I)$ of an ideal I equals $\text{in}_\omega(I)$ for some $\omega \in \mathbb{R}_{\geq 0}^n$.*

The proof of Proposition 1.7 is based on the Farkas lemma of linear programming [53, §1.4], so polyhedral theory is brought into play. Individual cases of the proposition are not hard to verify: lex order can be represented by $(1, \epsilon, \epsilon^2, \dots, \epsilon^{n-1})$ and grevlex by $(1 - \epsilon, 1 - \epsilon^2, \dots, 1 - \epsilon^{n-1}, 1)$ for ϵ sufficiently small.

If I is a homogeneous ideal and $\omega \in \mathbb{R}_{\geq 0}^n$, then adding any constant multiple of the all-ones vector to ω will not alter the initial ideal $\text{in}_{\omega}(I)$. It follows that we can now choose weight vectors arbitrarily from \mathbb{R}^n , not just from the positive orthant. The same is true in the more general case where I is homogeneous with respect to any positive grading. From now on, we will generally restrict ourselves to this case to avoid unnecessary technicalities.

Given a homogeneous ideal I , define an equivalence relation on \mathbb{R}^n by

$$\omega_1 \sim \omega_2 \text{ if } \text{in}_{\omega_1}(I) = \text{in}_{\omega_2}(I).$$

Proposition 1.8. [47, §2]

1. *The closure of the equivalence class of any $\omega \in \mathbb{R}^n$ is a polyhedral cone, denoted C_{ω} .*
2. *These cones form a complete fan in \mathbb{R}^n , called the Gröbner fan of I .*
3. *Each full-dimensional cone contains a strictly positive vector and indexes a monomial initial ideal of I .*
4. *The Gröbner fan is the normal fan of a polytope. Any such polytope is called a state polytope of I .*
5. *If $I = \langle f \rangle$ is a principal ideal, then the Newton polytope of f is a state polytope of I .*

The Gröbner fan and the state polytope index all of the reduced Gröbner bases of an ideal. An algorithm ([47, Algorithm 3.6]) to compute Gröbner fans is implemented in the software package `Gfan` [34], which we use in many of our examples.

1.2.4 Toric ideals and integer programming

Chapters 3 and 4 rely on the method of Conti and Traverso [13] of using toric ideals and their state polytopes to solve families of integer programs. Fix $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{Z}^d$ and let A be the matrix whose columns are $\mathbf{a}_1, \dots, \mathbf{a}_n$. The *toric ideal* of \mathcal{A} is defined as

$$I_{\mathcal{A}} := \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} = \mathbf{u}^+ - \mathbf{u}^- \in \ker(A) \cap \mathbb{Z}^n \rangle.$$

Toric ideals are exactly the binomial prime ideals in the polynomial ring and give the defining equations of (affine, not necessarily normal) *toric varieties* [25]: varieties equipped with an action of a dense open algebraic torus $(\mathbb{C}^*)^d$. Often we will assume that the row space of A contains a strictly positive vector such as $(1, \dots, 1)$; then the toric ideal $I_{\mathcal{A}}$ will be homogeneous and can be used to define a *projective* toric variety $X_{\mathcal{A}}$. In this case, the stratification of $X_{\mathcal{A}}$ into torus orbits is captured by the faces of a polytope: the convex hull of the columns of A . This allows algebro-geometric problems to be solved by polyhedral methods. For example, resolution of singularities in $X_{\mathcal{A}}$ is achieved by certain subdivisions of the normal fan of the polytope; see [25, §2.6].

However, we will focus on the relation of toric ideals to optimization more than algebraic geometry. Given $\mathbf{b} \in \mathbb{Z}^d$, define a polyhedron $P_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}_{\geq 0}^n : A\mathbf{x} = \mathbf{b}\}$. Assume now that A does have a strictly positive vector in its row span; this will imply that each polyhedron $P_{\mathbf{b}}$ is in fact a *polytope*. The toric ideal $I_{\mathcal{A}}$ can be used to study the integer hull $P_{\mathbf{b}}^I$ of this polytope as follows.

Define a homomorphism $\pi : \mathbb{N}^n \rightarrow \mathbb{Z}^d$ by $\mathbf{u} \mapsto A\mathbf{u}$. A fiber $\pi^{(-1)}(\mathbf{b})$ consists exactly of the set $P_{\mathbf{b}} \cap \mathbb{N}^n$, so $P_{\mathbf{b}}^I = \text{conv}(\pi^{(-1)}(\mathbf{b}))$. Given a finite set $\mathcal{F} \subset \ker(\pi)$, define a graph $\pi^{(-1)}(\mathbf{b})_{\mathcal{F}}$ on each fiber $\pi^{(-1)}(\mathbf{b})$ by connecting two points \mathbf{u} and \mathbf{u}' if $\pm(\mathbf{u} - \mathbf{u}') \in \mathcal{F}$. To solve an integer program over a fiber, we need to move from an arbitrary point in the fiber, a given *feasible solution*, to another point, the *optimal solution*.

Proposition 1.9. [47, Theorem 5.3] *Let $\mathcal{F} \subset \ker(\pi)$. The graphs $\pi^{(-1)}(\mathbf{b})_{\mathcal{F}}$ are connected for all \mathbf{b} if and only if the set $\{\mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{F}\}$ generates the toric ideal $I_{\mathcal{A}}$.*

In fact, we want to do better: not only should the graph be connected, but also we should know which way to travel along each edge in order to reach the optimal solution without backtracking. If ω is the cost vector of the integer program, a Gröbner basis for $I_{\mathcal{A}}$ with respect to ω allows us to do just this [47, Theorem 5.5].

In Chapter 3, we develop analogues to this theory for certain subideals of toric ideals. We also use Gröbner bases to understand the primary decomposition of the subideals. In Chapter 4, we investigate the state polytopes of toric ideals; these too have an interpretation in integer programming, namely that their vertices index the different equivalence classes of cost vectors ω for the family of integer programs

$$\text{IP}_{A,\omega}(\mathbf{b}) := \text{minimize } \{\omega \cdot \mathbf{u} : A\mathbf{u} = \mathbf{b}, \mathbf{u} \in \mathbb{N}^n\}$$

for all \mathbf{b} [49]. Chapter 5 was also motivated by this application of toric Gröbner fans.

Chapter 2

TROPICAL VARIETIES

Tropical geometry, the geometry of the min-plus semiring, has flourished in recent years. Its wide-ranging applications include dynamics [17], and statistics [39]. Tropical geometry can also be used to prove results in ordinary algebraic geometry: two examples are compactifications of affine algebraic varieties [51] and enumeration of algebraic curves in toric surfaces [37].

In the first section of this chapter, we briefly overview the subject. The remainder consists of a joint paper [8] with Anders Jensen, David Speyer, Bernd Sturmfels, and Rekha Thomas on computational aspects of tropical geometry and implementations of the algorithms in the program `Gfan` [34]. Most of our paper is included either verbatim or with only minor changes in this chapter.

2.1 Tropical Algebra and Geometry

The *tropical semiring* is $(\mathbb{R}, \oplus, \otimes)$, where $x \oplus y = \min(x, y)$ and $x \otimes y = x + y$. Sometimes it is convenient to adjoin ∞ as an additive identity, but there can be no additive inverses. The tropical operations extend to give \mathbb{R}^n the structure of a semi-vector space. For example, $3 \oplus 5 = 3$, $3 \otimes 5 = 8$, and

$$(1, 4, 5) \oplus (2 \otimes (1, 1, -6)) = (1, 4, 5) \oplus (3, 3, -4) = (1, 3, -4).$$

Any arithmetic not written with the symbols \oplus or \otimes should be interpreted as ordinary, not tropical.

A tropical polynomial F in n variables X_1, \dots, X_n (in ordinary arithmetic, a minimum of sums) thus represents a piecewise linear function from \mathbb{R}^n to \mathbb{R} . Its

tropical hypersurface $\mathcal{T}(F)$ is its non-differentiability locus, which is the subset of \mathbb{R}^n where the minimum is achieved at least twice. This set consists of the cones of positive codimension in the inner normal fan of the Newton polytope of F . Unlike in ordinary algebraic geometry, it does not make sense to define a tropical hypersurface as the points where the polynomial takes the value zero, as zero is not the tropical additive identity.

Example 2.1. Consider the tropical polynomial $F(X, Y) = X \oplus (2 \otimes Y) \oplus 3 = \min(X, 2 + Y, 3)$. Its tropical hypersurface is the subset of \mathbb{R}^2 where the minimum occurs at least twice, which is

$$\{(X, Y) \in \mathbb{R}^2 : X = 2 + Y \leq 3 \text{ or } X = 3 \leq 2 + Y \text{ or } 2 + Y = 3 \leq X\}.$$

This set (Figure 2.1) is the union of three rays emanating from the point (3,1) where the minimum is achieved by all three terms. The Newton polygon of F (Figure 2.2) is the triangle with vertices (1, 0), (0, 1), and (0, 0).

Varying the coefficients in this example, we see that all hyperplanes in \mathbb{R}^2 - in fact all hyperplanes in \mathbb{R}^n for a fixed n - are translates of each other. Unlike ordinary Euclidean space, tropical space has only translations, not rotations or reflections, as its natural symmetries.

To extend tropical algebraic geometry from hypersurfaces to more general varieties, we must take a more indirect approach as follows. Given an algebraically closed field K with a non-archimedean valuation $\deg : K^* \mapsto \mathbb{R}$ and a polynomial $f \in K[\mathbf{x}]$ given as

$$f(x_1, \dots, x_n) = \sum_{\mathbf{a}} c_{\mathbf{a}} x_1^{a_1} \cdots x_n^{a_n},$$

we can define a tropical polynomial F by the formula

$$F(X_1, \dots, X_n) = \bigoplus_{\mathbf{a}} \deg(c_{\mathbf{a}}) \otimes a_1 X_1 \otimes \cdots \otimes a_n X_n.$$




Figure 2.1: The tropical hyperplane in \mathbb{R}^2 of $F(X, Y) = X \oplus (2 \otimes Y) \oplus 3$.




Figure 2.2: The Newton polygon of Example 2.1 with the rays of its inner normal fan.

We can then associate to f the tropical hypersurface $\mathcal{T}(f) := \mathcal{T}(F)$, as in Example 2.2 below.

The usual choices for K are the Puiseux series field $\mathbb{C}\{\{t\}\}$ with the valuation $\deg(ct^b + \text{higher order terms}) = b$, or the algebraic closure of the p -adic field \mathbb{Q}_p with the p -adic valuation. The Puiseux series field, defined as the set of power series in rational powers of t such that the exponents are bounded below and have only bounded denominators, arises naturally in solving a system of polynomial equations with time-dependent coefficients. The first condition for a parametric curve $\phi(t) = (x_1(t), x_2(t), \dots, x_n(t))$ to satisfy such a system is that the lowest-order terms in t cancel out for all t , which is to say that the system is satisfied tropically (see below.) In our current work, however, we focus on the case in which the polynomials have only complex coefficients, and the program `Gfan` accepts only rational input.

Example 2.2. Take $K = \mathbb{C}\{\{t\}\}$ and define $f \in K[x, y]$ by $f(x, y) = x + 100t^2y + (t^3 - 2t^4)$. Its tropicalization is the polynomial F of Example 2.1, so $\mathcal{T}(f)$ is the union of rays shown in Figure 2.1. If we use constant coefficients, then there is no polynomial whose tropicalization is F because the coefficients 2 and 3 cannot be obtained from the zero valuation, the only non-archimedean valuation on \mathbb{C} .

The next step is to tropicalize not just individual polynomials, but whole ideals in $K[\mathbf{x}]$. The cornerstone of tropical geometry is the following theorem that states the equivalence of three definitions of a tropical variety $\mathcal{T}(I)$.

Theorem 2.3. [46, Theorem 2.1] *For an ideal $I \subseteq K[\mathbf{x}]$, the following subsets of \mathbb{R}^n coincide.*

1. *The closure of the set $\{(\deg(u_1), \dots, \deg(u_n)) : (u_1, \dots, u_n) \in V(I)\}$, where $V(I)$ is the variety of I in the algebraic torus $(K^*)^n$;*
2. *The intersections of the tropical hypersurfaces $\mathcal{T}(f)$ where $f \in I$;*

3. The set of all vectors $\omega \in \mathbb{R}^n$ such that the initial ideal $\text{in}_\omega(I)$ contains no monomial.

The first definition connects tropical to ordinary varieties and has the important consequence that $\mathcal{T}(I)$ and $V(I)$ have the same dimension. The second definition shows that tropical varieties are the intersections of tropical hypersurfaces. The third is the most useful for computation, and implies that in the constant-coefficient case, $\mathcal{T}(I)$ is a subfan of the Gröbner fan of I .

More precisely, the tropical variety is a subfan of the *negative* of the Gröbner fan, rather than the Gröbner fan itself. To avoid a profusion of minus signs, in the remainder of this chapter we will take the initial form of a polynomial f with respect to a vector ω to consist of the sum of the terms whose exponent vectors have the *smallest* inner product with ω , rather than the largest. Since all of the ideals in this chapter are homogeneous, this will not cause any problems. (Alternatively, tropical arithmetic can be defined with maximum rather than minimum as its additive operation.)

To compute an ordinary algebraic variety $V(I)$, we do not need to intersect the hypersurfaces defined by every polynomial in I ; any generating set will do. The situation for $\mathcal{T}(I)$ is different. For example, take $K = \mathbb{C}\{\{t\}\}$ and $I = \langle x + y + t, x + y + t^3 \rangle$. As we observed in Example 2.1, each of the tropical hypersurfaces of the given generators consists of three rays whose direction vectors are $(1, 0)$, $(0, 1)$, and $(-1, -1)$. The intersection of these two hypersurfaces (Figure 2.3) is a ray. The difference of the two generators, however, is $t - t^3$, so it follows from Theorem 2.3 (3) that $\mathcal{T}(I)$ is empty.

This motivates the following definitions that have no analogue for ordinary algebraic varieties.

Definition 2.4.

1. The *tropical prevariety* of a finite set of polynomials is the intersection of their tropical hypersurfaces.




Figure 2.3: Two tropical hyperplanes $\mathcal{T}(x+y+t)$ and $\mathcal{T}(x+y+t^3)$. Their intersection, shown as a thick ray, is not a tropical variety.

2. A finite set $\{f_1, \dots, f_r\}$ of polynomials is called a *tropical basis* of the ideal I that they generate if the tropical prevariety $\mathcal{T}(f_1) \cap \dots \cap \mathcal{T}(f_r)$ equals $\mathcal{T}(I)$.

2.2 Computing Tropical Varieties

2.2.1 Introduction

Our main contribution to tropical geometry is a practical algorithm, along with its implementation, for computing the tropical variety $\mathcal{T}(I)$ from any generating set of its ideal I . The emphasis lies on the geometric and algebraic features of this computation. We do not address issues of computational complexity, which have been studied by Theobald [52].

In Section 2.2.2 we give precise specifications of the algorithmic problems we are dealing with, including the computation of a tropical basis. We show that a finite tropical basis exists for every ideal I , and we give tight bounds on the size of a tropical basis for linear ideals, thereby answering the question raised in [45, §5, page 13]. In Section 2.2.3 we prove that the tropical variety $\mathcal{T}(I)$ of a prime ideal is connected in codimension one. This result is the foundation of Algorithm 2.30 for

computing $\mathcal{T}(I)$. Section 2.2.4 also describes methods for computing tropical bases and tropical prevarieties. These algorithms have been implemented in the software package `Gfan`. In Section 2.2.5 we compute the tropical variety of several non-trivial ideals using `Gfan`. The tropical variety $\mathcal{T}(I)$ is a subfan of the Gröbner fan of I , but the Gröbner fan is generally much more complicated and harder to compute than $\mathcal{T}(I)$. In Section 2.2.6 we compare these two fans, and we exhibit a family of curves for which the tropical variety of each member consists of four rays but the number of rays in the Gröbner fan grows arbitrarily. The proof of this result appears in more detail in this thesis than in the paper.

A note on the choice of ground field \mathbb{C} is in order. In the implementation of our algorithm (Section 2.2.5), we have required our polynomials to have rational coefficients so that they may be given by finite input, but our algorithms do not use any particular properties of \mathbb{Q} . It is important, however, that we work over a field of characteristic 0, as our proof of correctness uses the Kleiman-Bertini theorem at one point.

In most papers on tropical algebraic geometry (cf. [17, 37, 41, 46, 52]), tropical varieties are defined from polynomials with coefficients in a field K with a non-archimedean valuation. These tropical varieties are not fans but polyhedral complexes. We close the introduction by illustrating how our algorithms can be applied to this situation. Consider the field $\mathbb{C}(\epsilon)$ of rational functions in the unknown ϵ . Then $\mathbb{C}(\epsilon)$ is a subfield of the algebraically closed field $\mathbb{C}\{\{\epsilon\}\}$ of Puiseux series with real exponents, which is an example of a field K as in the above cited papers. Suppose we are given an ideal I in $\mathbb{C}(\epsilon)[x_1, \dots, x_n]$. Let $I' \subset \mathbb{C}\{\{\epsilon\}\}[x_1, \dots, x_n]$ be the ideal generated by I . The tropical variety $\mathcal{T}(I')$, in the sense of the papers above, is a finite polyhedral complex in \mathbb{R}^n which may have both bounded and unbounded faces. To study this complex, we consider the polynomial ring in $n+1$ variables, $\mathbb{C}[\epsilon, x_1, \dots, x_n]$ and we let J denote the intersection of I with this subring of $\mathbb{C}(\epsilon)[x_1^{\pm}, \dots, x_n^{\pm}]$. Generators of J are computed from generators of I by clearing denominators and saturating

with respect to ϵ . The tropical variety of I' is related to the tropical variety of J as follows.

Lemma 2.5. *A vector $\omega \in \mathbb{R}^n$ lies in the polyhedral complex $\mathcal{T}(I')$ if and only if the vector $(1, \omega) \in \mathbb{R}^{n+1}$ lies in the polyhedral fan $\mathcal{T}(J)$.*

2.2.2 Algorithmic Problems and Tropical Bases

For all computational problems addressed in this paper we fix the ambient ring to be the polynomial ring over the complex numbers, $\mathbb{C}[\mathbf{x}] := \mathbb{C}[x_1, \dots, x_n]$. The most basic computational problem in tropical geometry is the following:

Problem 2.6. *Given a finite list of polynomials $f_1, \dots, f_r \in \mathbb{C}[\mathbf{x}]$, compute the tropical prevariety $\mathcal{T}(f_1) \cap \dots \cap \mathcal{T}(f_r)$ in \mathbb{R}^n .*

The geometry of this problem is best understood by considering the Newton polytopes $\text{New}(f_1), \dots, \text{New}(f_r)$ of the given polynomials. Recall that $\text{New}(f_i)$ is the convex hull in \mathbb{R}^n of the exponent vectors which appear with non-zero coefficient in f_i . The tropical hypersurface $\mathcal{T}(f_i)$ is the $(n-1)$ -skeleton of the inner normal fan of the polytope $\text{New}(f_i)$. Our problem is to intersect these normal fans. The resulting tropical prevariety can be a fairly general polyhedral fan. Its maximal cones may have different dimensions.

In contrast, Bieri and Groves [6] proved that $\mathcal{T}(I)$ is a d -dimensional fan when d is the Krull dimension of $\mathbb{C}[\mathbf{x}]/I$. The fan is pure if I is unmixed. In Section 3 we shall prove that $\mathcal{T}(I)$ is connected in codimension one if I is prime.

We first note that it suffices to devise algorithms for computing tropical varieties of homogeneous ideals. Let ${}^h I \subset \mathbb{C}[x_0, x_1, \dots, x_n]$ be the homogenization of an ideal I in $\mathbb{C}[\mathbf{x}]$ and ${}^h f$ the homogenization of $f \in \mathbb{C}[\mathbf{x}]$.

Lemma 2.7. *Fix an ideal $I \subset \mathbb{C}[\mathbf{x}]$ and a vector $\omega \in \mathbb{R}^n$. The initial ideal $\text{in}_\omega(I)$ contains a monomial if and only if $\text{in}_{\prec(0, \omega)}({}^h I)$ contains a monomial.*

Proof: Suppose $\mathbf{x}^{\mathbf{u}} \in \text{in}_{\omega}(I)$. Then $\mathbf{x}^{\mathbf{u}} = \text{in}_{\omega}(f)$ for some $f \in I$. The $(0, \omega)$ -weight of a term in ${}^h f$ equals the ω -weight of the corresponding term in f . Hence $\text{in}_{\prec(0, \omega)}({}^h f) = x_0^a \mathbf{x}^{\mathbf{u}} \in \text{in}_{\prec(0, \omega)}({}^h I)$ where a is some non-negative integer.

Conversely, if $\mathbf{x}^{\mathbf{u}} \in \text{in}_{\prec(0, \omega)}({}^h I)$ then $\mathbf{x}^{\mathbf{u}} = \text{in}_{\prec(0, \omega)}(f)$ for some $f \in {}^h I$. Substituting $x_0 = 1$ in f gives a polynomial in I . The $(0, \omega)$ -weight of any term in f equals the ω -weight of the corresponding term in $f|_{x_0=1}$. Since $\text{in}_{\prec(0, \omega)}(f)$ is a monomial, only one term in f has minimal $(0, \omega)$ -weight. This term cannot be canceled during the substitution. Hence it lies in $\text{in}_{\omega}(I)$. \square

Our main goal is to solve the following computational task.

Problem 2.8. *Given a finite list of homogeneous polynomials $f_1, \dots, f_r \in \mathbb{C}[\mathbf{x}]$, compute the tropical variety $\mathcal{T}(I)$ of their ideal $I = \langle f_1, \dots, f_r \rangle$.*

It is important to note that the two problems stated so far are of a fundamentally different nature. Problem 2.6 is a problem of polyhedral geometry. It involves only polyhedral computations: no algebraic computations are required. Problem 2.8, on the other hand, combines the polyhedral aspect with an algebraic one. To solve Problem 2.8 we must perform algebraic operations (e.g. Gröbner bases) with polynomials. In Problem 2.6 we do not assume that the input polynomials f_1, \dots, f_r are homogeneous as the polyhedral computations can be performed easily without this assumption.

Proposition 2.9. *Let I be an ideal in $\mathbb{C}[\mathbf{x}]$ and let $\omega \in \mathbb{R}^n$. The following are equivalent:*

1. *The ideal I is ω -homogeneous; i.e. I is generated by a set S of ω -homogeneous polynomials, meaning that $\text{in}_{\omega}(f) = f$ for all $f \in S$.*
2. *The initial ideal $\text{in}_{\omega}(I)$ is equal to I .*

Proof: If I has a ω -homogeneous generating set then $I \subset \text{in}_\omega(I)$. Any maximal ω -homogeneous component of $f \in I$ is in I . In particular $\text{in}_\omega(f) \in I$. Conversely, the ideal $\text{in}_\omega(I)$ is generated by ω -homogeneous elements by definition, so if $I = \text{in}_\omega(I)$, then I is generated by ω -homogeneous elements. \square

The set of $\omega \in \mathbb{R}^n$ for which the above equivalent conditions hold is a vector subspace of \mathbb{R}^n . Its dimension is called the *homogeneity* of I and is denoted $\text{homog}(I)$. This space is contained in every cone of the fan $\mathcal{T}(I)$ and can be computed from the Newton polytopes of the polynomials that form any reduced Gröbner basis of I . Passing to the quotient of \mathbb{R}^n modulo that subspace and then to a sphere around the origin, $\mathcal{T}(I)$ can be represented as a polyhedral complex of dimension $n - \text{codim}(I) - \text{homog}(I) - 1 = \dim(I) - \text{homog}(I) - 1$. Here $\text{codim}(I)$ and $\dim(I)$ are the codimension and dimension of I . In what follows, $\mathcal{T}(I)$ is always presented in this way, and every ideal I is presented by a finite list of generators together with the three numbers n , $\dim(I)$, and $\text{homog}(I)$.

Example 2.10. Let I denote the ideal which is generated by the 3×3 -minors of a symmetric 4×4 -matrix of unknowns. This ideal has $n = 10$, $\dim(I) = 7$ and $\text{homog}(I) = 4$. Hence $\mathcal{T}(I)$ is a two-dimensional polyhedral complex, which we regard as the tropicalization of the secant variety of the Veronese threefold in \mathbb{P}^9 , i.e., the variety of symmetric 4×4 -matrices of rank at most two. Applying our `Gfan` implementation (see Example 2.35), we find that $\mathcal{T}(I)$ is a simplicial complex consisting of 75 triangles, 75 edges and 20 vertices. \square

Our next problem concerns tropical bases. Recall that a finite set $\{f_1, \dots, f_t\}$ is a tropical basis of I if $\langle f_1, \dots, f_t \rangle = I$ and $\mathcal{T}(I) = \mathcal{T}(f_1) \cap \dots \cap \mathcal{T}(f_t)$.

Problem 2.11. *Compute a tropical basis of a given ideal $I \subset \mathbb{C}[\mathbf{x}]$.*

A priori, it is not clear that every ideal I has a finite tropical basis, but we shall prove this below. First, here is one case where this is easy:

Example 2.12. If $I = \langle f \rangle$ is a principal ideal, then $\{f\}$ is a tropical basis. \square

In [46] it was claimed that any universal Gröbner basis of I is a tropical basis. Unfortunately, this claim is false as the following example shows.

Example 2.13. Let I be the intersection of the three linear ideals $\langle x+y, z \rangle$, $\langle x+z, y \rangle$, and $\langle y+z, x \rangle$ in $\mathbb{C}[x, y, z]$. Then I contains the monomial xyz , so $\mathcal{T}(I)$ is empty. A minimal universal Gröbner basis of I is

$$\mathcal{U} = \{x + y + z, x^2y + xy^2, y^2z + yz^2, x^2z + xz^2\},$$

and the intersection of the four corresponding tropical surfaces in \mathbb{R}^3 is the line $\omega_1 = \omega_2 = \omega_3$. Thus \mathcal{U} is not a tropical basis of I . \square

We now prove that every ideal $I \subset \mathbb{C}[\mathbf{x}]$ has a tropical basis. By Lemma 2.7, one tropical basis of a non-homogeneous ideal I is the dehomogenization of a tropical basis for ${}^h I$. Hence we shall assume that I is a homogeneous ideal.

Tropical bases can be constructed from the Gröbner fan of I . The tropical variety $\mathcal{T}(I)$ consists of all Gröbner cones $C_\omega(I)$ such that $\text{in}_\omega(I)$ does not contain a monomial. From the description of $\mathcal{T}(I)$ as $\bigcap_{f \in I} \mathcal{T}(f)$, it is clear that $\mathcal{T}(I)$ is closed. Thus we deduce that $\mathcal{T}(I)$ is a closed subfan of the Gröbner fan. This endows the tropical variety $\mathcal{T}(I)$ with the structure of a polyhedral fan.

Theorem 2.14. *Every ideal $I \subset \mathbb{C}[\mathbf{x}]$ has a tropical basis.*

Proof: Let \mathcal{F} be any finite generating set of I which is not a tropical basis. Pick a Gröbner cone $C_\omega(I)$ whose relative interior intersects $\bigcap_{f \in \mathcal{F}} \mathcal{T}(f)$ non-trivially and whose initial ideal $\text{in}_\omega(I)$ contains a monomial \mathbf{x}^m . Compute the reduced Gröbner basis $\mathcal{G}_{\prec_\omega}(I)$ for a refinement \prec_ω of ω , and let h be the normal form of \mathbf{x}^m with respect to $\mathcal{G}_{\prec_\omega}(I)$. Let $f := \mathbf{x}^m - h$. Since the normal form of \mathbf{x}^m with respect to $\mathcal{G}_{\prec_\omega}(\text{in}_\omega(I)) = \{\text{in}_\omega(g) : g \in \mathcal{G}_{\prec_\omega}(I)\}$ is 0 and h is the normal form of \mathbf{x}^m with respect to $\mathcal{G}_{\prec_\omega}(I)$, every monomial occurring in h has higher ω -weight than \mathbf{x}^m .

Moreover, h depends only on the reduced Gröbner basis $\mathcal{G}_{\prec_\omega}(I)$ and is independent of the particular choice of ω in $C_\omega(I)$. Hence for any ω' in the relative interior of $C_\omega(I)$, we have $\mathbf{x}^m = \text{in}_{\omega'}(f)$. This implies that the polynomial $f := \mathbf{x}^m - h$ is a *witness* for the cone $C_\omega(I)$ not being in the tropical variety $\mathcal{T}(I)$.

We now add the witness f to the current basis \mathcal{F} and repeat the process. Since the Gröbner fan has only finitely many cones, this process terminates. It removes all cones of the Gröbner fan which violate the condition for \mathcal{F} to be a tropical basis. \square

We next show that tropical bases can be very large even for linear ideals. Let I be the ideal in $\mathbb{C}[\mathbf{x}]$ generated by d linear forms $\sum_{j=1}^n a_{ij}x_j$ where $i = 1, \dots, d$ and (a_{ij}) is an integer $d \times n$ matrix of rank d . The tropical variety $\mathcal{T}(I)$ depends only on the matroid associated with I , and it is known as the *Bergman fan* of that matroid. The results on the Bergman fan proved in [3, 48] imply that the circuits in I form a tropical basis. A *circuit* of I is a non-zero linear polynomial $f \in I$ of minimal support. The following result answers the question which was posed in [45, §5].

Theorem 2.15. *For any $1 \leq d \leq n$, there is a linear ideal I in $\mathbb{C}[\mathbf{x}]$ such that any tropical basis of linear forms in I has size at least $\frac{1}{n-d+1} \binom{n}{d}$.*

Proof: Suppose that all $d \times d$ -minors of the coefficient matrix (a_{ij}) are non-zero. Equivalently, the matroid of I is uniform. There are $\binom{n}{n-d+1}$ circuits in I , each supported on a different $(n-d+1)$ -subset of $\{x_1, \dots, x_n\}$. Since the circuits form a tropical basis of I and each circuit has support of size $n-d+1$, the tropical variety $\mathcal{T}(I)$ consists of all vectors $\omega \in \mathbb{R}^n$ whose smallest $d+1$ components are equal. The latter condition is necessary and sufficient to ensure that no single variable in a circuit becomes the initial form of the circuit with respect to ω . Consider any vector $\omega \in \mathbb{R}^n$ satisfying

$$\omega_{i_1} = \omega_{i_2} = \dots = \omega_{i_d} < \min(\omega_j : j \in \{1, \dots, n\} \setminus \{i_1, i_2, \dots, i_d\}).$$

Since $\omega \notin \mathcal{T}(I)$, any tropical basis of linear forms in I contains an element f such that $\text{in}_\omega(f) \in \{x_{i_1}, \dots, x_{i_d}\}$. This implies that f is one of the d circuits whose support

contains the $n - d$ variables x_j with $j \notin \{i_1, \dots, i_d\}$. The support of each circuit has size $n - d + 1$, hence contains $n - d + 1$ distinct $(n - d)$ -subsets. There are $\binom{n}{n-d}$ $(n - d)$ -subsets of $\{x_1, \dots, x_n\}$ to be covered. Hence any tropical basis consisting of linear forms has size at least $\frac{1}{n-d+1} \binom{n}{n-d}$. \square

Example 2.16. Let $d = 3, n = 5$. The Bergman fan $\mathcal{T}(I)$ corresponds to the line in tropical projective 4-space which consists of the five rays in the coordinate directions. We have $\frac{1}{n-d+1} \binom{n}{n-d} = 10/3$. Hence this line is not a complete intersection of three tropical hyperplanes, but requires four. \square

2.2.3 Transversality and Connectivity

In this section we assume that I is a prime ideal of dimension d in $\mathbb{C}[\mathbf{x}]$. Then its tropical variety $\mathcal{T}(I)$ is called *irreducible*. It is a subfan of the Gröbner fan of I and, by the Bieri-Groves Theorem [6, 48], all facets of $\mathcal{T}(I)$ are cones of dimension d . A *ridge path* is a sequence of facets F_1, F_2, \dots, F_k such that $F_i \cap F_{i+1}$ is a ridge for all $i \in \{1, 2, \dots, k - 1\}$. Our objective is to prove the following result, which is crucial for the algorithms.

Theorem 2.17. *Any irreducible tropical variety $\mathcal{T}(I)$ is connected in codimension one; i.e., any two facets are connected by a ridge path.*

The proof of this theorem will be based on the following important lemma.

Lemma 2.18. *Transverse Intersection Lemma*

Let I and J be ideals in $\mathbb{C}[x_1, \dots, x_n]$ whose tropical varieties $\mathcal{T}(I)$ and $\mathcal{T}(J)$ meet transversally at a point $\omega \in \mathbb{R}^n$. Then $\omega \in \mathcal{T}(I + J)$.

By “meet transversely” we mean that if F and G are the cones of $\mathcal{T}(I)$ and $\mathcal{T}(J)$ which contain ω in their relative interior, then $\mathbb{R}F + \mathbb{R}G = \mathbb{R}^n$.

This lemma implies that any transverse intersection of tropical varieties is a tropical variety. In particular, any transverse intersection of tropical hypersurfaces is a

tropical variety, and such a tropical variety is defined by an ideal which is a complete intersection in the ordinary commutative algebra sense.

Corollary 2.19. *For any two ideals I and J in $\mathbb{C}[\mathbf{x}]$ we have*

$$\mathcal{T}(I + J) \subseteq \mathcal{T}(I) \cap \mathcal{T}(J).$$

Equality holds if the latter intersection is transverse at every point except the origin and the two fans meet in at least one point other than the origin.

Proof: We have $\mathcal{T}(I) \cap \mathcal{T}(J) = \bigcap_{f \in I} \mathcal{T}(f) \cap \bigcap_{f \in J} \mathcal{T}(f) = \bigcap_{f \in I+J} \mathcal{T}(f)$. Clearly, this contains $\mathcal{T}(I + J) = \bigcap_{f \in I+J} \mathcal{T}(f)$. If $\mathcal{T}(I)$ and $\mathcal{T}(J)$ intersect transversally and ω is a point of $\mathcal{T}(I) \cap \mathcal{T}(J)$ other than the origin then the preceding lemma tells us that $\omega \in \mathcal{T}(I + J)$. Thus $\mathcal{T}(I + J)$ contains every point of $\mathcal{T}(I) \cap \mathcal{T}(J)$ except possibly the origin. In particular, $\mathcal{T}(I + J)$ is not empty. Every nonempty fan contains the origin, so we see that the origin is in $\mathcal{T}(I + J)$ as well. \square

We first derive Theorem 2.17 from Lemma 2.18, which will be proved later. We must at this point address a technical detail. The subset $\mathcal{T}(I) \subset \mathbb{R}^n$ depends only on the ideal $IC[\mathbf{x}^{\pm 1}]$ generated by I in the Laurent polynomial ring $\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. From a theoretical perspective, then, it would be better to directly work with ideals in $\mathbb{C}[\mathbf{x}^{\pm 1}]$. Computationally, however, it is much better to deal with ideals in $\mathbb{C}[\mathbf{x}]$ as it is for such ideals that Gröbner basis techniques have been developed and this is the approach we take in the rest of the paper.

There is one great advantage to working with $\mathbb{C}[\mathbf{x}^{\pm 1}]$ however: we have available the symmetry group $GL_n(\mathbb{Z})$ of the multiplicative group of monomials. The action of this group transforms $\mathcal{T}(I)$ by the obvious action on \mathbb{R}^n . This symmetry will prove invaluable for simplifying the arguments in this section. Therefore, in the rest of this section, we will work with ideals in $\mathbb{C}[\mathbf{x}^{\pm 1}]$. Note that, if $I \subset \mathbb{C}[\mathbf{x}]$ is prime then so is the ideal it generates in $\mathbb{C}[\mathbf{x}^{\pm 1}]$. We will signify an application of the $GL_n(\mathbb{Z})$ symmetry by the phrase “making a multiplicative change of variables”. The

polyhedral structure on $\mathcal{T}(I)$ induced by the Gröbner fan of I may change under a multiplicative change of variables of $\mathbb{C}[\mathbf{x}^\pm]$ in $\mathbb{C}[\mathbf{x}^{\pm 1}]$, but all of the properties of $\mathcal{T}(I)$ that are of interest to us depend only on the underlying point set.

Proof of Theorem 2.17. As described above, we replace I by the ideal it generates in $\mathbb{C}[\mathbf{x}^{\pm 1}]$ and, by abuse of notation, continue to denote this ideal as I . The proof proceeds by induction on $d = \dim(\mathcal{T}(I))$. If $d \leq 1$ then the statement is trivially true. We now explain why the result holds when $d = 2$. By a multiplicative change of coordinates it suffices to check that $\mathcal{T}(I) \cap \{x_n = 1\}$ is connected.

Let K be the Puiseux series field over \mathbb{C} . Let $I' \subset K[x_1, \dots, x_{n-1}]$ be the prime ideal generated by I via the inclusion $\mathbb{C}[x_n] \rightarrow K$. By Lemma 2.5, the tropical variety of I' is $\mathcal{T}(I) \cap \{x_n = 1\}$. In [17] it was shown that the tropical variety of I' is connected whenever I' is prime. We conclude that $\mathcal{T}(I) \cap \{x_n = 1\}$ is connected, so our result holds for $d = 2$.

We now suppose that $d \geq 3$. Let F and F' be facets of $\mathcal{T}(I)$. We can find

$$H = \{ (u_1, \dots, u_n) \in \mathbb{R}^n : a_1 u_1 + \dots + a_n u_n = 0 \}$$

such that a_1, \dots, a_n are relatively prime integers, both $H \cap F$ and $H \cap F'$ are cones of dimension $d-1$, and H intersects every cone of $\mathcal{T}(I)$ except for the origin transversally. To see this, select rays ω and ω' in the relative interiors of F and F' . By perturbing ω and ω' slightly, we may arrange that the span of ω and ω' does not meet any ray of $\mathcal{T}(I)$ – here it is important that $d \geq 3$. Now, taking H to be the span of ω , ω' and a generic $(n-3)$ -plane, we get that H also does not contain any ray of $\mathcal{T}(I)$ and hence does not contain any positive dimensional face of $\mathcal{T}(I)$. So H is transverse to $\mathcal{T}(I)$ everywhere except at the origin. Since $H \cap F$ and $H \cap F'$ are positive dimensional (as $d \geq 2$) H does intersect $\mathcal{T}(I)$ at points other than the origin. The hyperplane H is the tropical hypersurface of a binomial, namely, $H = \mathcal{T}(\langle f_u \rangle)$, where

$$f_u = \prod_{i:a_i>0} (u_i x_i)^{a_i} - \prod_{j:a_j<0} (u_j x_j)^{-a_j},$$

and $u = (u_1, u_2, \dots, u_n)$ is an arbitrary point in the algebraic torus $(\mathbb{C}^*)^n$. Our transversality assumption regarding H and Lemma 2.18 imply that

$$H \cap \mathcal{T}(I) = \mathcal{T}(\langle f_u \rangle) \cap \mathcal{T}(I) = \mathcal{T}(I + \langle f_u \rangle). \quad (2.1)$$

Since I is prime of dimension d , and $f_u \notin I$, the ideal $I + \langle f_u \rangle$ has dimension $d - 1$ by Krull's Principal Ideal Theorem [19, Theorem 10.1]. If $I + \langle f_u \rangle$ were a prime ideal then we would be done by induction. Indeed, this would imply that there is a ridge path between the facets $H \cap F$ and $H \cap F'$ in the $(d - 1)$ -dimensional tropical variety (2.1). Since $d \geq 3$, the $(d - 1)$ - and $(d - 2)$ -dimensional faces of $H \cap \mathcal{T}(I)$ arise uniquely from the intersections of H with d - and $(d - 1)$ -dimensional faces of $\mathcal{T}(I)$. Hence this path is also a ridge path considered as a path in $\mathcal{T}(I)$.

Let $V(J)$ denote the subvariety of the algebraic torus $(\mathbb{C}^*)^n$ defined by an ideal $J \subset \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. The tropical variety in (2.1) depends only on the subvariety of $(\mathbb{C}^*)^n$ defined by our ideal $I + \langle f_u \rangle$. This subvariety is

$$V(I + \langle f_u \rangle) = V(I) \cap V(f_u) = V(I) \cap u^{-1} \cdot V(f_{\mathbf{1}}). \quad (2.2)$$

Here $\mathbf{1}$ denotes the identity element of $(\mathbb{C}^*)^n$. For generic choices of the group element $u \in (\mathbb{C}^*)^n$, the intersection (2.2) is an irreducible subvariety of dimension $d - 1$ in $(\mathbb{C}^*)^n$. This follows from Kleiman's version of Bertini's Theorem [29, Theorem III.10.8], applied to the algebraic group $(\mathbb{C}^*)^n$. Hence (2.1) is indeed an irreducible tropical variety of dimension $d - 1$, defined by the prime ideal $I + \langle f_u \rangle$. This completes the proof by induction. \square

Proof of Lemma 2.18: Again, we replace $I \subset \mathbb{C}[\mathbf{x}]$ by the ideal it generates in $\mathbb{C}[\mathbf{x}^{\pm 1}]$ and continue to denote this ideal by I .

Let F be the cone of $\mathcal{T}(I)$ which contains ω in its relative interior and G the cone of $\mathcal{T}(J)$ which contains ω in its relative interior. Our hypothesis is that F and G meet transversally at ω , that is,

$$\mathbb{R}F + \mathbb{R}G = \mathbb{R}^n.$$

We claim that the ideal $\text{in}_\omega(I)$ is homogeneous with respect to any weight vector $v \in \mathbb{R}F$ or, equivalently (see Proposition 2.9), that $\text{in}_{\prec_v}(\text{in}_\omega(I)) = \text{in}_\omega(I)$. According to Proposition 1.13 in [47], for ϵ a sufficiently small positive number, $\text{in}_{\prec_{\omega+\epsilon v}}(I) = \text{in}_{\prec_v}(\text{in}_\omega(I))$. The vector $\omega + \epsilon v$ is in the relative interior of F so $\text{in}_{\prec_{\omega+\epsilon v}}(I) = \text{in}_\omega(I)$. By the same argument, the ideal $\text{in}_\omega(J)$ is homogeneous with respect to any weight vector in $\mathbb{R}G$.

After a multiplicative change of variables in $\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ we may assume that $\omega = e_1$, $\mathbb{R}\{e_1, e_2, \dots, e_s\} \subseteq \mathbb{R}F$ and $\mathbb{R}\{e_1, e_{s+1}, \dots, e_n\} \subseteq \mathbb{R}G$. We change the notation for the variables as follows:

$$t = x_1, y = (y_2, \dots, y_s) = (x_2, \dots, x_s), z = (z_{s+1}, \dots, z_n) = (x_{s+1}, \dots, x_n).$$

The homogeneity properties of the two initial ideals ensure that we can pick generators $f_1(z), \dots, f_a(z)$ for $\text{in}_\omega(I)$ and generators $g_1(y), \dots, g_b(y)$ for $\text{in}_\omega(J)$. Since $\text{in}_\omega(I)$ is not the unit ideal, the Laurent polynomials $f_i(z)$ have a common zero $Z = (Z_{s+1}, \dots, Z_n) \in (\mathbb{C}^*)^{n-s}$, and likewise the Laurent polynomials $g_j(y)$ have a common zero $Y = (Y_2, \dots, Y_s) \in (\mathbb{C}^*)^{s-1}$.

Next we consider the following general chain of inclusions of ideals:

$$\text{in}_\omega(I) \cdot \text{in}_\omega(J) \subseteq \text{in}_\omega(I \cdot J) \subseteq \text{in}_\omega(I \cap J) \subseteq \text{in}_\omega(I) \cap \text{in}_\omega(J). \quad (2.3)$$

The product of two ideals which are generated by (Laurent) polynomials in disjoint sets of variables equals the intersection of the two ideals. Since the set of y -variables is disjoint from the set of z -variables, it follows that the first ideal in (2.3) equals the last ideal in (2.3). In particular, we conclude that

$$\text{in}_\omega(I \cap J) = \text{in}_\omega(I) \cap \text{in}_\omega(J). \quad (2.4)$$

We next claim that

$$\text{in}_\omega(I + J) = \text{in}_\omega(I) + \text{in}_\omega(J). \quad (2.5)$$

The left hand side is an ideal which contains both $\text{in}_\omega(I)$ and $\text{in}_\omega(J)$, so it contains their sum. We must prove that the right hand side contains the left hand side.

Consider any element $f + g \in I + J$ where $f \in I$ and $g \in J$. Let $f = f_0(y, z) + t \cdot f_1(t, y, z)$ and $g = g_0(y, z) + t \cdot g_1(t, y, z)$. We also have the following representation for some integer $a \geq 0$ and non-zero polynomial h_0 :

$$f + g = t^a \cdot h_0(y, z) + t^{a+1} \cdot h_1(t, y, z).$$

If $a = 0$ then we conclude

$$\text{in}_\omega(f + g) = h_0(y, z) = f_0(y, z) + g_0(y, z) \in \text{in}_\omega(I) + \text{in}_\omega(J).$$

If $a \geq 1$ then $f_0 = -g_0$ lies in $\text{in}_\omega(I) \cap \text{in}_\omega(J)$. In view of (2.4), there exists $p \in I \cap J$ with $f_0 = -g_0 = \text{in}_\omega(p)$. Then $f + g = (f - p) + (g + p)$ and replacing f by $(f - p)/t$ and g by $(g + p)/t$ puts us in the same situation as before, but with a reduced by 1. By induction on a , we conclude that $\text{in}_\omega(f + g)$ is in $\text{in}_\omega(I) + \text{in}_\omega(J)$, and the claim (2.5) follows.

For any constant $T \in \mathbb{C}^*$, the vector $(T, Y_2, \dots, Y_s, Z_{s+1}, \dots, Z_n)$ is a common zero in $(\mathbb{C}^*)^n$ of the ideal (2.5). We conclude that $\text{in}_\omega(I + J)$ is not the unit ideal, so it contains no monomial, and hence $\omega \in \mathcal{T}(I + J)$. \square

2.2.4 Algorithms

In this section we describe algorithms for solving the computational problems raised in Section 2.2.2. The emphasis is on algorithms leading to a solution of Problem 2.8 for prime ideals, taking advantage of Theorem 2.17. Recall that we only need to consider the case of homogeneous ideals in $\mathbb{C}[\mathbf{x}]$.

In order to state our algorithms we must first explain how polyhedral cones and polyhedral fans are represented. A polyhedral cone is represented by a canonical minimal set of inequalities and equations. Given arbitrary defining linear inequalities and equations, the task of bringing these to a canonical form involves linear programming. Representing a polyhedral fan requires a little thought. We are rarely interested in all faces of all cones.

Definition 2.20. A set S of polyhedral cones in \mathbb{R}^n is said to *represent* a fan \mathcal{F} in \mathbb{R}^n if the set of all faces of cones in S is exactly \mathcal{F} .

A representation may contain non-maximal cones, but each cone is represented minimally by its canonical form. A Gröbner cone $C_\omega(I)$ is represented by the pair $(\mathcal{G}_{\prec_\omega}(\text{in}_\omega(I)), \mathcal{G}_{\prec_\omega}(I))$ of marked reduced Gröbner bases, where \prec is some globally fixed term order. In a *marked* Gröbner basis the initial terms are distinguished. The advantage of using marked Gröbner bases is that the weight vector ω need not be stored: we can deduce defining inequalities for its cone from the marked reduced Gröbner bases themselves; see Example 2.32. This is done as follows; see [47, proof of Proposition 2.3]:

Lemma 2.21. *Let $I \subset \mathbb{C}[\mathbf{x}]$ be a homogeneous ideal, \prec a term order and $\omega \in \mathbb{R}^n$ a vector. For any other vector $\omega' \in \mathbb{R}^n$:*

$$\omega' \in C_\omega(I) \iff \forall f \in \mathcal{G}_{\prec_\omega}(I) : \text{in}_{\prec_\omega}(\text{in}_{\omega'}(f)) = \text{in}_{\prec_\omega}(f).$$

Our first two algorithms perform polyhedral computations and between them solve Problem 2.6. By the *support* of a fan we mean the union of its cones. Recall that for a polynomial f , the tropical hypersurface $\mathcal{T}(f)$ is the union of the normal cones of the edges of the Newton polytope $\text{New}(f)$. The first algorithm computes these cones.

Algorithm 2.22. *Tropical Hypersurface*

Input: $f \in \mathbb{C}[\mathbf{x}]$.

Output: A representation S of a polyhedral fan whose support is $\mathcal{T}(f)$.

{

$S := \emptyset;$

For every vertex $v \in \text{New}(f)$

{

Compute the normal cone C of v in $\text{New}(f)$;

$S := S \cup \{\text{the facets of } C\};$

}

}

}
}

Let \mathcal{F}_1 and \mathcal{F}_2 be polyhedral fans in \mathbb{R}^n . Their *common refinement* is the fan

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

To compute a common refinement we simply run through all pairs of cones in the fan representations and bring the set of their intersections to canonical form. The canonical form makes it easy to remove duplicates.

Algorithm 2.23. *Common Refinement*

Input: Representations S_1 and S_2 for polyhedral fans \mathcal{F}_1 and \mathcal{F}_2 .

Output: A representation S for the common refinement $\mathcal{F}_1 \wedge \mathcal{F}_2$.

{
 $S := \emptyset$;
 For every pair $(C_1, C_2) \in S_1 \times S_2$
 $S := S \cup \{C_1 \cap C_2\}$;
}

If refinements of more than two fans are needed, Algorithm 2.23 can be applied successively. Note that the intersection of the support of two fans is the support of their common refinement. Hence Algorithm 2.23 can be used for computing intersections of tropical hypersurfaces. This solves Problem 2.6, but the output may be a highly redundant representation.

Recall (from the proof of Theorem 2.14) that a witness $f \in I$ is a polynomial which certifies $\mathcal{T}(f) \cap \text{rel int}(C_\omega(I)) = \emptyset$. Computing witnesses is essential for solving Problem 2.8 and Problem 2.11. The first step of constructing a witness is to check if the ideal $\text{in}_\omega(I)$ contains monomials, and, if so, to compute one such monomial. The check for monomial containment can be implemented by saturating the ideal with respect to the product of the variables (cf. [47, Lemma 12.1]). Knowing that the ideal

contains a monomial, a simple way to find one is to repeatedly reduce powers of the product of the variables by applying the division algorithm until the remainder is 0.

Algorithm 2.24. *Monomial in Ideal*

Input: A set of generators for an ideal $I \subset \mathbb{C}[\mathbf{x}]$.

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

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

Remark 2.25. To pick the smallest monomial in I with respect to a term order, we first compute the largest monomial ideal contained in I using [42, Algorithm 4.2.2] and then pick the smallest monomial generator of this ideal.

Constructing a witness from a monomial was already explained in the proof of Theorem 2.14. Here, we only state the input and output of this algorithm.

Algorithm 2.26. *Witness*

Input: A set of generators for an ideal $I \subset \mathbb{C}[\mathbf{x}]$ and a vector $\omega \in \mathbb{R}^n$ with $\text{in}_\omega(I)$ containing a monomial.

Output: A polynomial $f \in I$ such that the tropical hypersurface $\mathcal{T}(f)$ and the relative interior of $C_\omega(I)$ have empty intersection.

Combining Algorithm 2.24 and Algorithm 2.26 with known methods (e.g. [47, Algorithm 3.6]) for computing Gröbner fans, we can now compute the tropical variety $\mathcal{T}(I)$ and a tropical basis of I . This solves Problem 2.8 and Problem 2.11. However, this approach is not at all practical, as shown in Section 2.2.6

We will now develop a practical algorithm for computing $\mathcal{T}(I)$ when I is prime. An ideal $I \subset \mathbb{C}[\mathbf{x}]$ is said to define a *tropical curve* if $\dim(I) = 1 + \text{homog}(I)$. Our problems are easier in this case because a tropical curve consists of only finitely many rays and the origin modulo the homogeneity space.

Algorithm 2.27. *Tropical Basis of a Curve*

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

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

```

{
  Compute a representation  $S$  of  $\bigwedge_{g \in \mathcal{G}} \mathcal{T}(g)$ ;
  For every  $C \in S$ 
  {
    Let  $\omega$  be a generic relative interior point in  $C$ ;
    If  $(\text{in}_\omega(I))$  contains a monomial
      then add a witness to  $\mathcal{G}$  and restart the algorithm;
  }
   $\mathcal{G}' := \mathcal{G}$ ;
}

```

Proof of correctness. The algorithm terminates because I has only finitely many initial ideals and at least one is excluded in every iteration. If a vector ω passes the monomial test (which verifies $\omega \in \mathcal{T}(I)$) then C has dimension 0 or 1 modulo the homogeneity space since we are considering a tropical curve and ω is generic in C . Any other relative interior point of C would also pass the monomial test. (This property fails in higher dimensions, when $\mathcal{T}(I)$ is no longer a tropical curve). Hence, when we terminate only points in the tropical variety are covered by S . Thus \mathcal{G}' is a tropical basis. \square

In the curve case, combining Algorithms 2.22 and 2.23 with Algorithm 2.27 we get a reasonable method for solving Problem 2.8. This method is used as a subroutine in

Algorithm 2.29 below. In the remainder of this section we concentrate on providing a better algorithm for Problem 2.8 for problem. The idea is to use the connectivity shown in Section 2.2.3 in order to traverse the tropical variety.

The next algorithm is one step in the Gröbner walk [12] and is an important subroutine here. We only specify the input and output.

Algorithm 2.28. *Lift*

Input: Marked reduced Gröbner bases $\mathcal{G}_{\prec'}(I)$ and $\mathcal{G}_{\prec_\omega}(\text{in}_\omega(I))$ where $\omega \in C_{\prec'}(I)$ is an unspecified vector and \prec and \prec' are unspecified term orders.

Output: The marked reduced Gröbner basis $\mathcal{G}_{\prec_\omega}(I)$.

We now suppose that I is a monomial-free prime ideal with $d = \dim(I)$, and \prec is a globally fixed term order. We first describe the local computations needed for a traversal of the d -dimensional Gröbner cones contained in $\mathcal{T}(I)$.

Algorithm 2.29. *Neighbors*

Input: A pair $(\mathcal{G}_{\prec_\omega}(\text{in}_\omega(I)), \mathcal{G}_{\prec_\omega}(I))$ such that $\text{in}_\omega(I)$ is monomial-free and $C_\omega(I)$ has dimension d .

Output: The collection N of pairs of the form $(\mathcal{G}_{\prec_{\omega'}}(\text{in}_{\omega'}(I)), \mathcal{G}_{\prec_{\omega'}}(I))$ where one ω' is taken from the relative interior of each d -dimensional Gröbner cone contained in $\mathcal{T}(I)$ that has a facet in common with $C_\omega(I)$.

{

$N := \emptyset;$

Compute the set \mathcal{F} of facets of $C_\omega(I)$;

For each facet $F \in \mathcal{F}$

{

Compute the initial ideal $J := \text{in}_{\prec_{\mathbf{u}}}(I)$

where \mathbf{u} is a relative interior point in F ;

Use Algorithm 2.27 and Algorithm 2.23 to produce a relative

interior point \mathbf{v} of each ray in the curve $\mathcal{T}(J)$;

For each such \mathbf{v}

{

Compute $(\mathcal{G}_{\prec_{\mathbf{v}}}(\text{in}_{\prec_{\mathbf{v}}}(J)), \mathcal{G}_{\prec_{\mathbf{v}}}(J)) = (\mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(\text{in}_{\prec_{\mathbf{v}}}(J)), \mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(J))$;

Apply Algorithm 2.28 to $\mathcal{G}_{\prec_{\omega}}(I)$ and $\mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(J)$ to get $\mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(I)$;

$N := N \cup \{(\mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(\text{in}_{\prec_{\mathbf{v}}}(J)), \mathcal{G}_{\prec_{\mathbf{v}\mathbf{u}}}(I))\}$;

}

}

}

Proof of correctness. Facets and relative interior points are computed using linear programming. Figure 2.4 illustrates the choices of vectors in the algorithm. The initial ideal $\text{in}_{\prec_{\mathbf{u}}}(I)$ is homogeneous with respect to the span of F . Hence its homogeneity space has dimension $d - 1$. The Krull dimension of $\mathbb{C}[\mathbf{x}]/\text{in}_{\prec_{\mathbf{u}}}(I)$ is d , so $\text{in}_{\prec_{\mathbf{u}}}(I)$ defines a curve and $\mathcal{T}(\text{in}_{\prec_{\mathbf{u}}}(I))$ can be computed using Algorithm 2.27. The identity $\text{in}_{\prec_{\mathbf{v}}}(\text{in}_{\prec_{\mathbf{u}}}(I)) = \text{in}_{\prec_{\mathbf{u}+\varepsilon\mathbf{v}}}(I)$ for small $\varepsilon > 0$ [47, Proposition 1.13] implies that we run through all the desired $\text{in}_{\omega'}(I)$ where $\omega' = \mathbf{u} + \varepsilon\mathbf{v}$ for small $\varepsilon > 0$. The lifting step can be carried out since $\mathbf{u} \in C_{\prec_{\omega}}(I)$. \square

Algorithm 2.30. *Traversal of an Irreducible Tropical Variety*

Input: A pair $(\mathcal{G}_{\prec_{\omega}}(\text{in}_{\omega}(I)), \mathcal{G}_{\prec_{\omega}}(I))$ such that $\text{in}_{\omega}(I)$ is monomial-free and $C_{\omega}(I)$ has dimension d .

Output: The collection T of pairs of the form $(\mathcal{G}_{\prec_{\omega'}}(\text{in}_{\omega'}(I)), \mathcal{G}_{\prec_{\omega'}}(I))$ where one ω' is taken from the relative interior of each d -dimensional Gröbner cone contained in $\mathcal{T}(I)$. The union of all the $C_{\omega'}(I)$ is $\mathcal{T}(I)$.

{

$T := \{(\mathcal{G}_{\prec_{\omega}}(\text{in}_{\omega}(I)), \mathcal{G}_{\prec_{\omega}}(I))\}$;

$Old := \emptyset$;

While $(T \neq Old)$




Figure 2.4: A projective drawing of the situation in Algorithm 2.29, with $\mathcal{T}(I)$ on the left and $\mathcal{T}(\text{in}_{\prec_u}(I))$ on the right.

```

{
  Old := T;
  T := T ∪ Neighbors(T);
}
}

```

Proof of correctness. By $\text{Neighbors}(T)$ we mean the union of all the output of Algorithm 2.29 applied to all pairs in T . The algorithm computes the connected component of the starting pair. Since I is a prime ideal, Theorem 2.17 implies that the union of all the computed $C_{\omega'}(I)$ is $\mathcal{T}(I)$. \square

To use Algorithm 2.30 we must know a starting d -dimensional Gröbner cone contained in the tropical variety. One inefficient method for finding one would be to compute the entire Gröbner fan. Instead we currently use heuristics, based on the following probabilistic recursive algorithm:

Algorithm 2.31. *Starting Cone*

Input: A marked reduced Gröbner basis \mathcal{G} for an ideal I whose tropical variety is

pure of dimension $d = \dim(I)$. A term order \prec for tie-breaking.

Output: Two marked reduced Gröbner bases:

- One for an initial ideal $\text{in}_{\omega'}(I)$ without monomials, where the homogeneity space of $\text{in}_{\omega'}(I)$ has dimension d . The term order is $\prec_{\omega'}$.
- A marked reduced Gröbner basis for I with respect to $\prec_{\omega'}$.

```

{
  If ( $\dim(I) = \text{homog}(I)$ )
    Return ( $\mathcal{G}_{\prec}(I), \mathcal{G}_{\prec}(I)$ );
  If not
  {
    Repeat
    {
      Compute a random reduced Gröbner basis of  $I$ ;
      Compute a random extreme ray  $\omega$  of its Gröbner cone;
    }
    Until ( $\text{in}_{\omega}(I)$  is monomial free);
    Compute  $\mathcal{G}_{\prec_{\omega}}(I)$ ;
    ( $\mathcal{G}_{Init}, \mathcal{G}_{Full}$ ) := Starting Cone( $\mathcal{G}_{\prec_{\omega}}(\text{in}_{\omega}(I))$ );
    Apply Algorithm 2.28 to  $\mathcal{G}_{\prec_{\omega}}(I)$  and  $\mathcal{G}_{Full}$ 
      to get a marked reduced Gröbner basis  $\mathcal{G}'$  for  $I$ ;
    Return ( $\mathcal{G}_{Init}, \mathcal{G}'$ );
  }
}

```

2.2.5 Software and Examples

The algorithms of Section 2.2.4 are implemented in the software package `Gfan` [34]. `Gfan` uses the library `cddlib` [23] for polyhedral computations such as finding facets and extreme rays of cones and bringing cones to canonical form. The library `gmp` [21] supplies both polyhedral computations and efficient exact arithmetic in $\mathbb{Q}[\mathbf{x}]$. (The input is required to have only rational coefficients.) In this section we illustrate the use of `Gfan` in computing various tropical varieties.

Example 2.32. We consider the prime ideal $I \subset \mathbb{C}[a, b, c, d, e, f, g]$ which is generated by the 3×3 minors of the generic *Hankel matrix* of size 4×4 :

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

Its tropical variety is a 4-dimensional fan in \mathbb{R}^7 with 2-dimensional homogeneity space. Its combinatorics is given by the graph in Figure 2.5. To compute $\mathcal{T}(I)$ in `Gfan`, we write the ideal generators on a file `hankel.in`:

```
% more hankel.in
{-c^3+2*b*c*d-a*d^2-b^2*e+a*c*e, -c^2*d+b*d^2+b*c*e-a*d*e-b^2*f+a*c*f,
-c*d^2+c^2*e+b*d*e-a*e^2-b*c*f+a*d*f, -d^3+2*c*d*e-b*e^2-c^2*f+b*d*f,
-c^2*d+b*d^2+b*c*e-a*d*e-b^2*f+a*c*f, -c*d^2+2*b*d*e-a*e^2-b^2*g+a*c*g,
-d^3+c*d*e+b*d*f-a*e*f-b*c*g+a*d*g, -d^2*e+c*e^2+c*d*f-b*e*f-c^2*g+b*d*g,
-c*d^2+c^2*e+b*d*e-a*e^2-b*c*f+a*d*f, -d^3+c*d*e+b*d*f-a*e*f-b*c*g+a*d*g,
-d^2*e+2*c*d*f-a*f^2-c^2*g+a*e*g, -d*e^2+d^2*f+c*e*f-b*f^2-c*d*g+b*e*g,
-d^3+2*c*d*e-b*e^2-c^2*f+b*d*f, -d^2*e+c*e^2+c*d*f-b*e*f-c^2*g+b*d*g,
-d*e^2+d^2*f+c*e*f-b*f^2-c*d*g+b*e*g, -e^3+2*d*e*f-c*f^2-d^2*g+c*e*g}
```

We then run the command

```
gfan_tropicalstartingcone < hankel.in > hankel.start
```

which applies Algorithm 2.31 to produce a pair of marked Gröbner bases. This output represents a maximal cone in $\mathcal{T}(I)$, as explained prior to Lemma 2.21.

```
% more hankel.start
{
c*f^2-c*e*g,
b*f^2-b*e*g,
b*e*f+c^2*g,
b*e^2+c^2*f,
b^2*g-a*c*g,
b^2*f-a*c*f,
b^2*e-a*c*e,
a*f^2-a*e*g,
a*e*f+b*c*g,
a*e^2+b*c*f}
{
c*f^2+e^3-2d*e*f+d^2*g-c*e*g,
b*f^2+d*e^2-d^2*f-c*e*f+c*d*g-b*e*g,
b*e*f+d^2*e-c*e^2-c*d*f+c^2*g-b*d*g,
b*e^2+d^3-2c*d*e+c^2*f-b*d*f,
b^2*g+c^2*e-b*d*e-b*c*f+a*d*f-a*c*g,
b^2*f+c^2*d-b*d^2-b*c*e+a*d*e-a*c*f,
b^2*e+c^3-2b*c*d+a*d^2-a*c*e,
a*f^2+d^2*e-2c*d*f+c^2*g-a*e*g,
a*e*f+d^3-c*d*e-b*d*f+b*c*g-a*d*g,
a*e^2+c*d^2-c^2*e-b*d*e+b*c*f-a*d*f}
```

Using Lemma 2.21 we can easily read off the canonical equations and equalities for

the corresponding Gröbner cone $C_\omega(I)$. For example, the polynomials $cf^2 - ceg$ and $cf^2 + e^3 - 2def + d^2g - ceg$ represent the equation

$$\omega_c + 2\omega_f = \omega_c + \omega_e + \omega_g$$

and the inequalities

$$\omega_c + 2\omega_f \leq \min\{3\omega_e, \omega_d + \omega_e + \omega_f, 2\omega_d + \omega_g, \omega_c + \omega_e + \omega_g\}.$$

At this point, we could run Algorithm 2.30 using the following command:

```
gfan_tropicaltraverse < hankel.start > hankel.out
```

However, we can save computing time and get a better idea of the structure of $\mathcal{T}(I)$ by instructing `Gfan` to take advantage of symmetries of I as it produces cones. The only symmetries that can be exploited in `Gfan` are those that simply permute the variables. The output will show which cones of $\mathcal{T}(I)$ lie in the same orbit under the action of the symmetry group we provide.

Our ideal I is invariant under reflection of the 4×4 -matrix across the anti-diagonal. This reverses the order of the variables a, b, \dots, g . To specify this permutation, we add the following line to the bottom of the file `hankel.start`:

```
{(6,5,4,3,2,1,0)}
```

We can add more symmetries by listing them one after another, separated by commas, inside the curly braces. `Gfan` will compute and use the group generated by the set of permutations we provide, and it will return an error if we input any permutation which is not a symmetry of the ideal.

After adding the symmetries, we run the command

```
gfan_tropicaltraverse --symmetry < hankel.start > hankel.out
```

to compute the tropical variety. We show the output with some annotations:

```
% more hankel.out
```

```
Ambient dimension: 7
Dimension of homogeneity space: 2
Dimension of tropical variety: 4
Simplicial: true
Order of input symmetry group: 2
F-vector: (16,28)
```

A short list of basic data: the dimensions of the ambient space, of $\mathcal{T}(I)$, and of its homogeneity space, and also the face numbers (f -vector) of $\mathcal{T}(I)$ and the order of symmetry group specified in the input.

Modulo the homogeneity space:

```
{(6,5,4,3,2,-1,0),
 (5,4,3,2,1,0,-1)}
```

A basis for the homogeneity space. The following rays are considered in the quotient of \mathbb{R}^7 modulo this 2-dimensional subspace.

Rays:

```
{0: (-1,0,0,0,0,0,0),
 1: (-5,-4,-3,-2,-1,0,0),
 2: (1,0,0,0,0,0,0),
 3: (5,4,3,2,1,0,0),
 4: (2,1,0,0,0,0,0),
 5: (4,3,2,1,0,0,0),
 6: (0,-1,0,0,0,0,0),
 7: (6,5,4,3,2,0,0),
 8: (3,2,1,0,0,0,0),
```

9: $(0, 0, -1, 0, 0, 0, 0)$,
 10: $(0, 0, 0, 0, -1, 0, 0)$,
 11: $(0, 0, 0, -1, 0, 0, 0)$,
 12: $(-6, -4, -3, -3, -1, 0, 0)$,
 13: $(-3, -2, -2, -1, -1, 0, 0)$,
 14: $(3, 2, 2, 1, 1, 0, 0)$,
 15: $(3, 2, 2, 0, 1, 0, 0)$

The direction vectors of the tropical rays. Because the homogeneity space is positive-dimensional, the directions are not uniquely specified. For instance the vectors $(-5, -4, -3, -2, -1, 0, 0)$ and $(0, 0, 0, 0, 0, 0, -1)$ represent the same ray. Notice that Gfan uses negated weight vectors.

Rays incident to each

dimension 2 cone:

$\{2, 6\}$, $\{3, 7\}$,
 $\{2, 4\}$, $\{3, 5\}$,
 $\{4, 9\}$, $\{5, 10\}$,
 $\{4, 8\}$, $\{5, 8\}$,
 $\{8, 11\}$,
 $\{0, 12\}$, $\{1, 12\}$,
 $\{0, 1\}$,
 $\{1, 6\}$, $\{0, 7\}$,
 $\{1, 9\}$, $\{0, 10\}$,
 $\{0, 13\}$, $\{1, 13\}$,
 $\{6, 14\}$, $\{7, 14\}$,
 $\{9, 13\}$, $\{10, 13\}$,
 $\{6, 10\}$, $\{7, 9\}$,
 $\{6, 7\}$,




Figure 2.5: The tropical variety of the ideal generated by the 3×3 minors of the generic 4×4 Hankel matrix.

$\{11, 12\}$,

$\{11, 15\}$,

$\{14, 15\}$

The cones in $\mathcal{T}(I)$ are listed from highest to lowest dimension. Each cone is named by the set of rays on it. There are 28 2-dimensional cones, broken down into 11 orbits of size 2 and 6 orbits of size 1.

The further output, which is not displayed here, shows that the 16 rays break down into 5 orbits of size 2 and 6 orbits of size 1.

Using the same procedure, we now compute several more examples.

Example 2.33. Let I be the ideal generated by the 3×3 minors of the generic 5×5 Hankel matrix. We again use the symmetry group $\mathbb{Z}/2$. The tropical variety is a graph with vertex degrees ranging from 2 to 7.

Ambient dimension: 9

Dimension of homogeneity space: 2

Dimension of tropical variety: 4

Simplicial: true

F-vector: (28,53)

Example 2.34. Let I be the ideal generated by the 3×3 minors of a generic 3×5 matrix. We use the symmetry group $S_5 \times S_3$, where S_5 acts by permuting the columns and S_3 by permuting the rows.

Ambient dimension: 15

Dimension of homogeneity space: 7

Dimension of tropical variety: 12

Simplicial: true

F-vector: (45,315,930,1260,630)

Example 2.35. Let I be the ideal generated by the 3×3 minors of a generic 4×4 symmetric matrix. We use the symmetry group S_4 which acts by simultaneously permuting the rows and the columns.

Ambient dimension: 10

Dimension of homogeneity space: 4

Dimension of tropical variety: 7

Simplicial: true

F-vector: (20,75,75)

If we take the 3×3 minors of a generic 5×5 symmetric matrix then we get

Ambient dimension: 15

Dimension of homogeneity space: 5

Dimension of tropical variety: 9

Simplicial: true

F-vector: (75, 495, 1155, 855)

Example 2.36. Let I be the prime ideal of a pair of commuting 2×2 matrices. That is, $I \subset \mathbb{C}[a, b, \dots, h]$ is defined by the matrix equation

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix} \begin{pmatrix} e & g \\ f & h \end{pmatrix} - \begin{pmatrix} e & g \\ f & h \end{pmatrix} \begin{pmatrix} a & c \\ b & d \end{pmatrix} = 0.$$

The tropical variety is the graph K_4 , which Gfan reports as follows:

Ambient dimension: 8

Dimension of homogeneity space: 4

Dimension of tropical variety: 6

Simplicial: true

F-vector: (4,6)

If I is the ideal of 3×3 commuting symmetric matrices then we get:

Ambient dimension: 12

Dimension of homogeneity space: 2

Dimension of tropical variety: 9

Simplicial: false

F-vector: (66,705,3246,7932,10888,8184,2745)

2.2.6 Tropical variety versus Gröbner fan

In computing the tropical variety of a prime ideal I , we took advantage of the fact that, since I is homogeneous, the set $\mathcal{T}(I)$ has naturally the structure of a polyhedral fan, namely, $\mathcal{T}(I)$ is the collection of all cones in the Gröbner fan of I whose corresponding initial ideal is monomial-free. A naive algorithm would be to compute the Gröbner fan of I and then retain only those d -dimensional cones who survive the monomial test (Algorithm 2.24). The software Gfan also computes the full Gröbner fan of I , and so we tested this naive algorithm. We found it to be inefficient. The reason is that typically the vast majority of d -dimensional cones in the Gröbner fan of I are not in the tropical variety $\mathcal{T}(I)$.

Example 2.37. Consider the ideal I in Example 2.32 which is generated by the 3×3 -minors of a generic 4×4 Hankel matrix. Let J be its initial ideal with respect to the first vector ω in the list of rays. The initial ideal J defines a tropical curve consisting of 5 rays and the origin. The curve is a subfan of the much more complicated Gröbner fan of J . The Gröbner fan is full-dimensional in \mathbb{R}^7 with $C_0(J)$ being three-dimensional. Its f-vector equals $(1, 7167, 32656, 45072, 19583)$. Of the 7167 rays only 5 are in the tropical variety. The Gröbner fan of J is the link of the Gröbner fan of I at ω . We were unable to compute the full Gröbner fan of I .

Example 2.38. Toric Ideals. Let $I = \langle \mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}} : A\mathbf{u} = A\mathbf{v} \rangle$ be the toric ideal of a matrix $A \in \mathbb{Z}^{d \times n}$ of rank d . The ideal I is a prime of dimension d . The tropical variety $\mathcal{T}(I)$ coincides with the homogeneity space $C_0(I)$ which is just the row space of A . Hence $\mathcal{T}(I)$ modulo $C_0(I)$ is a single point. Yet, the Gröbner fan of I can be very complicated, as it encodes the sensitivity information for an infinite family of integer programs [47, Chapter 7].

We next exhibit an infinite family of ideals such that the number of rays in $\mathcal{T}(I)$ is constant while the number of rays in the Gröbner fan of I grows linearly.

Theorem 2.39. Fix $n=3, d=1$ and for any positive integer p consider the ideal



$$I_p = \langle \underline{a} - (c+1)^{p+2}, \underline{b} - (c-1)^p \rangle.$$

Then the tropical variety $\mathcal{T}(I_p)$ consists of 4 rays but

the Gröbner fan of I_p has $\geq \frac{1}{4}(p+1)$ rays.

The ingredients of the proof are Gröbner basis calculations to show that the Gröbner fan has many cells, an Euler characteristic argument to show that many cells imply many rays, and a general fact about constant-coefficient parametrized tropical curves that shows that $\mathcal{T}(I)$ has only four rays.

Figures 2.6 and 2.7 show the intersection of the Gröbner fan of I with the positive orthant (the *restricted Gröbner fan*) for some small values of p .

Figure 2.6: $p = 1, 3$ Figure 2.7: $p = 5, 7$

Lemma 2.40. *The set $\mathcal{G}_0 := \{\underline{c} - a^{p+2} - a^{p+1}, \underline{b} - a^p\}$ is the reduced Gröbner basis of I_p with respect to the weight vector $\omega_0 = (1, 1 + p, p + 3)$ which is in the same Gröbner cone as $(0, 1, 1)$.*

Proof: Clearly $\mathcal{G}_0 \subset I_p$ and the underlined terms are the leading terms with respect to ω_0 . Since the leading terms are relatively prime and do not divide any of the trailing terms, the two polynomials form a reduced Gröbner basis of the ideal they generate. Thus it only remains to show that \mathcal{G}_0 generates I_p . If $f \in I_p = \ker(a \mapsto t, b \mapsto t^p, c \mapsto t^{p+1} + t^{p+2})$, then the normal form of f with respect to \mathcal{G}_0 is a polynomial in a . This polynomial has to be zero since the kernel described above contains no

univariate polynomial in a . □

Remark 2.41. The pictures in Figures 2.6 and 2.7 are drawn in the triangles with vertices $(1, 0, 0)$ (the right bottom vertex), $(0, 1, 0)$ (the left bottom vertex) and $(0, 0, 1)$ the top vertex. The Gröbner cone of \mathcal{G}_0 is the unique cell on the left of each picture with the segment $[(0, 1, 0), (0, 0, 1)]$ as a face. The description of \mathcal{G}_0 implies that this Gröbner cone is given as

$$K_0 = \{(x, y, z) \geq \mathbf{0} : y \geq px, z \geq (p+2)x\}.$$

In particular the extreme rays of K_0 are generated by $(0, 1, 0)$, $(0, 0, 1)$ and $(1, p, p+2) =: \mathbf{v}$.

Lemma 2.42. *The restricted Gröbner fan of I_p has at least $(p+1)/2$ full-dimensional Gröbner cones.*

Proof: We will describe $(p+1)/2$ Gröbner cones $K_1, \dots, K_{(p+1)/2}$ arranged counter-clockwise around $\mathbf{v} = (1, p, p+2)$, each one adjacent to the previous, with K_1 adjacent to K_0 . Let \mathcal{G}_i be the reduced Gröbner basis of I_p whose Gröbner cone is K_i . Let the weight vector ω_i induce \mathcal{G}_i . We will show that b^i is a minimal generator of the initial ideal $\text{in}_{\prec_{\omega_i}}(I_p)$.

The facet inequalities of K_0 are $z \geq (p+2)x$, $y \geq px$ and $x \geq 0$. Flipping across $z = (p+2)x$, and choosing ω_1 very close to this wall but satisfying $z < (p+2)x$, we may write $f_1 = \underline{a^{p+2}} + a^{p+1} - c$ and $f_2 = \underline{b} - a^p$ where the underlined terms are the leading terms with respect to ω_1 . Then $\mathcal{G}_1 := \{\underline{a^{p+2}} + a^{p+1} - c, \underline{b} - a^p\}$ is the reduced Gröbner basis of I_p with respect to ω_1 . Its Gröbner cone is

$$K_1 = \{(x, y, z) \geq \mathbf{0} : (p+2)x \geq z, y \geq px\}.$$

Next we flip across the wall $y = px$ and choose ω_2 very close to this wall but satisfying $y < px$. Reordering the elements of \mathcal{G}_1 with respect to ω_2 and reducing wherever

possible, we get $\mathcal{G}' = \{f_1 := \underline{a^p} - b, f_2 := \underline{a^2b} + ab - c\}$. To see that the leading term of $a^2b + ab - c$ is a^2b , note that $\omega_2 = (x, y, z)$ satisfies the inequalities $x, y, z > 0$, $(p+2)x > z$ and $y < px$ with y very close to px . Thus $(p+2)x \sim 2x + y > x + y, z$. Clearly \mathcal{G}' generates I_p since it is obtained from \mathcal{G}_1 by reordering terms and reduction. The S-polynomial of f_1 and f_2 is $b^2 + a^{p-1}b - a^{p-2}c$ which can be reduced by f_2 . The normal form obtained is

$$f_3 := b^2 - ab + c(1 - a + a^2 - a^3 + \dots - a^{p-2})$$

and its leading term is b^2 . We claim that

$$\mathcal{G}_2 := \{f_1 := \underline{a^p} - b, f_2 := \underline{a^2b} + ab - c, f_3 := \underline{b^2} - ab + c(1 - a + a^2 - a^3 + \dots - a^{p-2})\}$$

is the reduced Gröbner basis of I_p with respect to ω_2 . To prove this we need to show that the S-polynomial of f_2 and f_3 reduces to zero modulo \mathcal{G}_2 . Check that the S-polynomial is

$$h := -ab^2 + bc - a^3b + a^2c(1 - a + a^2 - a^3 + \dots - a^{p-2}).$$

Reducing modulo f_2 we get

$$\begin{aligned} & -ab^2 + bc - a(c - ab) + a^2c(1 - a + a^2 - a^3 + \dots - a^{p-2}) = \\ & -ab^2 + bc - ac + a^2b + a^2c(1 - a + a^2 - a^3 + \dots - a^{p-2}) \xrightarrow{f_2} \\ & -ab^2 + bc - ac + (c - ab) + a^2c(1 - a + a^2 - a^3 + \dots - a^{p-2}) = \\ & -ab^2 + bc - ac + c - ab + a^2c - a^3c + a^4c - \dots - a^pc \xrightarrow{f_1} \\ & -ab^2 + bc - ac + c - ab + a^2c - a^3c + a^4c - \dots + a^{p-1}c - bc = \\ & -ab^2 - ab + c(1 - a + a^2 - a^3 \dots + a^{p-1}) \xrightarrow{f_3} \\ & -a(ab - c(1 - a + a^2 - a^3 \dots - a^{p-2})) - ab + c(1 - a + a^2 - a^3 \dots + a^{p-1}) = \\ & -a^2b + ac(1 - a + a^2 - a^3 \dots - a^{p-2}) - ab + c - ac(1 - a + a^2 - a^3 \dots - a^{p-2}) = \\ & -a^2b - ab + c \xrightarrow{f_3} 0. \quad \square \end{aligned}$$

Proof of Theorem 2.39. The ideal I_p is prime. Its variety is the parametric curve $z \mapsto ((z+1)^{p+2}, (z-1)^p, z)$. The poles and zeros of this map are $0, -1, +1, \infty$. The

tropical variety of I_p consists of the four rays defined by the valuations at these points. These rays are generated by the columns of

$$\begin{pmatrix} 0 & 0 & p+2 & -p-2 \\ 0 & p & 0 & -p \\ 1 & 0 & 0 & -1 \end{pmatrix}$$

so there are exactly four of them for any p .

To estimate the number of rays of the Gröbner fan, consider the 2-dimensional polyhedral complex obtained by slicing the restricted Gröbner fan of I with the hyperplane $x_1 + x_2 + x_3 = 1$. This complex is a planar graph, so by Euler's formula we have $v - e + f = 2$ where v is the number of vertices of the complex, e the number of edges, and f the number of faces.

Since each edge is incident to exactly two faces and each face to at least three edges, we get $3f \leq 2e$. Combining this with Euler's formula yields $v \geq 2 + f/2$. Finally, we plug in the lower bound for $f - 1$ given in Lemma 2.42 to obtain the desired result. \square

Chapter 3

CIRCUIT IDEALS

This chapter consists of the paper [9], coauthored with Anders Jensen and Rekha Thomas. The only modifications are to standardize notation and to shorten background definitions and material that appear in previous chapters of this thesis.

Given a vector configuration $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{Z}^d$, a basis ideal of \mathcal{A} is an ideal $J_{\mathcal{B}} = \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{B} \rangle \subset k[x_1, \dots, x_n]$ where \mathcal{B} spans the lattice $\mathcal{L}_{\mathcal{A}} = \{\mathbf{u} \in \mathbb{Z}^n : \sum \mathbf{a}_i u_i = \mathbf{0}\}$. Our main interest is to understand when the toric ideal, $I_{\mathcal{A}}$, of \mathcal{A} equals a given basis ideal $J_{\mathcal{B}}$ with radical $I_{\mathcal{A}}$. The circuit ideal, $J_{\mathcal{C}_{\mathcal{A}}}$, of \mathcal{A} is an example of such a basis ideal.

We study such a $J_{\mathcal{B}}$ in relation to $I_{\mathcal{A}}$ from various algebraic and combinatorial perspectives with a special focus on $J_{\mathcal{C}_{\mathcal{A}}}$. We prove that the obstruction to equality of the ideals is the existence of certain polytopes. This result is based on a complete characterization of the standard pairs/associated primes of a monomial initial ideal of $J_{\mathcal{B}}$ and their differences from those for the corresponding toric initial ideal. Eisenbud and Sturmfels proved that the embedded primes of $J_{\mathcal{B}}$ are indexed by certain faces of the cone spanned by \mathcal{A} . We provide a necessary condition for a particular face to index an embedded prime and a partial converse. Finally, we compare various polyhedral fans associated to $I_{\mathcal{A}}$ and $J_{\mathcal{C}_{\mathcal{A}}}$. The Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ is shown to refine that of $I_{\mathcal{A}}$ when the codimension of the ideals is at most two.

3.1 Introduction

Fix an ordered vector configuration $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{Z}^d$. Assume that the $d \times n$ integer matrix $A = [\mathbf{a}_1 \ \dots \ \mathbf{a}_n]$ whose columns are the elements of \mathcal{A} has rank d . Let

$\mathcal{L}_{\mathcal{A}}$ be the $(n - d)$ -dimensional saturated lattice $\{\mathbf{u} \in \mathbb{Z}^n : A\mathbf{u} = \mathbf{0}\}$. Also assume that $\mathcal{L}_{\mathcal{A}} \cap \mathbb{N}^n = \{\mathbf{0}\}$ or equivalently that A has a strictly positive vector in its row span.

The *support* of a vector $\mathbf{u} \in \mathbb{Z}^n$ is defined to be $\text{supp}(\mathbf{u}) := \{i : u_i \neq 0\}$ and \mathbf{u} is *primitive* if the greatest common divisor of its components is one.

Definition 3.1. A vector $\mathbf{c} \in \mathcal{L}_{\mathcal{A}}$ is a **circuit** of \mathcal{A} if (1) \mathbf{c} is a non-zero primitive vector and (2) there does not exist a non-zero $\mathbf{d} \in \mathcal{L}_{\mathcal{A}}$ with $\text{supp}(\mathbf{d}) \subsetneq \text{supp}(\mathbf{c})$.

Let $\mathcal{C}_{\mathcal{A}}$ denote the set of all circuits of \mathcal{A} . Write $\mathbf{c} = \mathbf{c}^+ - \mathbf{c}^-$ where $c_j^+ = c_j$ if $c_j > 0$ and 0 otherwise, and $c_j^- = -c_j$ if $c_j < 0$ and 0 otherwise. Identify $\mathbf{c} \in \mathcal{C}_{\mathcal{A}}$ with the binomial $\mathbf{x}^{\mathbf{c}^+} - \mathbf{x}^{\mathbf{c}^-} \in k[x_1, \dots, x_n] =: k[\mathbf{x}]$ where k is an algebraically closed field and $\mathbf{x}^{\mathbf{u}} := x_1^{u_1} x_2^{u_2} \cdots x_n^{u_n}$. We refer to both \mathbf{c} and $\mathbf{x}^{\mathbf{c}^+} - \mathbf{x}^{\mathbf{c}^-}$ as a circuit of \mathcal{A} and denote both lists by $\mathcal{C}_{\mathcal{A}}$.

Definition 3.2. The **circuit ideal** of \mathcal{A} is the binomial ideal $J_{\mathcal{C}_{\mathcal{A}}} := \langle \mathcal{C}_{\mathcal{A}} \rangle \subseteq k[\mathbf{x}]$.

The circuit ideal $J_{\mathcal{C}_{\mathcal{A}}}$ is a subideal of the binomial prime *toric ideal* of \mathcal{A}

$$I_{\mathcal{A}} := \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{L}_{\mathcal{A}} \rangle$$

as defined in Chapter 1. Toric ideals are the defining ideals of *toric varieties* [25] and have numerous applications in combinatorics, optimization, algebra and algebraic geometry [47]. These connections make the computability of $I_{\mathcal{A}}$ an important practical concern.

Proposition 3.3. [47, Lemma 12.2] *Given a finite subset \mathcal{B} of $\mathcal{L}_{\mathcal{A}}$, define the ideal $J_{\mathcal{B}} := \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{B} \rangle \subset k[\mathbf{x}]$. A set \mathcal{B} spans $\mathcal{L}_{\mathcal{A}}$ if and only if $(J_{\mathcal{B}} : (x_1 x_2 \cdots x_n)^\infty) = I_{\mathcal{A}}$.*

When \mathcal{B} spans $\mathcal{L}_{\mathcal{A}}$, $J_{\mathcal{B}}$ is called a **basis ideal** of \mathcal{A} . Proposition 3.3 is the starting point of the best algorithms to compute $I_{\mathcal{A}}$ since a spanning set \mathcal{B} of $\mathcal{L}_{\mathcal{A}}$ can be

computed easily and each saturation in

$$(J_{\mathcal{B}} : (x_1 x_2 \cdots x_n)^\infty) = (((J_{\mathcal{B}} : x_1^\infty) : x_2^\infty) \cdots) : x_n^\infty$$

can be achieved by a Gröbner basis calculation ([31], [47, Chapter 12].) It can be checked that $\mathcal{C}_{\mathcal{A}}$ spans $\mathcal{L}_{\mathcal{A}}$ and hence $J_{\mathcal{C}_{\mathcal{A}}}$ is a basis ideal of \mathcal{A} and $I_{\mathcal{A}} = (J_{\mathcal{C}_{\mathcal{A}}} : (x_1 x_2 \cdots x_n)^\infty)$. Further, $I_{\mathcal{A}}$ is the radical of $J_{\mathcal{C}_{\mathcal{A}}}$. Our main motivation is to understand how close the circuit ideal is to the toric ideal, and in particular, when they are equal. The majority of our theorems hold for basis ideals $J_{\mathcal{B}}$ with the property that $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$ and we use $J_{\mathcal{C}_{\mathcal{A}}}$ as our running example. The main question we address is the following.

Problem 3.4. *When does a basis ideal $J_{\mathcal{B}}$ such that $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$ equal $I_{\mathcal{A}}$?*

We investigate Problem 3.4 from several different angles. Let $\mathbb{N}\mathcal{A}$ denote the semigroup $\{\mathbf{A}\mathbf{u} : \mathbf{u} \in \mathbb{N}^n\} \subset \mathbb{Z}^d$. Both $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ are homogeneous under multi-grading by $\mathbb{N}\mathcal{A}$ with $k[\mathbf{x}]/I_{\mathcal{A}}$ having Hilbert function value one for all $\mathbf{b} \in \mathbb{N}\mathcal{A}$. In Section 3.2 we recall conditions for the equality of $I_{\mathcal{A}}$ and a basis ideal $J_{\mathcal{B}}$ and then exhibit various properties of $J_{\mathcal{B}}$ that contrast those of toric ideals. We interpret the multi-graded Hilbert function values of $k[\mathbf{x}]/J_{\mathcal{B}}$.

From the point of view of Gröbner basis theory, it is natural to investigate $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ by examining the difference between their initial ideals with respect to a fixed weight vector ω . In Section 3.3, we give a complete characterization of the associated primes of a monomial initial ideal of a basis ideal $J_{\mathcal{B}}$ with $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$ (Theorem 3.22) extending previously known characterizations of the associated primes of a monomial initial ideal of $I_{\mathcal{A}}$ [32]. The associated primes and the difference between the two monomial initial ideals are described in terms of certain polytopes that depend on \mathcal{A} and ω . Using this we answer Problem 3.4 by showing that the obstruction to equality of the ideals is the existence of certain polytopes of the above type (Theorem 3.29).

A second natural measure of the difference between the two ideals in Problem 3.4 is an understanding of the embedded primes of $J_{\mathcal{B}}$. Let $\text{cone}(\mathcal{A})$ denote the d -

dimensional cone spanned by \mathcal{A} . Record a face σ of $\text{cone}(\mathcal{A})$ as the set of indices, j , of all \mathbf{a}_j that lie on σ . Eisenbud and Sturmfels [20] proved that the associated primes of $J_{\mathcal{B}}$ are all of the form $P_{\sigma} + I_{\mathcal{A}}$ where σ is some face of $\text{cone}(\mathcal{A})$ and $P_{\sigma} := \langle x_j : j \notin \sigma \rangle$. In particular, $I_{\mathcal{A}} = P_{[n]} + I_{\mathcal{A}}$ is indexed by the full face $[n] := \{1, 2, \dots, n\}$ of $\text{cone}(\mathcal{A})$. However, not all faces of $\text{cone}(\mathcal{A})$ need index an associated prime of $J_{\mathcal{B}}$ and Eisenbud and Sturmfels raise the following problem for the special case of $J_{\mathcal{C}_{\mathcal{A}}}$.

Problem 3.5. [20, §7] *“It remains an interesting combinatorial problem to characterize the embedded primary components of the circuit ideal $J_{\mathcal{C}_{\mathcal{A}}}$. In particular, which faces of $\text{cone}(\mathcal{A})$ support an associated prime of $J_{\mathcal{C}_{\mathcal{A}}}$? An answer to this question might be valuable for the applications of binomial ideals to integer programming and statistics.”*

In Section 3.4, we give a necessary condition for a prime $P_{\sigma} + I_{\mathcal{A}}$ to be an embedded prime of a basis ideal $J_{\mathcal{B}}$ with $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$ (Theorem 3.32) using the results in Section 3.3. We also provide a partial converse to Theorem 3.32. As an application, we derive connections between the smoothness of the toric variety defined by a face σ of $\text{cone}(\mathcal{A})$ and $P_{\sigma} + I_{\mathcal{A}}$ being an embedded prime of $J_{\mathcal{C}_{\mathcal{A}}}$ when \mathcal{A} is a graded vector configuration.

Given a homogeneous ideal I and a weight vector $\omega \in \mathbb{R}^n$, let $\text{in}_{\omega}(I)$ be the initial ideal of I with respect to ω , $\sqrt{\text{in}_{\omega}(I)}$ the radical of $\text{in}_{\omega}(I)$, and $\text{top}(\text{in}_{\omega}(I))$ the intersection of the top-dimensional primary components of $\text{in}_{\omega}(I)$. These entities define three equivalence relations on \mathbb{R}^n as follows.

1. The *initial ideal* equivalence relation: $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \text{in}_{\mathbf{u}}(I) = \text{in}_{\mathbf{v}}(I)$,
2. the *top* equivalence relation: $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \text{top}(\text{in}_{\mathbf{u}}(I)) = \text{top}(\text{in}_{\mathbf{v}}(I))$, and
3. the *radical* equivalence relation: $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \sqrt{\text{in}_{\mathbf{u}}(I)} = \sqrt{\text{in}_{\mathbf{v}}(I)}$.

Recall that for any homogeneous ideal I , the initial ideal equivalence classes form the cells of the Gröbner fan of I . For $I_{\mathcal{A}}$ it is well known that the other two equivalence classes also form polyhedral fans — the radical equivalence relation gives the *secondary fan* of \mathcal{A} [7], [47, Chapter 8], and the top equivalence relation gives the *hypergeometric fan* of \mathcal{A} [42]. In Section 3.5 we prove that for $J_{\mathcal{C}_{\mathcal{A}}}$, the equivalence classes of the radical and top equivalence relations coincide with those for $I_{\mathcal{A}}$ (Theorem 3.59 and Proposition 3.60). However, the Gröbner fans of $I_{\mathcal{A}}$ and $J_{\mathcal{C}_{\mathcal{A}}}$ do not coincide in general. Corollary 3.62 proves that when the codimension of the ideals is at most two, the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ refines that of $I_{\mathcal{A}}$. Again, the theorems hold in greater generality than for circuit ideals.

3.2 Properties of $I_{\mathcal{A}}$ versus $J_{\mathcal{B}}$

Consider $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathbb{Z}^d$ and the lattice $\mathcal{L}_{\mathcal{A}}$ as in the introduction. Let $\mathcal{B} \subseteq \mathbb{Z}^n$ be a spanning set of $\mathcal{L}_{\mathcal{A}}$ and consider the basis ideal

$$J_{\mathcal{B}} := \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{B} \rangle \subseteq k[\mathbf{x}].$$

The toric ideal $I_{\mathcal{A}}$ and the circuit ideal $J_{\mathcal{C}_{\mathcal{A}}}$ are of the form $J_{\mathcal{B}}$. In particular, since $I_{\mathcal{A}}$ is $J_{\mathcal{B}}$ for $\mathcal{B} = \mathcal{L}_{\mathcal{A}}$, every basis ideal $J_{\mathcal{B}}$ is contained in $I_{\mathcal{A}}$.

In this section we first collect conditions equivalent to the equality of $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$. Many of these stem from combinatorics and optimization and most are well known [15], [47]. We then contrast $J_{\mathcal{B}}$ with $I_{\mathcal{A}}$ in light of these conditions, using $J_{\mathcal{B}} = J_{\mathcal{C}_{\mathcal{A}}}$ in our examples.

Consider the semigroup homomorphism $\pi : \mathbb{N}^n \rightarrow \mathbb{N}\mathcal{A}$ such that $\mathbf{u} \mapsto \mathbf{A}\mathbf{u}$. The ideal $J_{\mathcal{B}}$ is homogeneous under the \mathcal{A} -grading of $k[\mathbf{x}]$ by $\deg(x_i) = \mathbf{a}_i$ for $i = 1, \dots, n$ since every binomial of the form $\mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}}$ in $J_{\mathcal{B}}$ is \mathcal{A} -homogeneous with \mathcal{A} -degree $\pi(\mathbf{u}) = \mathbf{A}\mathbf{u} = \mathbf{A}\mathbf{v} = \pi(\mathbf{v})$. Let $\mathbf{H}_{J_{\mathcal{B}}} : \mathbb{N}\mathcal{A} \rightarrow \mathbb{N}$ be the \mathcal{A} -graded Hilbert function of $k[\mathbf{x}]/J_{\mathcal{B}}$ given by $\mathbf{b} \mapsto \dim_k(k[\mathbf{x}]/J_{\mathcal{B}})_{\mathbf{b}}$. Let $\mathbf{H}_{I_{\mathcal{A}}}$ be the same for $I_{\mathcal{A}}$.

Since $\mathcal{L}_{\mathcal{A}} \cap \mathbb{N}^n = \{\mathbf{0}\}$, for each $\mathbf{b} \in \mathbb{N}\mathcal{A}$, the polyhedron $P_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}_{\geq 0}^n : \mathbf{A}\mathbf{x} = \mathbf{b}\}$

is bounded [43] which implies that $\pi^{-1}(\mathbf{b}) := \{\mathbf{x} \in \mathbb{N}^n : A\mathbf{x} = \mathbf{b}\} = P_{\mathbf{b}} \cap \mathbb{N}^n$ is finite for all $\mathbf{b} \in \mathbb{N}A$. For a fixed $\mathbf{b} \in \mathbb{N}A$, the set $\pi^{-1}(\mathbf{b})$ admits two natural graphs as follows. First, choose a binomial generating set $G(J_{\mathcal{B}})$ of $J_{\mathcal{B}}$ and let $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ be the graph on $\pi^{-1}(\mathbf{b})$ such that \mathbf{u} is adjacent to \mathbf{v} in $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ if $\mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}}$ is a monomial multiple of a binomial in $G(J_{\mathcal{B}})$. Next, fix a generic weight vector $\omega \in \mathbb{R}^n$ in the sense that the initial ideal $\text{in}_{\omega}(J_{\mathcal{B}})$ is a monomial ideal. Let $G_{\omega}(J_{\mathcal{B}})$ be the marked reduced Gröbner basis of $J_{\mathcal{B}}$ with respect to ω . Elements of $G_{\omega}(J_{\mathcal{B}})$ are \mathcal{A} -homogeneous binomials and the Gröbner basis being marked means that the first term in each binomial f is its initial term $\text{in}_{\omega}(f)$. Construct the *directed* graph $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ on $\pi^{-1}(\mathbf{b})$ by drawing an arrow from \mathbf{u} to \mathbf{v} in $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ if and only if $\mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}}$ is a monomial multiple of some marked binomial in $G_{\omega}(J_{\mathcal{B}})$. In the special case of $J_{\mathcal{B}} = I_{\mathcal{A}}$, we denote $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ by just $\mathcal{F}(\mathbf{b})$ and $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ by $\mathcal{F}_{\omega}(\mathbf{b})$.

Lemma 3.6. [15, Theorem 1.1] *Vectors $\mathbf{u}, \mathbf{v} \in \pi^{-1}(\mathbf{b})$ are in the same component of $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ if and only if $\mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}}$ lies in $J_{\mathcal{B}}$.*

Lemma 3.6 shows that while the edges in $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ depend on the choice of generating set $G(J_{\mathcal{B}})$, the components, and in particular the number of components, do not depend on this choice. Further, $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ and $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ partition $\pi^{-1}(\mathbf{b})$ identically into components. The following theorem collects results from [15] and [47].

Theorem 3.7. *The following statements are equivalent.*

1. *The ideals $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ are equal.*
2. *For every $\mathbf{b} \in \mathbb{N}A$, the graph $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ is connected.*
3. *For every $\mathbf{b} \in \mathbb{N}A$, the digraph $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ has a unique sink. (In this case, the unique sink \mathbf{u} in $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ is the optimal solution of the integer program minimize $\{\omega \cdot \mathbf{x} : A\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathbb{N}^n\}$.)*

4. For every $\mathbf{b} \in \mathbb{N}\mathcal{A}$ and generic weight vector $\omega \in \mathbb{R}^n$, $\text{in}_\omega(J_{\mathcal{B}})$ has a unique standard monomial of \mathcal{A} -degree \mathbf{b} . (In this case, the standard monomial of \mathcal{A} -degree \mathbf{b} is $\mathbf{x}^{\mathbf{u}}$ where \mathbf{u} is the unique sink in $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$.)
5. For every $\mathbf{b} \in \mathbb{N}\mathcal{A}$, the Hilbert function value $\mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b})$ is one.

Proof: Statements (2)–(5) are all true if $J_{\mathcal{B}}$ equals $I_{\mathcal{A}}$; see [47, Chapters 4,5,10]. Further, $I_{\mathcal{A}} = J_{\mathcal{B}}$ if and only if $G_\omega(I_{\mathcal{A}}) = G_\omega(J_{\mathcal{B}})$ if and only if for each \mathbf{b} , $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$ equals $\mathcal{F}_\omega(\mathbf{b})$ and hence if and only if $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ and $\mathcal{F}(\mathbf{b})$ have the same components. Hence (1) is equivalent to (2) and (3). Since $J_{\mathcal{B}} \subseteq I_{\mathcal{A}}$, the two ideals are equal if and only if (5). The equivalence of (3) and (4) follows from Lemma 3.10 below. \square

Remark 3.8. 1. The connectivity of $\mathcal{F}(\mathbf{b})$ was used in [16], in the context of statistical sampling, to devise random walks on $\pi^{-1}(\mathbf{b})$. Of particular interest was the case where $I_{\mathcal{A}} = J_{\mathcal{C}_{\mathcal{A}}}$ which allows $\pi^{-1}(\mathbf{b})$ to be connected using circuits of \mathcal{A} . In Section 3.3 we will see that under further assumptions on $J_{\mathcal{B}}$, for most $\mathbf{b} \in \mathbb{N}\mathcal{A}$, $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ is in fact connected (Theorem 3.26), and that the set of $\mathbf{b} \in \mathbb{N}\mathcal{A}$ for which $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ is disconnected can be described precisely. See also [15].

2. The equality of $I_{\mathcal{A}}$ and $J_{\mathcal{C}_{\mathcal{A}}}$ will allow all integer programs of the form

$$\text{minimize } \{\omega \cdot \mathbf{x} : A\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathbb{N}^n\}$$

as \mathbf{b} and ω vary to be solved by reduced Gröbner bases of $J_{\mathcal{C}_{\mathcal{A}}}$. The significance of this is that the circuits of \mathcal{A} are precisely the primitive edge directions of the polyhedra $P_{\mathbf{b}}$ as \mathbf{b} varies in $\mathbb{N}\mathcal{A}$ and hence the directions taken by the simplex algorithm in solving linear programs of the form

$$\text{minimize } \{\omega \cdot \mathbf{x} : A\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathbb{R}_{\geq 0}^n\}.$$

Using circuit ideals, we now contrast various properties of $J_{\mathcal{B}}$ with those of $I_{\mathcal{A}}$.

Proposition 3.9. 1. *The graph $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ may have arbitrarily many components, even if we restrict to the case of $A \in \mathbb{Z}^{1 \times 3}$.*

2. *The standard monomials of $\text{in}_\omega(J_{\mathcal{B}})$ of \mathcal{A} -degree \mathbf{b} are not necessarily the cheapest monomials of that degree with respect to ω .*

Proof: (1) For any natural number $k \geq 2$, let $A_k = (k \ 2k+1 \ 3k+1)$ and let \mathcal{B}_k be the set of circuits of \mathcal{A}_k . Since the three entries of A_k are pairwise relatively prime, the circuits are $x^{2k+1} - y^k$, $x^{3k+1} - z^k$, and $y^{3k+1} - z^{2k+1}$ with \mathcal{A} -degrees $2k^2+k$, $3k^2+k$, and $6k^2+5k+1$ respectively. Thus the graph $\mathcal{F}^{\mathcal{B}_k}(b)$ has no edges when $b < 2k^2+k$. In particular, this holds if we take $b = m(3k+1)$ for $m = \lfloor k/2 \rfloor$. This particular $\mathcal{F}^{\mathcal{B}_k}(b)$ has at least $m+1$ vertices $\{(j, j, m-j) : 0 \leq j \leq m\}$, so it has at least $m+1 = \lfloor k/2 \rfloor + 1$ components.

(2) Consider $\mathcal{A} = \{3, 4, 5\}$. The graded reverse lexicographic Gröbner basis of $J_{\mathcal{A}}$ with $a \succ b \succ c$ is

$$\{a^4 - b^3, ab^3 - c^3, b^5 - c^4, b^2c^3 - ac^4, a^3c^3 - bc^4, a^2bc^4 - c^6\}.$$

The monomials of degree 17 are a^4c, a^3b^2, abc^2 and b^3c of which the last three are standard monomials of the above grevlex initial ideal of $J_{\mathcal{A}}$. However, we see that the non-standard monomial a^4c is cheaper than the standard monomial a^3b^2 . \square

We now prove that $\mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b})$ equals the number of components of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$, or equivalently, of $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$. Proposition 3.9 (1) then shows that the values of $\mathbf{H}_{J_{\mathcal{B}}}$ can be arbitrarily large even for d and n fixed. In contrast, $\mathbf{H}_{J_{\mathcal{A}}}(\mathbf{b}) = 1$ for all $\mathbf{b} \in \mathbb{N}^A$.

Lemma 3.10. *Each component of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$ has a unique sink \mathbf{u} and $\mathbf{x}^{\mathbf{u}}$ is the unique standard monomial of $\text{in}_\omega(J_{\mathcal{B}})$ among all monomials $\mathbf{x}^{\mathbf{v}}$ such that \mathbf{v} is in the same component as \mathbf{u} . In particular, a monomial $\mathbf{x}^{\mathbf{u}}$ of \mathcal{A} -degree \mathbf{b} is a standard monomial of $\text{in}_\omega(J_{\mathcal{B}})$ if and only if for all $\mathbf{v} \neq \mathbf{u}$ in the same component of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$ as \mathbf{u} , $\omega \cdot \mathbf{u} < \omega \cdot \mathbf{v}$.*

Proof: Let D be an arbitrary component of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$, \mathbf{v} be an arbitrary vertex in D , and $\mathbf{x}^{\mathbf{u}}$ be the normal form of $\mathbf{x}^{\mathbf{v}}$ with respect to $G_\omega(J_{\mathcal{B}})$. Then $\mathbf{x}^{\mathbf{u}}$ is a standard monomial of $\text{in}_\omega(J_{\mathcal{B}})$ and by Lemma 3.6, \mathbf{u} is in D . If $\mathbf{x}^{\mathbf{u}'}$ is another standard monomial of $\text{in}_\omega(J_{\mathcal{B}})$ with \mathbf{u}' in D , then by Lemma 3.6, $f := \mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{u}'} \in J_{\mathcal{B}}$ with $\text{in}_\omega(f)$ equal to either $\mathbf{x}^{\mathbf{u}}$ or $\mathbf{x}^{\mathbf{u}'}$, a contradiction. This implies that $\mathbf{x}^{\mathbf{u}}$ is the unique normal form of all $\mathbf{x}^{\mathbf{v}}$, $\mathbf{v} \in D$ and hence it is the unique sink in D . The remaining assertions now follow. \square

Proposition 3.11. *The Hilbert function value $\mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b})$ equals the number of components of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$.*

Proof: By Lemma 3.10 (2), each component of $\mathcal{F}_\omega^{\mathcal{B}}(\mathbf{b})$ contributes precisely one standard monomial of $\text{in}_\omega(J_{\mathcal{B}})$. The number of standard monomials of $\text{in}_\omega(J_{\mathcal{B}})$ of degree \mathbf{b} equals $\dim_k(k[\mathbf{x}]/\text{in}_\omega(J_{\mathcal{B}}))_{\mathbf{b}} = \dim_k(k[\mathbf{x}]/J_{\mathcal{B}})_{\mathbf{b}} = \mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b})$. \square

Example 3.12. When $I_{\mathcal{A}} \neq J_{\mathcal{B}}$, the distribution of values of $\mathbf{H}_{J_{\mathcal{B}}}$ can be quite complicated. In Figure 3.1, we plot these values for $\mathcal{B} = \mathcal{C}_{\mathcal{A}}$ of

$$A = \begin{pmatrix} 1 & 3 & 2 & 4 \\ 1 & 4 & 5 & 2 \end{pmatrix}.$$

The boundary of $\text{cone}(\mathcal{A})$ is shown by dashed lines. Notice that deep in the interior of the cone, all of the values are one. Theorem 3.26 proves this fact.

3.3 Monomial Initial Ideals of the Circuit Ideal

In this section let \mathcal{A} and \mathcal{B} be as in Section 3.2 with the further assumption that $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$. This assumption always holds when \mathcal{B} is the set of circuits of \mathcal{A} [20, Proposition 7.10]. Fix a generic weight vector $\omega \in \mathbb{R}^n$ such that $\text{in}_\omega(I_{\mathcal{A}})$ and $\text{in}_\omega(J_{\mathcal{B}})$ are both monomial ideals. The main result of this section is Theorem 3.22 which characterizes the associated primes of $\text{in}_\omega(J_{\mathcal{B}})$ in terms of certain polytopes defined




Figure 3.1: The distribution of values of \mathbf{H}_{J_B} for the circuit ideal of the matrix A in Example 3.12.

from \mathcal{A} and ω and their lattice points. This theorem generalizes Theorem 2.5 in [32] which gave a complete characterization of the associated primes of $\text{in}_\omega(I_{\mathcal{A}})$ in terms of certain *lattice-point-free* polytopes defined from \mathcal{A} and ω . Using Theorem 3.22, we describe the similarities and differences between the associated primes (standard pairs) of $\text{in}_\omega(I_{\mathcal{A}})$ and $\text{in}_\omega(J_{\mathcal{B}})$, and give an answer to Problem 3.4 (Theorem 3.29).

Proposition 3.13. *Let I and J be homogeneous ideals in $k[\mathbf{x}]$ with $\sqrt{J} = I$. Then $\sqrt{\text{in}_\omega(J)} = \sqrt{\text{in}_\omega(I)}$ for all $\omega \in \mathbb{R}^n$.*

Proof: Since $\sqrt{J} = I$, $J \subseteq I$ which implies that $\text{in}_\omega(J) \subseteq \text{in}_\omega(I)$ and hence $\sqrt{\text{in}_\omega(J)} \subseteq \sqrt{\text{in}_\omega(I)}$ for all $\omega \in \mathbb{R}^n$. To prove the other inclusion we first observe that $\sqrt{\text{in}_\omega(I)}$ is ω -homogeneous since $\text{in}_\omega(I)$ is. Hence it suffices to show that any homogeneous element in $\sqrt{\text{in}_\omega(I)}$ is also in $\sqrt{\text{in}_\omega(J)}$. Let $f \in \sqrt{\text{in}_\omega(I)}$ be ω -homogeneous. Then there exists some m such that $f^m \in \text{in}_\omega(I)$. The polynomial f^m is also ω -homogeneous, so $f^m = \text{in}_\omega(F)$ for some $F \in I$. Since $\sqrt{J} = I$, $F^k \in J$ for some k , and $\text{in}_\omega(F^k) = \text{in}_\omega(F)^k = f^{mk}$. Hence, $f \in \sqrt{\text{in}_\omega(J)}$. \square

Definition 3.14. [47, Chapter 8]

1. The *regular triangulation* of \mathcal{A} with respect to ω is the simplicial complex Δ_ω on the vertex set $[n] = \{1, \dots, n\}$ such that $\{i_1, \dots, i_r\} \subseteq [n]$ is a face of Δ_ω if and only if there exists a vector $\mathbf{c} \in \mathbb{R}^d$ such that $\mathbf{a}_j \cdot \mathbf{c} = \omega_j$ if $j \in \{i_1, \dots, i_r\}$ and $\mathbf{a}_j \cdot \mathbf{c} < \omega_j$ if $j \notin \{i_1, \dots, i_r\}$.
2. The *Stanley-Reisner ideal* of a simplicial complex Δ on $[n]$ is the ideal in $k[\mathbf{x}]$ generated by the monomials $\mathbf{x}_\sigma := \prod_{i \in \sigma} x_i$ for each minimal nonface σ of Δ .

Theorem 8.3 in [47] states that $\sqrt{\text{in}_\omega(I_{\mathcal{A}})}$ is the Stanley-Reisner ideal of the regular triangulation Δ_ω of \mathcal{A} . For a set $\sigma \subseteq [n]$ define $P_\sigma := \langle x_j : j \notin \sigma \rangle \subset k[\mathbf{x}]$. Note that P_σ is a monomial prime ideal such that $k[\mathbf{x}]/P_\sigma$ has Krull dimension $|\sigma|$.

Corollary 3.15. *For a basis ideal $J_{\mathcal{B}}$ with $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$, the following hold.*

1. *All the associated primes of $\text{in}_{\omega}(J_{\mathcal{B}})$ are monomial ideals of the form P_{σ} where σ is a face of the simplicial complex Δ_{ω} .*
2. *The prime P_{σ} is a minimal prime of $\text{in}_{\omega}(J_{\mathcal{B}})$ if and only if σ is a maximal face of Δ_{ω} .*

Proof: If I is the Stanley-Reisner ideal of a simplicial complex Δ on $[n]$, then I has the irredundant prime decomposition $I = \bigcap_{\sigma \in \max(\Delta)} P_{\sigma}$ where $\max(\Delta)$ is the set of maximal faces of Δ [47, Chapter 8]. Thus the minimal primes of $\text{in}_{\omega}(J_{\mathcal{B}})$, which equal the minimal primes of $\sqrt{\text{in}_{\omega}(J_{\mathcal{B}})} = \sqrt{\text{in}_{\omega}(I_{\mathcal{A}})}$ (by Proposition 3.13), are the primes P_{σ} as σ varies in $\max(\Delta_{\omega})$, proving (2). If P_{τ} is an embedded prime of $\text{in}_{\omega}(J_{\mathcal{B}})$, then $\tau \subset \sigma$ for some $\sigma \in \max(\Delta_{\omega})$. This implies that τ is a lower dimensional face of Δ_{ω} , proving (1). \square

If τ is a lower dimensional face of Δ_{ω} , P_{τ} may or may not be an embedded prime of $\text{in}_{\omega}(J_{\mathcal{B}})$. Theorem 3.22 characterizes the lower dimensional faces of Δ_{ω} that index embedded primes of $\text{in}_{\omega}(J_{\mathcal{B}})$.

Example 3.16. Let \mathcal{B} be the set of circuits of the matrix

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 \end{pmatrix}.$$

Using the program Gfan [34] we find that both $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ have eight distinct monomial initial ideals. Table 3.1 gives a representative weight vector ω for each pair of initial ideals and verifies Proposition 3.13.

Remark 3.17. If we drop the assumption that $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$, then it need not be that $\sqrt{\text{in}_{\omega}(J_{\mathcal{B}})}$ is the Stanley-Reisner ideal of any regular triangulation of \mathcal{A} . For the matrix A in Example 3.16, the set $\mathcal{B} = \{(1, -2, 1, 0), (2, -3, 0, 1)\}$ spans the lattice

Table 3.1: Comparison of initial ideals of $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ from Example 3.16.

ω	$\text{in}_{\omega}(I_{\mathcal{A}})$	$\text{in}_{\omega}(J_{\mathcal{B}})$	radical of both initial ideals
(10, 0, 1, 3)	$\langle ac, ad, bd \rangle$	$\langle ac, a^2d, bd, ad^2 \rangle$	$\langle ac, ad, bd \rangle$
(10, 0, 3, 1)	$\langle ac, c^2, ad \rangle$	$\langle ac, c^2, a^2d, abd, ad^2 \rangle$	$\langle c, ad \rangle$
(3, 1, 10, 0)	$\langle ac, bc, c^2, a^2d \rangle$	$\langle ac, b^2c, c^2, a^2d, bcd \rangle$	"
(1, 3, 10, 0)	$\langle b^3, ac, bc, c^2 \rangle$	$\langle b^3, ac, b^2c, c^2, bcd \rangle$	$\langle b, c \rangle$
(1, 5, 3, 0)	$\langle b^2, bc, c^2 \rangle$	$\langle b^2, abc, c^2, bcd \rangle$	"
(0, 10, 3, 1)	$\langle b^2, bc, c^3, bd \rangle$	$\langle b^2, abc, bc^2, c^3, bd \rangle$	"
(1, 3, 0, 10)	$\langle b^2, ad, bd \rangle$	$\langle b^2, a^2d, bd, acd, ad^2 \rangle$	$\langle b, ad \rangle$
(3, 10, 0, 1)	$\langle b^2, bc, bd, ad^2 \rangle$	$\langle b^2, abc, bc^2, bd, ad^2 \rangle$	"

$\mathcal{L}_{\mathcal{A}}$ and $J_{\mathcal{B}} = \langle ac - b^2, a^2d - b^3 \rangle$. The radical of $J_{\mathcal{B}}$ is $\langle bc - ad, b^2 - ac \rangle$, a proper subideal of $I_{\mathcal{A}}$. The grevlex initial ideal of $J_{\mathcal{B}}$ with $a \succ b \succ c \succ d$ is $\langle b^2, abc, a^2c^2 \rangle$ whose radical is $\langle b, ac \rangle$. This ideal is not listed in the last column of Table 3.1.

We now establish the necessary definitions and lemmas for Theorem 3.22. The associated primes of a monomial ideal M can be studied via a combinatorial construction introduced in [50] called the *standard pair decomposition* of M .

Definition 3.18. Let $M \subset k[\mathbf{x}]$ be a monomial ideal, $\mathbf{x}^{\mathbf{u}}$ a standard monomial of M and $\sigma \subseteq [n]$. Then $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is an *admissible pair* of M if:

1. $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$,
2. all monomials in $\mathbf{x}^{\mathbf{u}} \cdot k[x_j : j \in \sigma]$ are standard monomials of M .

An admissible pair $(\mathbf{x}^{\mathbf{u}}, \sigma)$ of M is called a *standard pair* of M if there does not exist another admissible pair $(\mathbf{x}^{\mathbf{v}}, \tau)$ such that $\mathbf{v} \leq \mathbf{u}$ and $\text{supp}(\mathbf{u} - \mathbf{v}) \cup \sigma \subseteq \tau$.

The (unique) decomposition of the standard monomials of M given by its standard pairs is the *standard pair decomposition* of M . Let $\text{Ass}(I)$ denote the set of associated

primes of an ideal I . Since M is a monomial ideal, all elements of $\text{Ass}(M)$ have the form P_σ for some $\sigma \subseteq [n]$. Standard pairs of M are related to $\text{Ass}(M)$ as follows.

Proposition 3.19. [50, Lemmas 3.3 and 3.5]

1. $P_\sigma \in \text{Ass}(M)$ if and only if M has a standard pair of the form (\cdot, σ) .
2. P_σ is a minimal prime of M if and only if $(1, \sigma)$ is a standard pair of M .

We now define the polytopes needed in Theorem 3.22. Fix a matrix $G \in \mathbb{Z}^{n \times (n-d)}$ whose columns form a basis for the lattice \mathcal{L}_A . Such a G is called an *integer Gale dual* of A . In particular, the columns of G span the kernel of A as an \mathbb{R} -vector space. For $\mathbf{u} \in \mathbb{N}^n$ let

$$Q_{\mathbf{u}} := \{\mathbf{z} \in \mathbb{R}^{n-d} : G\mathbf{z} \leq \mathbf{u}\}.$$

Recall that by assumption, $P_{\mathbf{b}} = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^n : A\mathbf{x} = \mathbf{b}\}$ is a polytope for all $\mathbf{b} \in \mathbb{N}A$. The polyhedron $Q_{\mathbf{u}}$ is the image of $P_{A\mathbf{u}}$ under the isomorphism

$$\phi_{\mathbf{u}} : \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = A\mathbf{u}\} \rightarrow \mathbb{R}^{n-d} \text{ such that } \mathbf{x} \mapsto \mathbf{z} \text{ where } \mathbf{u} - G\mathbf{z} = \mathbf{x}.$$

For each $\mathbf{x} \in \mathbb{R}^n$ such that $A\mathbf{x} = A\mathbf{u}$, $\mathbf{u} - \mathbf{x} = G\mathbf{z}$ for some $\mathbf{z} \in \mathbb{R}^{n-d}$ since $\mathbf{u} - \mathbf{x} \in \ker(A) = \{G\mathbf{z} : \mathbf{z} \in \mathbb{R}^{n-d}\}$. Further, this \mathbf{z} is unique since the columns of G are linearly independent. The vector \mathbf{u} maps to $\mathbf{0}$ under $\phi_{\mathbf{u}}$ and hence $\mathbf{0} \in Q_{\mathbf{u}}$.

Next, define

$$Q_{\mathbf{u}, \omega} := Q_{\mathbf{u}} \cap \{\mathbf{z} \in \mathbb{R}^{n-d} : (-\omega G)\mathbf{z} \leq \mathbf{0}\},$$

the subpolytope of $Q_{\mathbf{u}}$ created by adding one new inequality depending on ω . For σ a face of Δ_ω , further define

$$Q_{\mathbf{u}, \omega}^{\bar{\sigma}} := \{\mathbf{z} \in \mathbb{R}^{n-d} : (G\mathbf{z} \leq \mathbf{u})^{\bar{\sigma}}, (-\omega G)\mathbf{z} \leq \mathbf{0}\}$$

where $(G\mathbf{z} \leq \mathbf{u})^{\bar{\sigma}}$ denotes the subsystem of inequalities indexed by $\bar{\sigma}$ in $G\mathbf{z} \leq \mathbf{u}$. Theorem 1 in [33] proves that the relaxation $Q_{\mathbf{u}, \omega}^{\bar{\sigma}}$ is a polytope (i.e. that it is bounded.)

For a non-zero lattice point $\mathbf{z} \in Q_{\mathbf{u}, \omega}^{\bar{\sigma}}$, set $\mathbf{m}_{\mathbf{z}} := (\mathbf{u} - G\mathbf{z})^-$. Let G_i denote the i -th row of G .

- Remark 3.20.**
1. The i -th component $(\mathbf{m}_z)_i > 0$ if and only if \mathbf{z} violates the i -th inequality $G_i \mathbf{z} \leq u_i$ among the inequalities $G\mathbf{z} \leq \mathbf{u}$ defining $Q_{\mathbf{u},\omega}$.
 2. Since every $\mathbf{z} \in Q_{\mathbf{u},\omega}^{\bar{\sigma}}$ satisfies $G_i \mathbf{z} \leq u_i$ for $i \notin \sigma$, the support of \mathbf{m}_z is contained in σ .
 3. The vector \mathbf{m}_z is the component-wise smallest vector \mathbf{m} in \mathbb{N}^n with support in σ such that $\mathbf{z} \in Q_{\mathbf{u}+\mathbf{m},\omega}$.
 4. By the definition of \mathbf{m}_z , $\mathbf{u} + \mathbf{m}_z - G\mathbf{z} \in \mathbb{N}^n$.

Theorem 3.22 will generalize the following theorem for toric ideals.

Theorem 3.21. [32, Theorem 2.5] *Assume $\mathbf{u} \in \mathbb{N}^n$ and $\sigma \in \Delta_\omega$ such that $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$. Then $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_\omega(I_A)$ if and only if the following two conditions hold.*

1. *There are no non-zero lattice points in $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$.*
2. *For every $i \in \bar{\sigma}$ there is a non-zero lattice point in $Q_{\mathbf{u},\omega}^{\overline{\sigma \cup \{i\}}}$.*

Theorem 3.22 is analogous, but involves an algebraic component rather than being purely polyhedral. Recall that $\mathbf{x}_\sigma = \prod_{i \in \sigma} x_i$.

Theorem 3.22. *Assume $\mathbf{u} \in \mathbb{N}^n$ and $\sigma \in \Delta_\omega$ such that $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$. Then $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_\omega(J_B)$ if and only if the following two conditions hold.*

1. *For each non-zero lattice point \mathbf{z} in $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$, $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \notin (J_B : \mathbf{x}_\sigma^\infty)$.*
2. *For each $i \in \bar{\sigma}$, there exists some non-zero lattice point $\mathbf{z} \in Q_{\mathbf{u},\omega}^{\overline{\sigma \cup \{i\}}}$ such that $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \in (J_B : \mathbf{x}_{\sigma \cup \{i\}}^\infty)$.*

We first use Theorem 3.22 to reprove Theorem 3.21.

Proof of Theorem 3.21: Since I_A is prime and monomial free, $(I_A : \mathbf{x}_\tau^\infty) = I_A$ for all $\tau \subseteq [n]$. Thus if \mathbf{z} is a non-zero lattice point in $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$, then $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \in I_A = (I_A : \mathbf{x}_\sigma^\infty)$. Hence, Theorem 3.22 (1) holds if and only if there are no non-zero lattice points in $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$. Similarly, Theorem 3.22 (2) holds in the toric situation if and only if for every $i \in \bar{\sigma}$ there is a non-zero lattice point in $Q_{\mathbf{u},\omega}^{\overline{\sigma \cup \{i\}}}$. \square

Proof of Theorem 3.22: (\Rightarrow): Suppose $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_\omega(J_B)$. Then by Corollary 3.15 $\sigma \in \Delta_\omega$, and $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$. Suppose \mathbf{z} is a non-zero lattice point in $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$. Then $-(\omega G)\mathbf{z} \leq 0$, and because ω is generic, we may assume $-(\omega G)\mathbf{z} < 0$. For any $\mathbf{m} \in \mathbb{N}^n$ with support contained in σ , $\mathbf{x}^{\mathbf{u}+\mathbf{m}}$ is a standard monomial of $\text{in}_\omega(J_B)$ since $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair. If further, $\mathbf{m} \geq \mathbf{m}_z = (\mathbf{u} - G\mathbf{z})^-$, then $\mathbf{u} + \mathbf{m} - G\mathbf{z} \in \mathbb{N}^n$ and $A(\mathbf{u} + \mathbf{m} - G\mathbf{z}) = A(\mathbf{u} + \mathbf{m})$ since $AG = 0$. Also, $\omega \cdot (\mathbf{u} + \mathbf{m} - G\mathbf{z}) = \omega \cdot (\mathbf{u} + \mathbf{m}) - (\omega G)\mathbf{z} < \omega \cdot (\mathbf{u} + \mathbf{m})$ since $-(\omega G)\mathbf{z} < 0$. Therefore, by Lemma 3.10, $\mathbf{x}^{\mathbf{u}+\mathbf{m}} - \mathbf{x}^{\mathbf{u}+\mathbf{m}-G\mathbf{z}} \notin J_B$. In particular, $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \notin J_B$ and for all $\mathbf{m}' \in \mathbb{N}^n$ with support in σ , $\mathbf{x}^{\mathbf{m}'}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) \notin J_B$. Rewriting, this gives $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \notin (J_B : \mathbf{x}_\sigma^\infty)$ and (1) holds.

Suppose $i \notin \sigma$. Then there exists some $\mathbf{m} \in \mathbb{N}^n$ with $\text{supp}(\mathbf{m}) \subseteq \sigma$ and $p > 0$ such that $\mathbf{x}^{\mathbf{u}+\mathbf{m}}x_i^p \in \text{in}_\omega(J_B)$. Let \mathbf{q} be the unique sink in the same component of $\mathcal{F}_\omega^B(A(\mathbf{u} + \mathbf{m} + p\mathbf{e}_i))$ as $\mathbf{u} + \mathbf{m} + p\mathbf{e}_i$. Note that $\mathbf{q} \neq \mathbf{u} + \mathbf{m} + p\mathbf{e}_i$ since $\mathbf{x}^{\mathbf{q}} \notin \text{in}_\omega(J_B)$. Let $\mathbf{z} \in \mathbb{Z}^{n-d}$ be such that $\mathbf{q} = \mathbf{u} + \mathbf{m} + p\mathbf{e}_i - G\mathbf{z}$. Then $\mathbf{u} + \mathbf{m} + p\mathbf{e}_i$ maps to $\mathbf{0}$ and \mathbf{q} maps to $\mathbf{z} \neq \mathbf{0}$ in $Q_{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i}$ under the map $\phi_{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i}$. Since $\omega \cdot \mathbf{q} = \omega \cdot (\mathbf{u} + \mathbf{m} + p\mathbf{e}_i - G\mathbf{z}) < \omega \cdot (\mathbf{u} + \mathbf{m} + p\mathbf{e}_i)$, we see that $-(\omega G)\mathbf{z} < 0$. Therefore, \mathbf{z} is a lattice point in $Q_{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i,\omega}$ and hence in $Q_{\mathbf{u},\omega}^{\overline{\sigma \cup \{i\}}}$ obtained by throwing away the inequalities of $G\mathbf{z} \leq \mathbf{u}$ indexed by $\sigma \cup \{i\}$ from $Q_{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i,\omega}$. This is because $\text{supp}(\mathbf{m} + p\mathbf{e}_i) \subseteq \sigma \cup \{i\}$. By definition, $\mathbf{m}_z \leq \mathbf{m} + p\mathbf{e}_i$ since \mathbf{m}_z is the component-wise smallest vector \mathbf{m}' with support in $\sigma \cup \{i\}$ such that $\mathbf{z} \in Q_{\mathbf{u}+\mathbf{m}',\omega}^{\overline{\sigma \cup \{i\}}}$ and we know that $\mathbf{z} \in Q_{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i}$. Since $\mathbf{q} = \mathbf{u} + \mathbf{m} + p\mathbf{e}_i - G\mathbf{z}$ lies in the same component of $\mathcal{F}_\omega^B(A(\mathbf{u} + \mathbf{m} + p\mathbf{e}_i))$ as

$\mathbf{u} + \mathbf{m} + p\mathbf{e}_i$, by Lemma 3.6,

$$\mathbf{x}^{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i} - \mathbf{x}^{\mathbf{u}+\mathbf{m}+p\mathbf{e}_i-G\mathbf{z}} = \mathbf{x}^{\mathbf{m}+p\mathbf{e}_i-\mathbf{m}_z}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) \in J_{\mathcal{B}}.$$

This implies that $\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \in (J_{\mathcal{B}} : \mathbf{x}_{\sigma \cup \{i\}}^{\infty})$ and (2) holds.

(\Leftarrow): Suppose (1) and (2) hold for some $\sigma \in \Delta_{\omega}$ and some $\mathbf{u} \in \mathbb{N}^n$ with support in $\bar{\sigma}$. We first show that $\mathbf{x}^{\mathbf{u}+\mathbf{m}}$ is a standard monomial of $\text{in}_{\omega}(J_{\mathcal{B}})$ where $\mathbf{m} \in \mathbb{N}^n$ is an arbitrary vector with $\text{supp}(\mathbf{m}) \subseteq \sigma$. Suppose \mathbf{z} is a non-zero lattice point in $Q_{\mathbf{u}+\mathbf{m},\omega}$. Then \mathbf{z} is also a non-zero lattice point in the relaxation $Q_{\mathbf{u}+\mathbf{m},\omega}^{\bar{\sigma}} = Q_{\mathbf{u},\omega}^{\bar{\sigma}}$. Compute \mathbf{m}_z for this \mathbf{u} and \mathbf{z} . Since $\mathbf{z} \in Q_{\mathbf{u}+\mathbf{m},\omega}$, $\mathbf{m}_z \leq \mathbf{m}$. By (1), $(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) \notin (J_{\mathcal{B}} : \mathbf{x}_{\sigma}^{\infty})$ which implies that

$$\mathbf{x}^{\mathbf{m}-\mathbf{m}_z}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) = \mathbf{x}^{\mathbf{u}+\mathbf{m}} - \mathbf{x}^{\mathbf{u}+\mathbf{m}-G\mathbf{z}} \notin J_{\mathcal{B}}.$$

Thus for each $\mathbf{z} \neq \mathbf{0}$ in $Q_{\mathbf{u}+\mathbf{m},\omega}$, the vector $\mathbf{u} + \mathbf{m} - G\mathbf{z}$ does not lie in the same component as $\mathbf{u} + \mathbf{m}$. This implies that $\omega \cdot \mathbf{v} > \omega \cdot (\mathbf{u} + \mathbf{m})$ for all \mathbf{v} in the same component as $\mathbf{u} + \mathbf{m}$. By Lemma 3.10, $\mathbf{x}^{\mathbf{u}+\mathbf{m}}$ is a standard monomial of $\text{in}_{\omega}(J_{\mathcal{B}})$. Since $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$ and \mathbf{m} is an arbitrary vector with support contained in σ , we conclude that $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is an admissible pair of $\text{in}_{\omega}(J_{\mathcal{B}})$.

To show that $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair, we need to argue that the monomials of this pair are not properly contained in any other standard pair $(\mathbf{x}^{\mathbf{u}'}, \sigma')$ of $\text{in}_{\omega}(J_{\mathcal{B}})$. Suppose there is such a standard pair. We first argue that $\sigma = \sigma'$. By (2), if $i \notin \sigma$ then there exists some non-zero lattice point \mathbf{z} in $Q_{\mathbf{u},\omega}^{\overline{\sigma \cup \{i\}}}$ such that

$$\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}} \in (J_{\mathcal{B}} : \mathbf{x}_{\sigma \cup \{i\}}^{\infty}).$$

This implies that there exists some $p \in \mathbb{N}$ and $\mathbf{m} \in \mathbb{N}^n$ with support in σ such that $x_i^p \mathbf{x}^{\mathbf{m}}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) \in J_{\mathcal{B}}$. Since $(-\omega G)\mathbf{z} < 0$, $x_i^p \mathbf{x}^{\mathbf{m}}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z})$ is the leading term of $x_i^p \mathbf{x}^{\mathbf{m}}(\mathbf{x}^{\mathbf{u}+\mathbf{m}_z} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_z-G\mathbf{z}}) \in J_{\mathcal{B}}$ and hence is in $\text{in}_{\omega}(J_{\mathcal{B}})$. This construction shows that not all monomials of the form $\mathbf{x}^{\mathbf{u}}\mathbf{x}^{\mathbf{q}}$ where the support of \mathbf{q} is contained in $\sigma \cup \{i\}$ are standard monomials of $\text{in}_{\omega}(J_{\mathcal{B}})$ and hence $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is not contained in

any admissible pair $(\mathbf{x}^{\mathbf{u}}, \sigma')$ with $\sigma \subsetneq \sigma'$. To finish the argument, suppose $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is contained in a standard pair of form $(\mathbf{x}^{\mathbf{u}'}, \sigma)$. Then $\mathbf{u} = \mathbf{m} + \mathbf{u}'$ for some \mathbf{m} whose support is contained in σ . However, because $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair, the support of \mathbf{u} must also be disjoint from σ . Thus $\mathbf{m} = \mathbf{0}$ and so $\mathbf{u} = \mathbf{u}'$. \square

We now apply Theorems 3.22 and 3.21 to study the difference between the two monomial ideals $\text{in}_{\omega}(I_{\mathcal{A}})$ and $\text{in}_{\omega}(J_{\mathcal{B}})$. This difference will be the key to our study of the associated primes of $J_{\mathcal{B}}$ itself in Section 3.4.

Definition 3.23. A $J_{\mathcal{B}}$ -specific standard pair (JSP) is a standard pair of $\text{in}_{\omega}(J_{\mathcal{B}})$ that is not also a standard pair of $\text{in}_{\omega}(I_{\mathcal{A}})$.

Corollary 3.24. Assume $\mathbf{u} \in \mathbb{N}^n$ and $\sigma \in \Delta_{\omega}$ such that $\text{supp}(\mathbf{u}) \cap \sigma = \emptyset$. Then $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a JSP if and only if the two conditions of Theorem 3.22 hold and there exists at least one non-zero lattice point $\mathbf{z} \in Q_{\mathbf{u}, \omega}^{\bar{\sigma}}$.

Proof: If the two conditions of Theorem 3.22 hold then $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_{\omega}(J_{\mathcal{B}})$ and if there is a non-zero lattice point $\mathbf{z} \in Q_{\mathbf{u}, \omega}^{\bar{\sigma}}$, then by Theorem 3.21, $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is not a standard pair of $\text{in}_{\omega}(I_{\mathcal{A}})$. Thus it is a JSP. Conversely, if $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a JSP then the two conditions of Theorem 3.22 hold. Suppose there is no nonzero lattice point $\mathbf{z} \in Q_{\mathbf{u}, \omega}^{\bar{\sigma}}$. Then condition (1) of Theorem 3.21 is true. But since condition (2) of Theorem 3.22 holds for this JSP, there is a non-zero lattice point in $Q_{\mathbf{u}, \omega}^{\overline{\sigma \cup \{i\}}}$ for each $i \notin \sigma$, which is condition (2) of Theorem 3.21. This implies that $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_{\omega}(I_{\mathcal{A}})$, contradicting that it is a JSP. \square

Example 3.25. Consider the matrix A and weight vector ω given below:

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 3 & 4 & 5 & 6 \\ 0 & 0 & 7 & 8 & 9 \end{pmatrix}, \quad \omega = (1000, 100, 10, 1, 0).$$

Let \mathcal{B} be the set of circuits of \mathcal{A} . The matrix

$$G = \begin{pmatrix} 2 & 4 \\ -1 & -2 \\ -5 & -9 \\ 1 & 0 \\ 3 & 7 \end{pmatrix}$$

is a Gale dual of A and $\omega G = (1851, 3710)$. We have

$$J_{\mathcal{B}} = \langle b^2c^9 - a^4e^7, bd^9 - a^2e^8, bc^8 - a^2d^7, ce - d^2 \rangle$$

and its initial ideal

$$\text{in}_{\omega}(J_{\mathcal{B}}) = \langle a^2d^7, ce, a^4d^6e^4, a^4d^4e^5, a^2d^6e^5, a^4d^2e^6, a^2d^4e^6, a^4e^7, a^2d^2e^7, a^2e^8 \rangle$$

has 58 standard pairs. These ideals and standard pairs were computed using Macaulay 2 [27]. Consider the standard pair $(d^4e^3, \{1, 2\})$ for which $\mathbf{u} = (0, 0, 0, 4, 3)$ and $\sigma = \{1, 2\}$. The monomial d^4e^3 is a standard monomial for the toric initial ideal $\text{in}_{\omega}(I_{\mathcal{A}})$ as well and $Q_{\mathbf{u}, \omega} \cap \mathbb{Z}^2 = \{0\}$. However, the polytope

$$Q_{\mathbf{u}, \omega}^{\sigma} = \left\{ \mathbf{z} \in \mathbb{Z}^2 : \begin{pmatrix} -5 & -9 \\ 1 & 0 \\ 3 & 7 \end{pmatrix} \mathbf{z} \leq \begin{pmatrix} 0 \\ 4 \\ 3 \end{pmatrix}, 1851z_1 + 3710z_2 \geq 0 \right\}$$

contains two more lattice points: $(1, 0)$ and $(3, -1)$. Thus $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is not a standard pair of $\text{in}_{\omega}(I_{\mathcal{A}})$, so it is a JSP. Both points have $\mathbf{m}_{\mathbf{z}} = (2, 0, 0, 0, 0)$. For $(1, 0)$, $\mathbf{x}^{\mathbf{u}+\mathbf{m}_{\mathbf{z}}} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_{\mathbf{z}}-G\mathbf{z}} = a^2d^4e^3 - bc^5d^3$ is not in $(J_{\mathcal{B}} : (ab)^{\infty})$ but lies in $(J_{\mathcal{B}} : (aby)^{\infty})$ for each $y \in \{c, d, e\}$. Similarly, for $(3, -1)$, $\mathbf{x}^{\mathbf{u}+\mathbf{m}_{\mathbf{z}}} - \mathbf{x}^{\mathbf{u}+\mathbf{m}_{\mathbf{z}}-G\mathbf{z}} = a^2d^4e^3 - bc^6de$ does not lie in $(J_{\mathcal{B}} : (ab)^{\infty})$ but lies in $(J_{\mathcal{B}} : (aby)^{\infty})$ for each $y \in \{c, d, e\}$.

We now use JSPs to give a precise description of the set $\mathcal{H} := \{\mathbf{b} \in \mathbb{N}^{\mathcal{A}} : \mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b}) > 1\}$. This description gives a new proof of the following theorem alluded to in Section 3.2 (cf. Figure 3.1). The theorem also follows from [15, Corollary 5.3].

Theorem 3.26. *For all $\mathbf{b} \in \mathbb{N}\mathcal{A}$ sufficiently far from the boundary of $\text{cone}(\mathcal{A})$, $\mathbf{H}_{J_{\mathcal{B}}}(\mathbf{b}) = 1$ and hence the graphs $\mathcal{F}^{\mathcal{B}}(\mathbf{b})$ and $\mathcal{F}_{\omega}^{\mathcal{B}}(\mathbf{b})$ are connected.*

Recall that \mathbf{b} lies in \mathcal{H} if and only if for a generic ω , $\text{in}_{\omega}(J_{\mathcal{B}})$ has more than one standard monomial of degree \mathbf{b} . That is, $\mathcal{H} = \{\mathbf{A}\mathbf{u} : \mathbf{x}^{\mathbf{u}} \in \text{in}_{\omega}(I_{\mathcal{A}}) \setminus \text{in}_{\omega}(J_{\mathcal{B}})\}$. Since all standard monomials of degree \mathbf{b} other than the toric standard monomial lie on JSPs of $\text{in}_{\omega}(J_{\mathcal{B}})$, it follows that \mathcal{H} is contained in the union of the images in $\mathbb{N}\mathcal{A}$ of the JSPs of $\text{in}_{\omega}(J_{\mathcal{B}})$ under the map $\pi : \mathbb{N}^n \rightarrow \mathbb{N}\mathcal{A}$, $\mathbf{u} \mapsto \mathbf{A}\mathbf{u}$.

Lemma 3.27. *If $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a JSP of $\text{in}_{\omega}(J_{\mathcal{B}})$, then the set $\{\mathbf{a}_i : i \in \sigma\}$ is contained in a facet of $\text{cone}(\mathcal{A})$.*

Proof: Since $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a standard pair of $\text{in}_{\omega}(J_{\mathcal{B}})$, by Corollary 3.15, σ is a face of the triangulation Δ_{ω} . Suppose $\text{cone}(\mathcal{A}_{\sigma})$ intersects the interior of $\text{cone}(\mathcal{A})$. Choose a monomial \mathbf{x}^{α} on the JSP $(\mathbf{x}^{\mathbf{u}}, \sigma)$ such that $\mathbf{x}^{\alpha} \in \text{in}_{\omega}(I_{\mathcal{A}})$. Let \mathbf{x}^{β} be the standard monomial of $\text{in}_{\omega}(I_{\mathcal{A}})$ of degree $\mathbf{A}\alpha$. Then $\mathbf{x}^{\alpha} - \mathbf{x}^{\beta} \in I_{\mathcal{A}}$ with leading term \mathbf{x}^{α} . Since $\mathbf{x}_{\sigma}^{\mathbf{m}}\mathbf{x}^{\alpha} \notin \text{in}_{\omega}(J_{\mathcal{B}})$ for any \mathbf{m} , the binomial $\mathbf{x}_{\sigma}^{\mathbf{m}}(\mathbf{x}^{\alpha} - \mathbf{x}^{\beta}) \notin J_{\mathcal{B}}$ for any \mathbf{m} since its leading term $\mathbf{x}_{\sigma}^{\mathbf{m}}\mathbf{x}^{\alpha}$ would then be in $\text{in}_{\omega}(J_{\mathcal{B}})$. This implies that $(J_{\mathcal{B}} : \mathbf{x}_{\sigma}^{\infty}) \neq I_{\mathcal{A}}$. On the other hand, every embedded prime of $J_{\mathcal{B}}$ is of the form $P_{\tau} + I_{\mathcal{A}}$ where τ indexes some proper face of $\text{cone}(\mathcal{A})$ (see Proposition 3.30). The monomial \mathbf{x}_{σ} lies in each of these embedded primes since σ is not contained in any proper face of $\text{cone}(\mathcal{A})$. This implies that for \mathbf{m} large enough, $\mathbf{x}_{\sigma}^{\mathbf{m}}$ lies in every primary component of $J_{\mathcal{B}}$ except $I_{\mathcal{A}}$, which in turn implies that $(J_{\mathcal{B}} : \mathbf{x}_{\sigma}^{\infty}) = I_{\mathcal{A}}$, a contradiction. \square

Proof of Theorem 3.26: By Lemma 3.27, if $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a JSP of $\text{in}_{\omega}(J_{\mathcal{B}})$, then $\mathbf{A}\mathbf{u} + \mathbb{N}\mathcal{A}_{\sigma}$, its image under π in $\mathbb{N}\mathcal{A}$, is contained in a hyperplane parallel to a facet of $\text{cone}(\mathcal{A})$. Since there are finitely many JSPs of $\text{in}_{\omega}(J_{\mathcal{B}})$, \mathcal{H} is contained in finitely many hyperplanes parallel to the facets of $\text{cone}(\mathcal{A})$. This implies that the maximum distance of a point in \mathcal{H} from the boundary of $\text{cone}(\mathcal{A})$ is bounded which proves the theorem. \square

We conclude this section with an answer to Problem 3.4.

Definition 3.28. A polytope $Q_{\mathbf{u},\omega}^{\bar{\sigma}}$ corresponding to a JSP $(\mathbf{x}^{\mathbf{u}}, \sigma)$ of $\text{in}_{\omega}(J_{\mathcal{B}})$ is called a *JSP polytope* of \mathcal{A} .

Note that JSP polytopes can be defined independently of standard pairs by the conditions of Corollary 3.24.

Theorem 3.29. *The following are equivalent.*

1. *The ideals $I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ are not equal.*
2. *There is a generic $\omega \in \mathbb{R}^n$ for which \mathcal{A} has a JSP polytope.*
3. *For every generic $\omega \in \mathbb{R}^n$, \mathcal{A} has a JSP polytope.*

Proof: The ideal $I_{\mathcal{A}} = J_{\mathcal{B}}$ if and only if for any generic $\omega \in \mathbb{R}^n$, $\text{in}_{\omega}(I_{\mathcal{A}}) = \text{in}_{\omega}(J_{\mathcal{B}})$ which is if and only if $\text{in}_{\omega}(J_{\mathcal{B}})$ has no JSPs. \square

3.4 Associated Primes of the Circuit Ideal

In this section, we show how the associated primes of $J_{\mathcal{B}}$ relate to the JSP polytopes of its initial ideals discussed in Section 3.3. Again, we assume throughout this section that $J_{\mathcal{B}}$ is a basis ideal of \mathcal{A} and that $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$. Recall that a face F of $\text{cone}(\mathcal{A})$ is recorded as the set $\sigma := \{j : \mathbf{a}_j \in F\} \subseteq [n]$.

Proposition 3.30. [20, Proposition 7.8] *All associated primes of $J_{\mathcal{B}}$ are of the form $P_{\sigma} + I_{\mathcal{A}}$ for some face σ of $\text{cone}(\mathcal{A})$. The toric ideal $I_{\mathcal{A}} = P_{[n]} + I_{\mathcal{A}}$ is the unique minimal prime of $J_{\mathcal{B}}$. However, not all proper faces of $\text{cone}(\mathcal{A})$ need index an associated prime of $J_{\mathcal{B}}$.*

Definition 3.31. [4] Let I be any ideal in $k[\mathbf{x}]$ and let P be an ideal that contains I . Then P is an associated prime of I if P is prime and there exists some $f \in k[\mathbf{x}]$ such that $(I : f) = P$. We call f a *witness* for P .

Using Proposition 3.30, we can now state the main results of this section. We say that τ is the *type* of a standard pair of the form (\cdot, τ) .

Theorem 3.32. *Let τ be any (possibly empty) proper face of $\text{cone}(\mathcal{A})$ and ω be a generic weight vector. If $P_\tau + I_{\mathcal{A}}$ is associated to $J_{\mathcal{B}}$ (and hence embedded), then there exists $\gamma \subseteq \tau$ and a $J_{\mathcal{B}}$ -specific standard pair of $\text{in}_\omega(J_{\mathcal{B}})$ of type γ such that*

1. *if σ is a face of $\text{cone}(\mathcal{A})$ properly contained in τ , then γ is not contained in σ , and*
2. *$|\gamma| = \dim(\text{cone}(\mathcal{A}_\tau))$.*

Furthermore, there is a witness for the prime $P_\tau + I_{\mathcal{A}}$ whose leading term with respect to ω lies on such a JSP.

We also prove a partial converse.

Theorem 3.33. *For a generic ω , if $\text{in}_\omega(J_{\mathcal{B}})$ has a JSP of type γ , then $J_{\mathcal{B}}$ has an embedded prime $P_\sigma + I_{\mathcal{A}}$ for some face σ of $\text{cone}(\mathcal{A})$ such that $\sigma \supseteq \gamma$.*

Before proving the theorems, we consider a few implications. We say that a face τ of $\text{cone}(\mathcal{A})$ is *simplicial* if $|\tau| = \dim \text{cone}(\mathcal{A}_\tau)$. If τ is a simplicial face of $\text{cone}(\mathcal{A})$, then no binomial in $I_{\mathcal{A}}$ is supported entirely on τ , so $P_\tau + I_{\mathcal{A}}$ is just the monomial prime $P_\tau = \langle x_i : i \notin \tau \rangle$. Then Theorem 3.32 specializes as follows.

Corollary 3.34. *If $P_\tau + I_{\mathcal{A}}$ is an embedded prime of $J_{\mathcal{B}}$ and τ is a simplicial face of $\text{cone}(\mathcal{A})$, then for every generic ω , $\text{in}_\omega(J_{\mathcal{B}})$ has a JSP of type τ .*

The situation is more complicated when non-simplicial faces of $\text{cone}(\mathcal{A})$ index embedded primes. In particular, Theorem 3.32 does not specify a particular $\gamma \subseteq [n]$ such that every monomial initial ideal of $J_{\mathcal{B}}$ must have a JSP of type γ .




Figure 3.2: The point configuration of Example 3.35.

Example 3.35. Let \mathcal{B} be the set of circuits of the matrix

$$A = \begin{pmatrix} 5 & 0 & 0 & 2 & 1 & 0 \\ 0 & 5 & 0 & 1 & 4 & 2 \\ 0 & 0 & 5 & 2 & 0 & 3 \end{pmatrix}.$$

The configuration \mathcal{A} labeled $1, \dots, 6$ in Figure 3.2 spans the cone over a triangle in \mathbb{R}^3 , so by Proposition 3.30, there are seven possible embedded primes of $J_{\mathcal{B}}$ corresponding to the seven proper faces of $\text{cone}(\mathcal{A})$. All seven of these primes are indeed associated to $J_{\mathcal{B}}$. The two non-simplicial 2-dimensional faces, $\{1, 2, 5\}$ and $\{2, 3, 6\}$, index the primes $P_{\{1,2,5\}} + I_{\mathcal{A}} = \langle f, d, c, ab^4 - e^5 \rangle$ and $P_{\{2,3,6\}} + I_{\mathcal{A}} = \langle e, d, a, b^2c^3 - f^5 \rangle$. The third 2-face $\{1, 3\}$ is simplicial and indexes the prime $P_{\{1,3\}} = \langle b, d, e, f \rangle$. The remaining four primes $P_{\{1,2\}} = \langle c, d, e, f \rangle$, $P_{\{1,3\}} = \langle b, d, e, f \rangle$, $P_{\{2,3\}} = \langle a, d, e, f \rangle$, and $P_{\{\emptyset\}} = \langle a, b, c, d, e, f \rangle$ correspond to the three rays of $\text{cone}(\mathcal{A})$ and to the apex, all of which are trivially simplicial.

By Corollary 3.34, each initial ideal $\text{in}_{\omega}(J_{\mathcal{B}})$ has JSPs of types $\{1, 3\}$, $\{3\}$, $\{2\}$, $\{1\}$, and \emptyset corresponding to the five simplicial faces of $\text{cone}(\mathcal{A})$. Since $P_{\{1,2,5\}} + I_{\mathcal{A}}$ is associated, Theorem 3.32 requires that $\text{in}_{\omega}(J_{\mathcal{B}})$ has a JSP of type $\{1, 2\}$, $\{1, 5\}$, or $\{2, 5\}$. Similarly, because of $P_{\{2,3,6\}} + I_{\mathcal{A}}$, there must be a JSP of type $\{2, 3\}$, $\{2, 6\}$, or $\{3, 6\}$. We list the types of JSPs that appear for two term orders.

- For lexicographic order with $f \succ e \succ \dots \succ a$, $\text{in}_{\succ}(J_{\mathcal{B}})$ has the following types of JSPs: $\{1, 3\}$, $\{3\}$, $\{2\}$, $\{1\}$, \emptyset , $\{1, 2\}$, $\{2, 3\}$.

- For reverse lexicographic order with $a \succ b \succ \dots \succ f$, $\text{in}_\succ(J_{\mathcal{B}})$ has the following types of JSPs: $\{1, 3\}, \{3\}, \{2\}, \{1\}, \emptyset, \{6\}, \{1, 5\}, \{3, 6\}, \{2, 6\}, \{2, 5\}$.

We now prove Theorem 3.32. The idea is to find a witness for the embedded prime $P_\tau + I_{\mathcal{A}}$, compute its normal form with respect to the reduced Gröbner basis $G_\omega(J_{\mathcal{B}})$, and show that the result is a witness whose ω -initial term lies on a JSP satisfying all of the desired properties.

Lemma 3.36. *If f is a witness for an embedded prime $P_\tau + I_{\mathcal{A}}$ of $J_{\mathcal{B}}$, then the following hold.*

1. *The witness f is in the toric ideal $I_{\mathcal{A}}$.*
2. *For any $g \in J_{\mathcal{B}}$, $f + g$ is also a witness for $P_\tau + I_{\mathcal{A}}$. In particular, the normal form of f with respect to $G_\omega(J_{\mathcal{B}})$ is a witness.*
3. *If $\mathbf{x}^{\mathbf{m}}$ is a monomial with $\text{supp}(\mathbf{m}) \subseteq \tau$, then $\mathbf{x}^{\mathbf{m}}f$ is also a witness.*

Proof:

1. Since τ is a proper face of $\text{cone}(\mathcal{A})$, there is some variable $x_i \in P_\tau + I_{\mathcal{A}}$, so $x_i f \in J_{\mathcal{B}} \subset I_{\mathcal{A}}$. Since $I_{\mathcal{A}}$ is a prime ideal without monomials, $f \in I_{\mathcal{A}}$.
2. Since $g \in J_{\mathcal{B}}$, so is pg for any polynomial $p \in k[\mathbf{x}]$, and thus $p(f + g) \in J_{\mathcal{B}} \Leftrightarrow pf \in J_{\mathcal{B}}$. Thus $(J_{\mathcal{B}} : f + g) = (J_{\mathcal{B}} : f) = P_\tau + I_{\mathcal{A}}$.
3. If $h \in P_\tau + I_{\mathcal{A}}$, then $h(\mathbf{x}^{\mathbf{m}}f) = (\mathbf{x}^{\mathbf{m}}h)f$ is in $J_{\mathcal{B}}$ by the assumption that f is a witness. On the other hand, if $h \notin P_\tau + I_{\mathcal{A}}$, then neither is $\mathbf{x}^{\mathbf{m}}h$ because $\text{supp}(\mathbf{m}) \subseteq \tau$ and $P_\tau + I_{\mathcal{A}}$ is prime. Thus $\mathbf{x}^{\mathbf{m}}h \notin P_\tau + I_{\mathcal{A}} = (J_{\mathcal{B}} : f)$, so $h \notin (J_{\mathcal{B}} : \mathbf{x}^{\mathbf{m}}f)$. Thus $(J_{\mathcal{B}} : \mathbf{x}^{\mathbf{m}}f) = (J_{\mathcal{B}} : f) = P_\tau + I_{\mathcal{A}}$ as claimed.

□

Lemma 3.37. *If f is a witness for an embedded prime $P_\tau + I_{\mathcal{A}}$ of $J_{\mathcal{B}}$ and \bar{f} is the normal form of f with respect to $G_\omega(J_{\mathcal{B}})$, then $\text{in}_\omega(\bar{f})$ lies on a JSP (\cdot, γ) of $\text{in}_\omega(J_{\mathcal{B}})$ with $\gamma \subseteq \tau$.*

Proof: By Lemma 3.36, \bar{f} is also a witness for $P_\tau + I_{\mathcal{A}}$ and $\bar{f} \in I_{\mathcal{A}}$. This implies that $\text{in}_\omega(\bar{f}) \in \text{in}_\omega(I_{\mathcal{A}}) \setminus \text{in}_\omega(J_{\mathcal{B}})$, so $\text{in}_\omega(\bar{f})$ must lie on some JSP (\cdot, γ) of $\text{in}_\omega(J_{\mathcal{B}})$. Since $x_i \bar{f} \in J_{\mathcal{B}}$ whenever $i \notin \tau$ because \bar{f} witnesses $P_\tau + I_{\mathcal{A}}$, it follows that $x_i \text{in}_\omega(\bar{f}) \in \text{in}_\omega(J_{\mathcal{B}})$ for $i \notin \tau$, so $\gamma \subseteq \tau$. \square

Proof of Theorem 3.32: Suppose $P_\tau + I_{\mathcal{A}}$ is an embedded prime of $J_{\mathcal{B}}$ and $e := \dim \text{cone}(\mathcal{A}_\tau)$. We first claim the following: there is a constant C such that for all sufficiently large N , $P_\tau + I_{\mathcal{A}}$ has at least N^e witnesses whose normal forms with respect to $G_\omega(J_{\mathcal{B}})$ have distinct leading terms, and each such leading term $\mathbf{x}^{\mathbf{p}}$ has the property that every component of \mathbf{p} is bounded above by CN .

Suppose the claim is true. By Lemma 3.37, each such monomial $\mathbf{x}^{\mathbf{p}}$ must lie on a JSP of type γ with $\gamma \subseteq \tau$. Each such standard pair contains at most $C^{|\gamma|}(N+1)^{|\gamma|}$ monomials $\mathbf{x}^{\mathbf{p}}$ such that p_i is bounded above by CN . Since there are only finitely many standard pairs for $\text{in}_\omega(J_{\mathcal{B}})$, all the standard pairs of type γ with $|\gamma| < e$ together cover only at most $\mathcal{O}(N^{e-1})$ of the monomials which is not enough to contain the N^e leading terms $\mathbf{x}^{\mathbf{p}}$. Thus some of these leading terms must lie on standard pairs (\cdot, γ) with $|\gamma| \geq e$. Since by Corollary 3.15, each γ is a face of the triangulation Δ_ω of \mathcal{A} , this is only possible if $|\gamma| = e$ and γ is not contained in any face σ of $\text{cone}(\mathcal{A})$ whose dimension is less than e . These are exactly the types of standard pairs specified by Theorem 3.32.

Now we prove the claim. Since $P_\tau + I_{\mathcal{A}}$ and $J_{\mathcal{B}}$ are both \mathcal{A} -homogeneous, there exists an \mathcal{A} -homogeneous witness f for $P_\tau + I_{\mathcal{A}}$. Set $\mathbf{x}^{\mathbf{u}} := \text{in}_\omega(f)$. Since $e = \dim \text{cone}(\mathcal{A}_\tau)$, we can find an e -subset α of τ such that the columns of A indexed by α are linearly independent. Thus if $\mathbf{m}_1 \neq \mathbf{m}_2$ are supported only on α , then $A\mathbf{m}_1 \neq A\mathbf{m}_2$.

Consider all polynomials of the form $\mathbf{x}^{\mathbf{m}}f$ where $0 \leq m_i < N$ for $i \in \alpha$ and $m_i = 0$ for $i \notin \alpha$. Such a polynomial is \mathcal{A} -homogeneous of \mathcal{A} -degree $A\mathbf{m} + A\mathbf{u}$ and so is its normal form with respect to $G_\omega(J_{\mathcal{B}})$ since $J_{\mathcal{B}}$ is an \mathcal{A} -homogeneous ideal. Thus the normal forms of these N^e polynomials are all \mathcal{A} -homogeneous of different degrees, so in particular they all have distinct leading terms. Furthermore, by parts (2) and (3) of Lemma 3.36, each such normal form is a witness for $P_\tau + I_{\mathcal{A}}$.

It remains to establish that if $\mathbf{x}^{\mathbf{p}}$ is the leading term of one of the normal forms, then each component of \mathbf{p} is bounded by a constant multiple of N . Let \mathbf{a} be a strictly positive vector in the rowspan of A . Such a vector exists since $\mathcal{L}_{\mathcal{A}} \cap \mathbb{N}^n = \{0\}$. By scaling, we can assume that the minimum component of \mathbf{a} is 1. Let R be its maximum component. Since $A\mathbf{p} = A(\mathbf{u} + \mathbf{m})$, it follows that $\mathbf{a} \cdot \mathbf{p} = \mathbf{a} \cdot (\mathbf{u} + \mathbf{m})$. Then $\|\mathbf{p}\|_1 = \sum_{i=1}^n p_i \leq \sum_{i=1}^n a_i p_i =$

$$\sum_{i=1}^n a_i(u_i + m_i) \leq R \sum_{i=1}^n (u_i + m_i) = R \left(\sum_{i=1}^n u_i + \sum_{i=1}^n m_i \right) < R(\|\mathbf{u}\|_1 + nN).$$

It follows that for any i , we have

$$p_i \leq \|\mathbf{p}\|_1 < RnN + R\|\mathbf{u}\|_1$$

which is a bound of the desired form. \square

We now prove Theorem 3.33. Recall the following algebraic fact.

Lemma 3.38. *If I is an ideal in $k[\mathbf{x}]$ and g is any polynomial, then the associated primes of $(I : g^\infty)$ are exactly the associated primes of I that do not contain g .*

Proposition 3.39. *Recall that $\mathbf{x}_\tau = \prod_{i \in \tau} x_i$. The associated primes of $(J_{\mathcal{B}} : \mathbf{x}_\tau^\infty)$ are exactly the associated primes $P_\sigma + I_{\mathcal{A}}$ of $J_{\mathcal{B}}$ that satisfy $\sigma \supseteq \tau$.*

Proof: We get $\mathbf{x}_\tau \in P_\sigma + I_{\mathcal{A}}$ if and only if some x_i with $i \in \tau$ lies in $P_\sigma + I_{\mathcal{A}}$, which occurs if and only if τ is not contained in σ . Now apply Lemma 3.38. \square

Proof of Theorem 3.33: Suppose (\mathbf{x}^u, γ) is a JSP of $\text{in}_\omega(J_{\mathcal{B}})$. Choose $f \in I_{\mathcal{A}}$ such that $\text{in}_\omega(f) = \mathbf{x}^u \mathbf{x}_\gamma^m$ for some $m \geq 0$. This is possible since every JSP of $\text{in}_\omega(J_{\mathcal{B}})$ contains non-standard monomials of $\text{in}_\omega(I_{\mathcal{A}})$. Since no monomial of the form $\text{in}_\omega(f) \cdot \mathbf{x}_\gamma^*$ lies in $\text{in}_\omega(J_{\mathcal{B}})$, no polynomial of the form $f \cdot \mathbf{x}_\gamma^*$ lies in $J_{\mathcal{B}}$. This implies that $(J_{\mathcal{B}} : \mathbf{x}_\gamma^\infty)$ does not contain f and is hence not equal to $I_{\mathcal{A}}$. However, since $(J_{\mathcal{B}} : \mathbf{x}_\gamma^\infty) \subsetneq I_{\mathcal{A}}$, $(J_{\mathcal{B}} : \mathbf{x}_\gamma^\infty)$ must have an embedded prime. This prime is also embedded in $J_{\mathcal{B}}$ by Proposition 3.39, and it has the form $P_\sigma + I_{\mathcal{A}}$ for some $\sigma \supseteq \gamma$. \square

Theorem 3.33 is only a partial converse to Theorem 3.32. It is not true for a given weight vector ω that the existence of a JSP (\mathbf{x}^u, γ) of $\text{in}_\omega(J_{\mathcal{B}})$ satisfying the conditions of Theorem 3.32 with respect to some proper face τ of $\text{cone}(\mathcal{A})$ implies that $P_\tau + I_{\mathcal{A}}$ is associated to $J_{\mathcal{B}}$.

Example 3.40. Let \mathcal{B} be the set of circuits of the matrix

$$A = \begin{pmatrix} 1 & 3 & 2 & 4 \\ 1 & 4 & 5 & 2 \end{pmatrix}.$$

The values of the \mathcal{A} -graded Hilbert function of this \mathcal{A} are shown in Figure 3.1. The proper faces of $\text{cone}(\mathcal{A})$ are $\{3\}$, $\{4\}$, and \emptyset . Only the first two index associated primes of $J_{\mathcal{B}}$. However, if we take ω to represent the lexicographic term order with $a \succ b \succ c \succ d$, there are five JSPs of $\text{in}_\omega(J_{\mathcal{B}})$ of type \emptyset . On the other hand, if ω represents the \mathcal{A} -graded reverse lexicographic order with $a \succ b \succ c \succ d$, then there are no JSPs of type \emptyset .

Question 3.41. *If τ is a face of $\text{cone}(\mathcal{A})$ such that for every generic ω there is a JSP of $\text{in}_\omega(J_{\mathcal{B}})$ satisfying the conditions of Theorem 3.32 with respect to τ , then is $P_\tau + I_{\mathcal{A}}$ necessarily associated to $J_{\mathcal{B}}$?*

We conclude this section with an application of Theorem 3.32 to the specific case of circuit ideals of normal configurations. Let \mathcal{B} be the set of circuits of a configuration \mathcal{A} satisfying $\mathbb{Z}\mathcal{A} = \mathbb{Z}^d$ and whose vectors comprise the lattice points in a lattice

polytope R . Further assume that the first row of A is $(1, \dots, 1)$, so R is at height 1. The polytope R defines a projective toric variety $X_{\mathcal{A}}$ and the faces $\{\tau\}$ of R (which are in bijection with the faces of $\text{cone}(\mathcal{A})$) index a canonical collection of affine charts $\{U_{\tau}\}$ covering $X_{\mathcal{A}}$ [25]. We investigate how smoothness of U_{τ} determines whether $P_{\tau} + I_{\mathcal{A}}$ is associated to $J_{\mathcal{C}_{\mathcal{A}}} = J_{\mathcal{B}}$.

- Definition 3.42.**
1. Let K be a convex rational polyhedral cone in \mathbb{R}^t that does not contain a line. We say that K is *smooth* if it is generated by primitive vectors that form part of a basis for \mathbb{Z}^t .
 2. Let K_F denote the inner normal cone of a face F of a polytope Q . The face F is *smooth* if the restriction of K_F to the linear span of Q is smooth.

- Remark 3.43.**
1. If v is a smooth vertex of a polytope Q then there are exactly $\dim Q$ edges of Q incident to v . Further, the cone dual to K_v is also smooth [22, Theorem 2.10, Chapter V]. Note that this dual cone is the tangent cone of Q at v and contains Q .
 2. A face F of a polytope Q is smooth if and only if the affine toric variety U_F is smooth [25].

Theorem 3.44. *Let \mathcal{A} and R be as above. If \mathbf{a}_n is a smooth vertex of R , then $P_{\{n\}} (= P_{\{n\}} + I_{\mathcal{A}})$ is not an associated prime of $J_{\mathcal{C}_{\mathcal{A}}}$.*

Proof: Suppose $P_{\{n\}}$ is associated. Since $\{n\}$ is a simplicial face of $\text{cone}(\mathcal{A})$, by Corollary 3.34, every monomial initial ideal $\text{in}_{\omega}(J_{\mathcal{C}_{\mathcal{A}}})$ has a JSP of the form $(\mathbf{x}^{\mathbf{u}}, \{n\})$. In particular, let ω represent an elimination order with x_n most expensive. That is, any monomial containing x_n is more expensive than any monomial that does not. We may assume that each of $\mathbf{a}_1, \dots, \mathbf{a}_{d-1} \in \mathcal{A}$ is the first lattice point from \mathbf{a}_n along one of the $d - 1$ edges incident to \mathbf{a}_n . Let $\mathbf{y}_i := \mathbf{a}_i - \mathbf{a}_n$ for $i = 1, \dots, d - 1$.

Since \mathbf{a}_n is smooth and R is contained in the tangent cone at \mathbf{a}_n , for each lattice point in R (i.e. column of A), there are unique $m_i \in \mathbb{N}$ such that $\mathbf{a}_j = \mathbf{a}_n + \sum_{i=1}^{d-1} m_i \mathbf{y}_i$. Rearranging terms, and setting $M = -1 + \sum_{i=1}^{d-1} m_i$, we get

$$\mathbf{a}_j + M\mathbf{a}_n = \sum_{i=1}^{d-1} m_i \mathbf{a}_i \quad (3.1)$$

with all coefficients nonnegative. If $j = n$, (3.1) reduces to $0 = 0$, and if $1 \leq j \leq d-1$, it reduces to $\mathbf{a}_j = \mathbf{a}_j$. But in the nontrivial case where $d-1 < j < n$, (3.1) is a *circuit* because $\mathbf{a}_1, \dots, \mathbf{a}_{d-1}, \mathbf{a}_n$ form a maximal linearly independent set. Thus $x_j x_n^M - \prod_{i=1}^{d-1} x_i^{m_i} \in J_{\mathcal{C}_A}$. By choice of ω , its leading term is $x_j x_n^M$. Since $(\mathbf{x}^u, \{n\})$ is a JSP, this term must not divide $x_n^N \mathbf{x}^u$ for any N . Thus $j \notin \text{supp}(\mathbf{u})$.

For N sufficiently large, $x_n^N \mathbf{x}^u \in \text{in}_\omega(I_A)$, so we can choose $\mathbf{x}^v \notin \text{in}_\omega(I_A)$ such that $x_n^N \mathbf{x}^u - \mathbf{x}^v \in I_A$. Since I_A is prime, factor out any common monomial to get $\mathbf{x}^{\tilde{u}} - \mathbf{x}^{\tilde{v}} \in I_A$ where \tilde{u} and \tilde{v} have disjoint supports. Since $\tilde{u} - \tilde{v} \in \mathcal{L}_A$, the convex hulls of $\{\mathbf{a}_i : i \in \text{supp}(\tilde{u})\}$ and $\{\mathbf{a}_i : i \in \text{supp}(\tilde{v})\}$ must intersect.

Since \mathbf{a}_n is smooth, we can assume by applying an invertible \mathbb{Z} -affine transformation that \mathbf{a}_n is the origin and \mathbf{a}_i is the i -th standard basis vector in \mathbb{Z}^{d-1} for $1 \leq i \leq d-1$. That is, $\mathbf{a}_n, \mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_{d-1}$ are the vertices of the standard simplex S in \mathbb{R}^{d-1} . Since $j \notin \text{supp}(\tilde{u})$ for any $d-1 < j < n$, $U := \text{conv}(\mathbf{a}_i : i \in \text{supp}(\tilde{u}))$ is a face of S . Since $\text{supp}(\tilde{v}) \cap \text{supp}(\tilde{u}) = \emptyset$ and S contains no lattice points except its vertices, $V := \text{conv}(\mathbf{a}_i : i \in \text{supp}(\tilde{v}))$ consists only of vertices of S outside U along with lattice points in $R \setminus S$. Now S and $\overline{R \setminus S}$ are both convex, so U and V could intersect only on the boundary of S . But since the vertices in U and those in $V \cap S$ form disjoint faces of S , there is no intersection on this boundary, contradicting that $U \cap V \neq \emptyset$. \square

Example 3.45.

1. Non-smooth vertices of R may or may not index associated primes of $J_{\mathcal{C}_A}$. For

the A below, the polytope R is a triangle in \mathbb{R}^3 .

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 3 & 4 & 5 & 5 \\ 0 & 1 & 1 & 1 & 2 \end{pmatrix}.$$

None of the three vertices $(0, 0)$, $(5, 1)$, and $(5, 2)$ of R are smooth. The vertices $(0, 0)$ and $(5, 2)$ both index associated primes, while the vertex $(5, 1)$ does not.

2. Smooth edges of R may index associated primes of J_{C_A} . Consider the matrix

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

whose columns are the lattice points in a rectangular prism R . All edges of R are smooth, but we compute that four of them index associated primes of J_{C_A} .

3.5 Fans of Toric and Circuit Ideals

The main goal of this section is to compare I_A and J_{C_A} via three polyhedral fans that can be associated to them. These results rely on Corollary 3.52 which states that for a generic ω , the top-dimensional components of $\text{in}_\omega(I_A)$ and $\text{in}_\omega(J_{C_A})$ are the same. We begin by providing a more general result.

Definition 3.46. [42, page 112] Let $J \subseteq k[\mathbf{x}]$ be an ideal and d be the Krull dimension of $k[\mathbf{x}]/J$. We define $\text{top}(J)$ to be the intersection of all primary components of J of dimension d .

Note that $\text{top}(J)$ is well-defined since the d -dimensional primary components of J are minimal and are hence unique.

In projective space the degree of an irreducible variety is defined as the number of points in its intersection with a complementary dimensional subspace in general position. Recall the usual generalization to ideals.

Definition 3.47. [50, Section 1] Let $P \subseteq k[\mathbf{x}]$ be a prime ideal. The *multiplicity* $\text{mult}(Q)$ of a P -primary ideal Q is the length of a maximal strictly increasing chain of P -primary ideals $Q \subset \dots \subset P$. Let $I \subseteq k[\mathbf{x}]$ be a homogeneous ideal in the total degree grading. The *degree* of I is defined as

$$\deg(I) = \sum_Q \text{mult}(Q) \deg(V(Q))$$

where the sum is taken over the d -dimensional primary components Q in a minimal primary decomposition of I and $V(Q)$ denotes the variety defined by Q .

The degree is also characterized as the normalized leading coefficient of the Hilbert polynomial of $k[\mathbf{x}]/I$ and thus does not change when going to initial ideals.

Lemma 3.48. *Let $J \subseteq I \subseteq k[\mathbf{x}]$ be ideals with the Krull dimension of $k[\mathbf{x}]/J$ being d . Any d -dimensional associated prime P of I is an associated prime of J .*

Proof: Clearly, $J \subseteq P$. We wish to show that P is a minimal prime of J . By Proposition 4.6 in [4] it suffices to prove that P is minimal with respect to inclusion among all prime ideals containing J . Suppose P' is a prime ideal with $J \subseteq P' \subseteq P$ then, since J and P have the same dimension, $P = P'$ and P is indeed minimal. \square

Proposition 3.49. *Let $I, J \subseteq k[\mathbf{x}]$ be homogeneous ideals in the total degree grading with the Krull dimensions of $k[\mathbf{x}]/I$ and $k[\mathbf{x}]/J$ both being d and with the degrees of I and J being equal. If $J \subseteq I$ then $\text{top}(I) = \text{top}(J)$.*

Proof: We consider minimal primary decompositions of I and J

$$I = Q_1 \cap \dots \cap Q_t \cap \dots \cap Q_r$$

$$J = Q'_1 \cap \dots \cap Q'_s$$

where Q_1, \dots, Q_t are the d -dimensional components in the decomposition of I . By Lemma 3.48 we may assume that $\sqrt{Q_i} = \sqrt{Q'_i} =: P_i$ for $i = 1, \dots, t$. Let $i \leq t$ and consider the primary components Q_i and Q'_i . As $J \subseteq I$ we have $Q'_i \subseteq Q_i$. By the definition of multiplicity, $\text{mult}(Q'_i) \geq \text{mult}(Q_i)$. The ideal J may have other d -dimensional components in its decomposition. Hence using Definition 3.47 we get

$$\deg(J) \geq \sum_{i=1}^t \text{mult}(Q'_i) \deg(V(Q_i)) \geq \sum_{i=1}^t \text{mult}(Q_i) \deg(V(Q_i)) = \deg(I).$$

Our assumption $\deg(J) = \deg(I)$ now implies that $\text{mult}(Q'_i) = \text{mult}(Q_i)$ for all $i \leq t$. Furthermore, we see that J cannot have more d -dimensional components. According to the definition of multiplicity the inclusion $Q'_i \subseteq Q_i$ cannot be strict. This proves that the d -dimensional components are the same in the two decompositions. As top is defined as their intersection, $\text{top}(I) = \text{top}(J)$. \square

Corollary 3.50. *If $J \subseteq I \subseteq k[\mathbf{x}]$ are homogeneous ideals in the total degree grading with $\text{top}(I) = \text{top}(J)$ then for $\omega \in \mathbb{R}^n$*

$$\text{top}(\text{in}_\omega(I)) = \text{top}(\text{in}_\omega(J)).$$

Proof: Clearly, $\text{in}_\omega(J) \subseteq \text{in}_\omega(I)$. It follows from the definition of degree and top that I and J have the same degree and dimension. So do their initial ideals. We now apply Proposition 3.49 to $\text{in}_\omega(J)$ and $\text{in}_\omega(I)$. \square

Corollary 3.51. *Let $J \subseteq k[\mathbf{x}]$ be a homogeneous ideal in the total degree grading. For $\omega \in \mathbb{R}^n$ we have*

$$\text{top}(\text{in}_\omega(\text{top}(J))) = \text{top}(\text{in}_\omega(J)).$$

Proof: Let $I = \text{top}(J)$ and apply Corollary 3.50. \square

Corollary 3.52. *If $I_{\mathcal{A}}$ is homogeneous in the total degree grading and $J_{\mathcal{B}}$ is a basis ideal with $\sqrt{J_{\mathcal{B}}} = I_{\mathcal{A}}$, then for $\omega \in \mathbb{R}^n$*

$$\text{top}(\text{in}_{\omega}(I_{\mathcal{A}})) = \text{top}(\text{in}_{\omega}(J_{\mathcal{B}})).$$

In particular,

$$\text{top}(\text{in}_{\omega}(I_{\mathcal{A}})) = \text{top}(\text{in}_{\omega}(J_{\mathcal{C}_{\mathcal{A}}})).$$

Proof: By [20, Theorem 7.6], the unique minimal primary component of $J_{\mathcal{B}}$ is $(J_{\mathcal{B}} : (x_1 \cdots x_n)^{\infty})$ which equals $I_{\mathcal{A}}$ by Proposition 3.3. Thus $\text{top}(J_{\mathcal{B}}) = I_{\mathcal{A}}$ and the claim follows immediately from Corollary 3.51. \square

Corollary 3.53. *If $J \subseteq I \subseteq k[\mathbf{x}]$ are homogeneous ideals in the total degree grading with $\text{top}(I) = \text{top}(J)$ then for generic $\omega \in \mathbb{R}^n$ the d -dimensional standard pairs of $\text{in}_{\omega}(I)$ and of $\text{in}_{\omega}(J)$ are the same.*

Proof: The claim follows from Corollary 3.50 if we can prove that any d -dimensional monomial ideal M has the same d -dimensional standard pairs as $\text{top}(M)$. Consider a minimal primary decomposition $M = Q_1 \cap \dots \cap Q_t \cap \dots \cap Q_r$ where Q_1, \dots, Q_t are the d -dimensional components. Now $\text{top}(M) = Q_1 \cap \dots \cap Q_t$. Without loss of generality we may assume that each Q_i is a monomial primary ideal and hence of the form $Q_i = \langle x_j^{(\mathbf{v}_i)_j} \rangle_{j \notin \sigma_i} + \langle \mathbf{x}^{\mathbf{w}} \rangle_{\mathbf{w} \in S_i}$, for some $\sigma_i \subseteq [n]$, $\mathbf{v}_i \in \mathbb{N}^n$ and a collection S_i of vectors in \mathbb{N}^n with support of size at least two and contained in $\overline{\sigma_i}$. Here $(\mathbf{v}_i)_j$ denotes the j th entry of the vector \mathbf{v}_i . The exponent vectors of monomials not in Q_i are unbounded exactly on the entries indexed by σ_i .

Any d -dimensional standard pair of $\text{top}(M)$ is clearly admissible for M . Furthermore, since $\dim(M) = d$ it is also a standard pair of M . Conversely, if $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is a d -dimensional standard pair of M then the monomials it represents are contained in $\overline{Q_1 \cap \dots \cap Q_r} = \overline{Q_1} \cup \dots \cup \overline{Q_r}$. As the exponent vector of such a monomial may be arbitrary large at the entries indexed by σ , for some i we must have $\sigma \subseteq \sigma_i$ with

$\mathbf{x}^{\mathbf{u}} \notin Q_i$. Since $|\sigma_k| \leq d$ for all $k = 1, \dots, r$, we get $|\sigma_i| = d$ and $\dim(Q_i) = d$. Hence $(\mathbf{x}^{\mathbf{u}}, \sigma)$ is admissible for $\text{top}(M)$ and since $\dim(\text{top}(M)) = d$ it is also standard. \square

In particular, the above Corollary can be applied to the circuit ideal and the toric ideal of a point configuration. A longer variant of our proof, not presented here, proves the following statement for generic ω .

Conjecture 3.54. *The equality in Corollary 3.52 holds for any $\omega \in \mathbb{R}^n$ even if the ideals are not homogeneous in the total degree grading.*

3.5.1 Polyhedral Fans of J_{C_A}

An ideal in $k[\mathbf{x}]$ gives rise to several natural equivalence relations on \mathbb{R}^n some of which give rise to polyhedral fans. In this final part, we compare various equivalence relations and fans for toric and circuit ideals.

Definition 3.55. Let $I \subseteq k[\mathbf{x}]$ be an ideal homogeneous with respect to grading by a positive vector $\mathbf{a} \in \mathbb{N}_{>0}^n$. We define three equivalence relations on \mathbb{R}^n :

- The *initial ideal* equivalence relation $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \text{in}_{\mathbf{u}}(I) = \text{in}_{\mathbf{v}}(I)$.
- The *top* equivalence relation $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \text{top}(\text{in}_{\mathbf{u}}(I)) = \text{top}(\text{in}_{\mathbf{v}}(I))$.
- The *radical* equivalence relation $\mathbf{u} \sim \mathbf{v} \Leftrightarrow \sqrt{\text{in}_{\mathbf{u}}(I)} = \sqrt{\text{in}_{\mathbf{v}}(I)}$.

In all three cases, the equivalence classes are invariant under translation by \mathbf{a} .

Proposition 3.56. *Let I be as in Definition 3.55. Then*

1. *The initial ideal equivalence relation defines the Gröbner fan of I .*
2. *The radical equivalence relation does not define a fan in general.*

A proof of the first claim is given in [47, Chapter 2]. See also [38]. The following example demonstrates the second claim.




Figure 3.3: The Gröbner fan from Example 3.57 and the two radical equivalence classes. The fan is drawn in the standard simplex with $(1, 0, 0)$ at the right bottom, $(0, 1, 0)$ at the left bottom and $(0, 0, 1)$ at the top.

Example 3.57. The radical equivalence classes of the homogeneous ideal $I = \langle c^4 - ba^3, ab^3 - ba^3 \rangle \subset k[a, b, c]$ are not all convex. Four of the eight monomial initial ideals of I have radical $\langle ab, ac, bc \rangle$ and the others have radical $\langle ab, c \rangle$. The intersection of the Gröbner fan with the two-dimensional standard simplex is shown in Figure 3.3 and the two radical equivalence classes appear in gray and white.

However, for toric ideals, all three equivalence relations of Definition 3.55 give rise to polyhedral fans.

Proposition 3.58.

1. *The radical equivalence relation of $I_{\mathcal{A}}$ defines the secondary fan of \mathcal{A} .*
2. *The top equivalence relation of $I_{\mathcal{A}}$ defines the hypergeometric fan of \mathcal{A} .*

Furthermore, the Gröbner fan of $I_{\mathcal{A}}$ is a refinement of the hypergeometric fan of \mathcal{A} , which is a refinement of the secondary fan of \mathcal{A} .

Proposition 3.58 may be taken as the definition of the hypergeometric and secondary fans of \mathcal{A} . The proposition is a collection of several known results [42, Proposition 3.3.1 and Corollary 3.3.2], [7], and [47, Chapter 8]. We now study the three equivalence classes for $J_{\mathcal{C}_{\mathcal{A}}}$.

Proposition 3.59. *The radical equivalence classes of J_{C_A} form a polyhedral fan that coincides with the secondary fan of I_A .*

Proof: This follows from Proposition 3.13 since $I_A = \sqrt{J_{C_A}}$. \square

For the top equivalence relation of J_{C_A} we need the following proposition which follows from Corollary 3.52.

Proposition 3.60. *The top equivalence relation for I_A and J_B are the same if I_A is homogeneous in the total degree grading and $I_A = \sqrt{J_B}$. It follows that the top equivalence relation of J_B defines the hypergeometric fan of I_A . In particular, this holds if $B = C_A$.*

We conjecture that the condition that I_A is homogeneous in the total degree grading can be left out. In contrast to Theorem 3.59 and Proposition 3.60, we have the following result for the initial ideal equivalence relation for I_A and J_{C_A} .




Figure 3.4: The Gröbner fans in the proof of Proposition 3.61 intersected with the simplex with coordinates $(0, 1, 0, 0)$ (right), $(0, 0, 1, 0)$ (left) and $(0, 0, 0, 1)$ (top). The Gröbner fan of J_{C_A} is to the left, the Gröbner fan of I_A in the middle, and the hypergeometric fan at the right.

Proposition 3.61. *In general, neither is the Gröbner fan of I_A a refinement of the Gröbner fan of J_{C_A} , nor vice-versa.*

Proof: Let $A = (7 \ 9 \ 13 \ 15)$. It is easy to check that $\text{in}_{(0,16,27,1)}(J_{C_A}) = \text{in}_{(0,20,25,3)}(J_{C_A})$ while $\text{in}_{(0,16,27,1)}(I_A) \neq \text{in}_{(0,20,25,3)}(I_A)$. Hence $(0, 16, 27, 1)$ and $(0, 20, 25, 3)$ lie in the

same maximal cell of the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ but in different maximal cells of the Gröbner fan of $I_{\mathcal{A}}$. This proves that the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ does not refine the Gröbner fan of $I_{\mathcal{A}}$. On the other hand, $\text{in}_{(0,4,19,9)}(I_{\mathcal{A}}) = \text{in}_{(0,4,16,5)}(I_{\mathcal{A}})$ and $\text{in}_{(0,4,19,9)}(J_{\mathcal{C}_{\mathcal{A}}}) \neq \text{in}_{(0,4,16,5)}(J_{\mathcal{C}_{\mathcal{A}}})$. Hence the Gröbner fan of $I_{\mathcal{A}}$ does not refine the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$. \square

Corollary 3.62. *If $n - d = 2$ the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ refines that of $I_{\mathcal{A}}$.*

Proof: By Theorem 3.3.8 in [42], if $n - d = 2$ then the Gröbner fan of $I_{\mathcal{A}}$ equals the hypergeometric fan of \mathcal{A} . The corollary then follows from Proposition 3.60 and the fact that the Gröbner fan of $J_{\mathcal{C}_{\mathcal{A}}}$ refines the hypergeometric fan of \mathcal{A} . \square

Chapter 4

STATE POLYTOPES OF TORIC IDEALS

4.1 Introduction

This chapter studies toric state polytopes from the point of view of complexity, properties, and characterization. The *state polytope* of an ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$, as described in Chapter 1, is a polytope whose vertices index the distinct reduced Gröbner bases of I . The normal fan of a state polytope is the *Gröbner fan* of I . The Gröbner fan can be used to convert one reduced Gröbner basis of I to another [12]. The *tropical variety* of I [41], a discrete approximation to the ordinary algebraic variety and the subject of Chapter 2 of this thesis, is a subfan of the Gröbner fan.

When I is a *toric ideal*, its state polytope has important additional structure. A toric ideal is the defining ideal of a *toric variety*, an algebraic variety equipped with the action of a dense open algebraic torus $(\mathbb{C}^*)^n$. A projective toric variety V can itself be specified by a polytope [25], so the state polytope of the ideal I defining V can be used to construct another toric variety: the main component of the *toric Hilbert scheme* [40] of all ideals whose Hilbert functions agree with that of I .

Toric ideals also provide an important algebraic method in integer programming, due originally to Conti and Traverso [13]. The generators of a toric ideal $I_{\mathcal{A}}$ are given by integer vectors in the kernel of an integer matrix A . A Gröbner basis of $I_{\mathcal{A}}$ with respect to a weight vector ω can be used to solve integer programs of the form

$$\text{maximize } \{\omega \cdot \mathbf{u} : A\mathbf{u} = \mathbf{b}, \mathbf{u} \in \mathbb{N}^n\}.$$

The edge directions of a state polytope of $I_{\mathcal{A}}$ are exactly the elements of a *universal Gröbner basis* of $I_{\mathcal{A}}$: the union of all reduced Gröbner bases of $I_{\mathcal{A}}$ [49]. Knowing the

set of these directions makes it possible to solve integer programs as above for any right-hand-side \mathbf{b} and for any cost vector ω . The combinatorics and especially the complexity of a toric state polytope are thus relevant to both algebraic geometry and optimization. In Section 4.2 we explain how to check via Gröbner basis computation whether a given embedded lattice polytope is a toric state polytope. A corollary of this approach is that up to normal equivalence, there is only one toric state polytope in any linear subspace of \mathbb{R}^n . It follows that the property of being a toric state polytope is highly sensitive to the arithmetic data given by the vertices.

In Section 4.3 we focus on the case of polygons. We show that for every k , there is a simple two-by-four matrix A that has an k -gon as its toric state polygon; this shows that the diameter of the state polytope of a toric ideal $I_{\mathcal{A}}$ can be exponential in the size of A . On the other hand, we show that every toric state polygon is *smooth*, a strong restriction on the positions of its vertices in \mathbb{Z}^n .

In Section 4.4 we explore the combinatorics of toric state polytopes of arbitrary dimension. Our main result, Theorem 4.21, is that any toric state polytope that is not a simplex must be *decomposable*: that is, it must have nontrivial Minkowski summands. Using a result of Shephard [44] on decomposability, we derive several corollaries: polytopes that are simplicial or nearly simplicial cannot be toric state polytopes. This allows us to determine which 3-polytopes with at most six vertices arise as toric state polytopes.

4.2 Definitions and Properties

As in Chapter 3, fix $A \in \mathbb{Z}^{d \times n}$ of full rank and containing a strictly positive vector in its row span. Set $m = n - d$. Recall that the *toric ideal* of A is the ideal

$$I_{\mathcal{A}} := \langle \mathbf{x}^{\mathbf{u}^+} - \mathbf{x}^{\mathbf{u}^-} : \mathbf{u} \in \mathcal{L}_{\mathcal{A}} \rangle \subseteq \mathbb{C}[x_1, \dots, x_n]$$

where $\mathcal{L}_{\mathcal{A}}$ is the integer kernel of A , an m -dimensional saturated sublattice of \mathbb{Z}^n . Toric ideals are exactly the binomial prime ideals in the polynomial ring. The stipulation

that there is a positive vector in the row span of A implies that I_A is homogeneous with respect to a positive grading. Often, we simply choose $(1, \dots, 1)$ as a row of A , and then I_A is homogeneous in the usual sense.

Recall from Chapter 1 that the *Gröbner fan* [38] of a homogeneous ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ is the complete fan in \mathbb{R}^n whose relatively open cones are the equivalence classes of the equivalence relation

$$\omega_1 \sim \omega_2 \text{ if } \text{in}_{\omega_1}(I) = \text{in}_{\omega_2}(I).$$

A *state polytope* of I is any polytope whose normal fan is the Gröbner fan. Such a polytope always exists and can be constructed algorithmically [47, §2]. Since the state polytope of a principal ideal $\langle f \rangle$ is just the *Newton polytope* of f (that is, the convex hull of the exponent vectors of the monomials that appear in f), every lattice polytope arises trivially as a state polytope of a principal ideal.

However, state polytopes of *toric* ideals carry additional geometric meaning and are not completely arbitrary lattice polytopes. For a generic $\omega \in \mathbb{R}^n$ and any $\mathbf{b} \in \mathbb{Z}^d$, the integer program

$$\text{IP}_{A,\omega}(\mathbf{b}) := \text{minimize } \{\omega \cdot \mathbf{u} : A\mathbf{u} = \mathbf{b}, \mathbf{u} \in \mathbb{N}^n\}$$

can be solved by identifying a feasible solution \mathbf{u} and reducing it by a Gröbner basis for I_A with respect to ω [13]. Allowing ω to vary through \mathbb{R}^n gives the following description of toric Gröbner fans.

Theorem 4.1. [49, Theorem 3.10] *Two vectors $\omega_1, \omega_2 \in \mathbb{R}^n$ lie in the same relatively open cone of the Gröbner fan if and only if for every $\mathbf{b} \in \mathbb{Z}^d$, the integer programs $\text{IP}_{A,\omega_1}(\mathbf{b})$ and $\text{IP}_{A,\omega_2}(\mathbf{b})$ have the same optima or are both infeasible.*

This suggests a method for constructing toric state polytopes. Define $\pi : \mathbb{N}^n \rightarrow \mathbb{Z}^d$ by $\mathbf{u} \mapsto A\mathbf{u}$. We define the *fiber* of π over \mathbf{b} to be the polytope $P_{\mathbf{b}}^I := \text{conv}(\pi^{-1}(\mathbf{b}))$. Note that since the row space of A contains a strictly positive vector, the fibers are

indeed polytopes rather than unbounded polyhedra. A *Gröbner fiber* is a fiber $P_{\mathbf{b}}^I$ such that there exist $\mathbf{u}, \mathbf{v} \in \mathbb{N}^n$ with $A\mathbf{u} = A\mathbf{v} = \mathbf{b}$ and the binomial $\mathbf{x}^{\mathbf{u}} - \mathbf{x}^{\mathbf{v}}$ appears in some reduced Gröbner basis of I_A . The union of the reduced Gröbner bases of an ideal is a *universal Gröbner basis*, which for a toric ideal I_A we denote by UGB_A .

Proposition 4.2. [47, 49]

1. The Minkowski integral $\int_{\mathbf{b}} P_{\mathbf{b}}^I d\mathbf{b}$ is a state polytope for I_A .
2. The Minkowski sum of all Gröbner fibers of A is a state polytope for I_A .
3. A fiber $P_{\mathbf{b}'}^I$ is a Gröbner fiber if and only if it has an edge which is not parallel to any edge of any fiber $P_{\mathbf{b}}^I$, such that $\mathbf{b}' < \mathbf{b}$; that is, such that $\mathbf{b}'_i \leq \mathbf{b}_i$ for all i and $\mathbf{b}' \neq \mathbf{b}$.
4. A integer vector $\mathbf{e}^+ - \mathbf{e}^-$ (with $\mathbf{e}^+, \mathbf{e}^- \geq 0$) lies in the minimal universal Gröbner basis of I_A if and only if the line segment $[\mathbf{e}^+, \mathbf{e}^-]$ is an edge of the fiber $P_{A\mathbf{e}^+}^I$.

Corollary 4.3. [49, §5] The primitive edge directions $\mathbf{e}^+ - \mathbf{e}^-$ of the edges of a toric state polytope exactly correspond to the elements $\mathbf{x}^{\mathbf{e}^+} - \mathbf{x}^{\mathbf{e}^-}$ of the universal Gröbner basis UGB_A of the toric ideal.

The Gröbner fan of I_A is a complete fan in \mathbb{R}^n whose lineality space is the (d -dimensional) rowspace of A , so any state polytope of A is m -dimensional. Also note that I_A and hence its toric state polytope depend only on $\ker_{\mathbb{Z}}(A)$, not on A itself.

It is possible to algorithmically determine whether a given lattice polytope $P \subseteq \mathbb{R}^n$ is a toric state polytope. Let $E(P) \subset \mathbb{Z}^n$ be the set of primitive edge directions of P and $L(P)$ be the lattice in \mathbb{Z}^n spanned by $E(P)$. Also set $B(P) = \{\mathbf{x}^{\mathbf{e}^+} - \mathbf{x}^{\mathbf{e}^-} : \mathbf{e} \in E(P)\}$ and $I(P) = \langle B(P) \rangle \subseteq \mathbb{C}[\mathbf{x}]$. Note that $I(P)$ is a binomial ideal.

Proposition 4.4. [20, Theorem 2.1] The ideal $I(P)$ is prime if and only if $L(P)$ is a saturated sublattice of \mathbb{Z}^n .

Proposition 4.5. *If P is a toric state polytope, then $I(P)$ is prime and $B(P)$ is exactly the union of all reduced Gröbner bases of $I(P)$.*

Proof: Suppose that P is the state polytope of a toric ideal I . By Corollary 4.3, the generating set $B(P)$ for $I(P)$ is the union of the reduced Gröbner bases for I , so $I = I(P)$. Since $I(P) = I$ is a toric ideal, it is prime. \square

Using Propositions 4.4 and 4.5, one can test whether P is a toric state polytope. First check whether the lattice $L(P)$ is saturated by computing the Smith normal form of a matrix whose columns generate $L(P)$. If so, then compute the universal Gröbner basis for $I(P)$, using $B(P)$ as a generating set, and see whether it contains any polynomials not already in $B(P)$. If not, then P passes both tests, and we must actually compute the Gröbner fan of $I(P)$ and check if it is the normal fan of P . Universal Gröbner bases and Gröbner fans can be computed by the software package `Gfan` [34].

Example 4.6.

1. Let P be the parallelogram in \mathbb{R}^3 with vertices $(0, 0, 0)$, $(2, 0, -1)$, $(0, 4, -3)$, and $(2, 4, -4)$. Then $E(P) = \{(2, 0, -1), (0, 4, -3)\}$ and the lattice $L(P)$ that it spans is not saturated: $L(P)$ contains $(2, 4, 4)$ but not $(1, 2, 2)$. By the first condition in Proposition 4.4, P is not a toric state polytope.
2. Let H be a lattice hexagon in \mathbb{R}^3 with three pairs of opposite edges, arranged so that $E(H) = \{(2, 0, -1), (0, 4, -3), (3, -2, 0)\}$. Then $L(H)$ is saturated, being the kernel of the matrix $A = \begin{pmatrix} 2 & 3 & 4 \end{pmatrix}$. The set $B(H) = \{a^2 - c, b^4 - c^3, a^3 - b^2\}$ generates $I_A = I(P)$. However, $B(H)$ is not a universal Gröbner basis for $I(P)$. For example, the lexicographic Gröbner basis (with $a \succ b \succ c$) contains $ac - b^2 \notin B(H)$. So by the second condition in Proposition 4.5, H again is not a toric state polytope.

3. Let Z be the lattice zonotope (a 10-gon) that is the Minkowski sum of the vectors in the set $E = \{(2, 0, -1), (0, 4, -3), (3, -2, 0), (1, -2, 1), (1, 2, -2)\}$. The lattice $L(Z)$ generated by the vectors in E is saturated, and furthermore the binomials obtained from E are a universal Gröbner basis of $I(Z)$. However, the Gröbner fan of $I(Z)$ has the following six reduced Gröbner bases:

$$\{b^4 - c^3, ac - b^2, ab^2 - c^2, a^2 - c\}$$

$$\{c^3 - b^4, ac - b^2, ab^2 - c^2, a^2 - c\}$$

$$\{c^2 - ab^2, ac - b^2, a^2 - c\}$$

$$\{c - a^2, a^3 - b^2\}$$

$$\{c - a^2, b^2 - a^3\}$$

$$\{b^2 - ac, a^2 - c\}$$

and hence the state polytope of $I(E(Z))$ has six vertices. On the other hand, the zonotope Z has ten edges and hence ten vertices. This shows that Z is not a toric state polytope.

Corollary 4.7. *Given a rational linear subspace V of \mathbb{R}^n , there exists a toric state polytope P whose linear span equals V . Furthermore, P is unique up to normal equivalence: any two such polytopes have the same normal fan.*

Proof: Consider the lattice $L = V \cap \mathbb{Z}^n$. This lattice is the unique saturated lattice whose span is V , so the only toric ideal $I_{\mathcal{A}}$ whose state polytopes span exactly V is given by any matrix A whose integer kernel is L . All state polytopes of $I_{\mathcal{A}}$ have the Gröbner fan of $I_{\mathcal{A}}$ as their normal fan, so all are normally equivalent. \square

Remark 4.8. It follows from Proposition 4.7 that if P is a toric state polytope with linear span V and if Q is obtained from P by moving one vertex to another integer point of V , then Q cannot also be a toric state polytope. That is, the property of

being a toric state polytope is not at all combinatorial, but instead is highly sensitive to arithmetic data.

4.3 Toric state polygons

In this section we consider the case $m = 2$. We show that a k -gon arises as a toric state polytope for every k , but that there are strong restrictions on the placement of its vertices: every toric state polygon is *smooth*.

Given $A \in \mathbb{Z}^{d \times n}$ as before, fix a matrix $B \in \mathbb{Z}^{n \times m}$ whose columns form a basis for the lattice \mathcal{L}_A . Such a matrix B is called an *integer Gale dual* of A . For $\mathbf{u} \in \mathbb{N}^n$ let

$$Q_{\mathbf{u}} := \{\mathbf{z} \in \mathbb{R}^{n-d} : B\mathbf{z} \leq \mathbf{u}\}.$$

The polyhedron $Q_{\mathbf{u}}$ is the image of $P_{A\mathbf{u}}$ under the isomorphism

$$\phi_{\mathbf{u}} : \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = A\mathbf{u}\} \rightarrow \mathbb{R}^{n-d} \text{ such that } \mathbf{x} \mapsto \mathbf{z} \text{ where } \mathbf{u} - B\mathbf{z} = \mathbf{x}.$$

Furthermore, $\phi_{\mathbf{u}}$ maps the fiber $P_{A\mathbf{u}}^I$ isomorphically onto the polytope

$$Q_{\mathbf{u}}^I := \text{conv}\{\mathbf{z} \in \mathbb{Z}^m : B\mathbf{z} \leq \mathbf{u}\}.$$

It follows that the state polytope of I_A , the Minkowski sum of the Gröbner fibers, is isomorphic to $(\sum_{\mathbf{u}} Q_{\mathbf{u}}^I) \subset \mathbb{R}^m$, where the sum is taken over one \mathbf{u} from each Gröbner fiber. The Gröbner fan is mapped into \mathbb{R}^m in such a way that the combinatorics is preserved, while the lineality space (the row space of A) is sent to 0. Its image in \mathbb{R}^m is called the *pointed Gröbner fan* of A .

Definition 4.9. [7] The *secondary fan* of A is the complete fan in \mathbb{R}^n whose cones are the equivalence classes of the equivalence relation

$$\omega \sim \omega' \text{ if } \Delta_{\omega} = \Delta_{\omega'},$$

where Δ_{ω} denotes the regular subdivision of $\text{cone}(A)$ (using only columns of A) induced by ω .

Since Δ_ω is unchanged by adding any vector in the row space of A to ω , it follows that the secondary fan, like the Gröbner fan, is combinatorially isomorphic to a fan in \mathbb{R}^m , the *pointed secondary fan*. Furthermore, regular triangulations (i.e., subdivisions in which every face is a simplex) are in bijection with *radicals* of initial ideals of I_A [47, §8]. It follows that the Gröbner fan of A refines the secondary fan.

Proposition 4.10. [7] *The pointed secondary fan of A in \mathbb{R}^m is the chamber complex of B ; that is, the common refinement of all complete fans whose cones are spanned by subsets of the rows of B .*

In dimension two, this fan is easy to construct: it is formed by drawing the rays whose directions are given by the rows of B and filling in the two-dimensional cones between them. Even better, the Gröbner fan and a state polytope can also be described explicitly as follows.

Given a pointed rational polyhedral cone $C \subset \mathbb{R}^m$, its *Hilbert basis* is the (necessarily finite) minimal generating set of the semigroup $C \cap \mathbb{Z}^m$.

Proposition 4.11. [42] *Given $A \in \mathbb{N}^{(n-2) \times n}$, the pointed Gröbner fan of I_A is obtained by subdividing each cone of the pointed secondary fan along the rays generated by its Hilbert basis elements.*

A polytope $P \subset \mathbb{R}^n$ is said to be *smooth* if at each vertex v of P , the primitive integer direction vectors of the edges incident to v form part of a lattice basis of \mathbb{Z}^n . An equivalent condition is that the rays of the normal cone to P at v form part of a basis of \mathbb{Z}^n [22, Theorem 2.10, Chapter V]. A polytope P is smooth if and only if the projective toric variety defined by P is smooth; see [25, §2].

Theorem 4.12. *Every toric state polygon is smooth.*

Proof: Let v be any vertex of a toric state polygon. By Proposition 4.11, the pointed Gröbner cone C_v is generated by two adjacent rays spanned by Hilbert basis elements

of a secondary cone. By [30, Proposition 2.2], any two such elements must span \mathbb{Z}^2 , so v is smooth. \square

Example 4.13. Let

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 \end{pmatrix}.$$

An integer Gale dual is

$$B = \begin{pmatrix} 1 & 2 \\ -2 & -3 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

so the pointed secondary fan is the complete fan in \mathbb{R}^2 with rays in directions $(1, 0)$, $(1, 2)$, $(0, 1)$, and $(-2, -3)$ in counterclockwise order from the positive horizontal axis. Note that the secondary polytope is not smooth.

The Gröbner fan is constructed by computing Hilbert bases of each full-dimensional cone; these Hilbert bases are $\{(1, 0), (1, 1), (1, 2)\}$, $\{(1, 2), (0, 1)\}$, $\{(0, 1), (1, -1), (-2, -3)\}$, and $\{(-2, -3), (-1, -2), (0, -1), (1, 0)\}$. The pointed secondary fan and Gröbner fan are shown in Figure 4.1. A state polytope, not shown, would be any octagon whose normal fan is the Gröbner fan.

Remark 4.14.

1. Not every smooth polygon is a toric state polygon. This is an immediate consequence of Corollary 4.7, since many smooth polygons have the same affine span.
2. Higher-dimensional toric state polytopes need not be smooth; in fact they need not even be simple. For example, take

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 4 \end{pmatrix}.$$




Figure 4.1: The pointed secondary fan (left) and pointed Gröbner fan (right) in Example 4.13.

Using `Gfan` [34], we identify several vertices of the (three-dimensional) state polytope whose degrees are four rather than three. One of these vertices indexes the Gröbner basis $\{d^2 - ce, cd - be, c^2 - ae, bd - ae, bc - ad, b^2 - ac\}$. \square

Theorem 4.15. *For every $k \geq 3$, there is a matrix $A \in \mathbb{Z}^{2 \times 4}$ such that the toric state polytope of I_A is a k -gon in \mathbb{R}^4 .*

Proof: For $k = 3$, take $A = (1 \ 1 \ 1)$. The universal Gröbner basis of I_A consists of the three linear forms $x_1 - x_2$, $x_1 - x_3$, and $x_2 - x_3$. All three arise from the same Gröbner fiber P_1^I , a triangle, so this fiber is a state polytope for I_A by Proposition 4.2.

For $k \geq 4$, set $j = k - 3$ and take

$$A = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & j & j-1 & 0 \end{pmatrix}.$$

The sum of the two rows is a strictly positive vector, so the toric ideal I_A is positively

graded. An integer Gale dual of A is

$$B = \begin{pmatrix} 1 & j \\ -1 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Three of the four cones in the pointed secondary fan are unimodular. The fourth, spanned by the first and third rows of B , has as its Hilbert basis the set

$$\{(1, 0), (1, 1), \dots, (1, j - 1), (1, j)\},$$

so it contains j Gröbner cones. Thus the state polytope has $j + 3 = k$ vertices. \square

4.4 Combinatorial types

We now turn our attention from embedded polytopes to combinatorial types; i.e., *face lattices*. Face lattices and f -vectors (a coarser invariant that simply counts the number of faces of each dimension) are the fundamental combinatorial invariants of polytopes. Since the face lattice of the normal fan of a polytope P is obtained by inverting the face lattice of P with the empty face removed, our results can be easily adapted for combinatorial types of Gröbner fans.

Proposition 4.16. *The following (combinatorial) polytopes arise as toric state polytopes.*

1. *An m -simplex for every m .*
2. *The product in \mathbb{R}^{i+j} of any two toric state polytopes $P \in \mathbb{R}^i$, $Q \in \mathbb{R}^j$.*
3. *An m -cube for every m .*

Proof:

1. Let \mathcal{A} be the $1 \times (m + 1)$ configuration given by

$$A = (1 \ 1 \ \dots \ 1).$$

As in the first part of the proof of Theorem 4.15, the universal Gröbner basis is given by

$$\{x_j - x_1 : 1 \leq i < j \leq m + 1\}$$

so the only Gröbner fiber (over $\mathbf{b} = 1$) is the standard m -simplex. So by Proposition 4.2, the state polytope of \mathcal{A} is a standard simplex.

2. Suppose P and Q are toric state polytopes of the respective matrices A_1 and A_2 . Define

$$A = \begin{pmatrix} A_1 & \mathbf{0} \\ \mathbf{0} & A_2 \end{pmatrix}.$$

Given a weight vector (ω, η) , we have

$$\text{in}_{(\omega, \eta)}(I_{\mathcal{A}}) = \text{in}_{\omega}(I_{A_1}) + \text{in}_{\eta}(I_{A_2}),$$

the sum of two ideals generated by polynomials in disjoint sets of variables. In particular,

$$\text{in}_{\prec_{\omega, \eta}}(I_{\mathcal{A}}) = \text{in}_{\prec_{\omega', \eta'}}(I_{\mathcal{A}})$$

if and only if

$$\text{in}_{\omega}(I_{A_1}) = \text{in}_{\omega'}(I_{A_1}) \text{ and } \text{in}_{\eta}(I_{A_2}) = \text{in}_{\eta'}(I_{A_2}).$$

Thus the Gröbner fan of $I_{\mathcal{A}}$ is the common refinement of the Gröbner fans of I_{A_1} and I_{A_2} , so a state polytope of $I_{\mathcal{A}}$ is the Minkowski sum of the two state polytopes. Since the state polytopes live in orthogonal subspaces \mathbb{R}^i and \mathbb{R}^j of \mathbb{R}^{i+j} , this Minkowski sum is just the direct product.

3. The m -cube is the product of m line segments (1-dimensional simplices), so this follows immediately from (1) and (2). In particular, the state polytope of the $m \times 2m$ configuration

$$A = (\mathbf{e}_1 \ \mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_m \ \mathbf{e}_m)$$

is an empty m -cube, where $\mathbf{e}_1, \dots, \mathbf{e}_m$ form the standard basis of \mathbb{Z}^m .

□

Theorem 4.17. *Fix $m \geq 2$. The diameter of an m -dimensional toric state polytope can be arbitrarily large and can grow exponentially in the bit size of the matrix A .*

Proof: By Proposition 4.16 and Theorem 4.15, the product of a k -gon and $m - 2$ line segments is the toric state polytope of a certain $m \times 2m$ matrix A whose largest entry is $k - 3$. Its diameter is $\lfloor k/2 \rfloor + m - 2$ while the bit size of A is $\mathcal{O}(m^2 \log k)$. □

Remark 4.18. The diameter of the state polytope of a toric ideal $I_{\mathcal{A}}$ measures the complexity of the Gröbner walk algorithm used to convert one Gröbner basis of $I_{\mathcal{A}}$ to another [12, 24]. Theorem 4.17 shows that the algorithm may require exponential time, even apart from the complexity of computing individual Gröbner bases.

We now develop necessary conditions for a combinatorial polytope P to be realizable as a toric state polytope. For technical reasons, we now consider only matrices A with only nonnegative entries, though a similar argument works with the original assumptions on A . We say that a Gröbner degree \mathbf{b} is *basic* if there is no other Gröbner degree \mathbf{b}' that is less than \mathbf{b} with respect to the natural partial order on \mathbb{N}^n . If so, we also say that the Gröbner fiber $P_{\mathbf{b}}^I$ is basic.

Lemma 4.19. *Every basic Gröbner fiber $P_{\mathbf{b}}^I$ is a simplex.*

Proof: Fix any term order \prec . Let $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_r, \mathbf{v}$ be the elements of $\pi^{-1}(\mathbf{b})$, labeled so that $\mathbf{x}^{\mathbf{v}}$ is the unique standard monomial with respect to \prec . Then $f_i := \mathbf{x}^{\mathbf{u}_i} - \mathbf{x}^{\mathbf{v}} \in I_{\mathcal{A}}$ for each $i = 1, \dots, r$ and its \prec -leading term is $\mathbf{x}^{\mathbf{u}_i}$.

Fix $1 \leq i < j \leq r$ and set $\mathbf{x}^{\gamma} := \text{LCM}(\mathbf{x}^{\mathbf{u}_i}, \mathbf{x}^{\mathbf{u}_j})$. The S-pair

$$\begin{aligned} S(f_i, f_j) &= \mathbf{x}^{\gamma - \mathbf{u}_i}(\mathbf{x}^{\mathbf{u}_i} - \mathbf{x}^{\mathbf{v}}) - \mathbf{x}^{\gamma - \mathbf{u}_j}(\mathbf{x}^{\mathbf{u}_j} - \mathbf{x}^{\mathbf{v}}) \\ &= \mathbf{x}^{\gamma - \mathbf{u}_j + \mathbf{v}} - \mathbf{x}^{\gamma - \mathbf{u}_i + \mathbf{v}} \\ &= \mathbf{x}^{\mathbf{v}}(\mathbf{x}^{\gamma - \mathbf{u}_j} - \mathbf{x}^{\gamma - \mathbf{u}_i}) \end{aligned}$$

is in $I_{\mathcal{A}}$. Since $I_{\mathcal{A}}$ is prime and contains no monomials, it follows that $f := \mathbf{x}^{\gamma - \mathbf{u}_j} - \mathbf{x}^{\gamma - \mathbf{u}_i} \in I_{\mathcal{A}}$.

If $\mathbf{x}^{\mathbf{u}_i}$ and $\mathbf{x}^{\mathbf{u}_j}$ are not relatively prime, then γ is strictly less than $\mathbf{u}_i + \mathbf{u}_j$, so

$$A(\gamma - \mathbf{u}_i) = A(\gamma - \mathbf{u}_j) < A\mathbf{u}_i = \mathbf{b}.$$

That is, the \prec -leading term of f (whichever term it is) is a monomial in $\text{in}_{\omega}(I_{\mathcal{A}})$ whose A -degree is strictly less than \mathbf{b} , contradicting the assumption that \mathbf{b} is basic. So the monomials $\mathbf{x}^{\mathbf{u}_1}, \dots, \mathbf{x}^{\mathbf{u}_r}$ are pairwise relatively prime. This implies that no variable occurs in more than one of these monomials, so in particular the set of exponent vectors $\{\mathbf{u}_1, \dots, \mathbf{u}_r\}$ is linearly independent. Then the complete set $\{\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{v}\}$ is affinely independent, so indeed $P_{\mathbf{b}}^I$ is a simplex. \square

Definition 4.20. A polytope P is *indecomposable* if it cannot be expressed as the Minkowski sum of two polytopes not homothetic to P .

Theorem 4.21. *The only indecomposable state polytopes are simplices.*

Proof: If P is the state polytope of a toric ideal $I_{\mathcal{A}}$ and is indecomposable, then by parts (1) and (2) of Proposition 4.2, it follows that A has only one Gröbner fiber G and that $P \simeq G$. Then by Lemma 4.19, P is a simplex. \square

Decomposability of polytopes has been well studied. The first result is due to Gale [26]: any pyramid (i.e., a polytope that is the convex hull of a vertex and a facet) is indecomposable. Shephard showed that every simplicial polytope is indecomposable. More generally, we will apply the following result of Shephard to derive interesting corollaries of Theorem 4.21. Strengthenings of Shephard's result are available; see [36].

Proposition 4.22. [44] *Suppose P is a polytope containing a sequence of triangular 2-dimensional faces T_1, T_2, \dots, T_m such that*

1. *The intersection $T_i \cap T_{i-1}$ is an edge for each $2 \leq i \leq m$ and*
2. *the union $\bigcup_{i=1}^m T_i$ contains every vertex of P .*

Then P is indecomposable.

Corollary 4.23.

1. *If P is simplicial but not a simplex, then P cannot be realized as a toric state polytope.*
2. *If P has exactly one non-simplicial facet, then P cannot be realized as a toric state polytope.*
3. *For $d > 2$, there are no d -dimensional toric state polytopes with exactly $d + 2$ vertices.*
4. *The only three-dimensional toric state polytopes with fewer than seven vertices are the simplex and the triangular prism.*

Proof:

1. We apply Theorem 4.21 and Proposition 4.22. The statement follows immediately; in fact if P is even *2-simplicial* (that is, if all of its two-dimensional faces are simplices), then it cannot be a toric state polytope.
2. Suppose P has a unique non-simplicial facet F . The polyhedral complex obtained from the boundary of P by removing the relative interior of F is homeomorphic to a disk. Its two-dimensional faces, all triangles, must then be completely connected by ridge paths.
3. Any d -polytope with $d + 2$ vertices is either simplicial or a pyramid [53, §6.5]. In either case, it is indecomposable by previous results.
4. Steinitz's theorem [53, Theorem 4.1] tells us that a graph G is the 1-skeleton of a 3-polytope if and only if G is planar and 3-connected. This allows us to enumerate all 3-polytopes with a small number of vertices. The only 3-polytopes with five vertices are the square pyramid and the bipyramid over a triangle, both ruled out by (1).

There are seven 3-polytopes with six vertices. One is the triangular prism, the product of an edge and a triangle in orthogonal spaces, which is indeed a toric state polytope by Proposition 4.16. Of the remaining six, two are simplicial and three have exactly one non-triangular facet. The last has two square facets, but its four triangular facets can be ordered to satisfy the conditions of Proposition 4.22, as shown in Figure 4.2.

□

Remark 4.24. It follows from these results that, unlike products, joins and direct sums of toric state polytopes need not be toric state polytopes. The join of a line segment and a square is a square pyramid, which has only one non-simplicial facet. The direct sum of a line segment and a square (that is, the polytope obtained by

placing them in skew affine subspaces of \mathbb{R}^4 and taking the convex hull) is a 4-polytope with six vertices.




Figure 4.2: Schlegel diagrams of the seven combinatorially distinct 3-polytopes with six vertices. The first is a toric state polytope; the others are all decomposable. The last one has two non-triangular facets but its triangles are numbered to form a ridge path.

Chapter 5

SMALL CHVÁTAL RANK

This chapter consists of joint work with Rekha Thomas on the geometry of families of polytopes and their integer hulls. As in Chapter 4, these families are obtained by varying the right-hand-side of a system of equations and inequalities. However, the notation will now be different: instead of the polytopes $P_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$, associated with a toric ideal $I_{\mathcal{A}}$, we will consider polytopes of the form $Q_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$. Such polytopes arose as *Gale duals* in the preceding chapters. This form, however, is the most common one in integer programming, and we employ it here because we are no longer using toric ideals.

We introduce a new measure of complexity of integer hulls of rational polyhedra called the small Chvátal rank (SCR). The SCR of an integer matrix A is the number of rounds of a Hilbert basis procedure needed to generate all normals of a sufficient set of inequalities to cut out the integer hulls of all polyhedra $\{\mathbf{x} : A\mathbf{x} \leq \mathbf{b}\}$ as \mathbf{b} varies. The SCR of A is bounded above by the Chvátal rank of A and is hence finite. We exhibit examples where SCR is much smaller than Chvátal rank. When the number of columns of A is at least three, we show that SCR can be arbitrarily high proving that, in general, SCR is not a function of dimension alone. For polytopes in the unit cube we provide a lower bound for SCR that is comparable to the known lower bounds for Chvátal rank in that situation. Lastly, we establish the connection between SCR and the notion of supernormality.

5.1 Introduction

The study of integer hulls of rational polyhedra is a fundamental area of research in integer programming and discrete geometry. For a matrix $A \in \mathbb{Z}^{m \times n}$ and vector $\mathbf{b} \in \mathbb{Z}^m$, consider the rational polyhedron $Q_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$ and its integer hull $Q_{\mathbf{b}}^I := \text{convex hull}(Q_{\mathbf{b}} \cap \mathbb{Z}^n)$. An algorithm for computing $Q_{\mathbf{b}}^I$ from the inequality description of $Q_{\mathbf{b}}$ is given by the *Chvátal-Gomory* procedure [43, §23]. This method involves iteratively adding rounds of *cutting planes* to $Q_{\mathbf{b}}$ until $Q_{\mathbf{b}}^I$ is obtained. The *Chvátal rank* of $A\mathbf{x} \leq \mathbf{b}$ is the minimum number of rounds of cuts needed in the Chvátal-Gomory procedure to obtain $Q_{\mathbf{b}}^I$, and the *Chvátal rank* of A is the maximum over the Chvátal ranks of $A\mathbf{x} \leq \mathbf{b}$ as \mathbf{b} varies in \mathbb{Z}^m . The Chvátal-Gomory procedure and the Chvátal ranks defined above are all finite [43, §23].

In this paper we fix a matrix $A \in \mathbb{Z}^{m \times n}$ of rank n and look at the problem of finding normals of a sufficient set of inequalities to cut out all integer hulls $Q_{\mathbf{b}}^I$ as \mathbf{b} varies in \mathbb{Z}^m . By this we mean finding a matrix M such that for each $Q_{\mathbf{b}}^I$ there exists an integer vector \mathbf{d} such that $Q_{\mathbf{b}}^I = \{\mathbf{x} \in \mathbb{R}^n : M\mathbf{x} \leq \mathbf{d}\}$. Theorem 17.4 in [43] proves the existence of such an M . The method is to first prove that if Δ is the maximum absolute value of a minor of A , then every $Q_{\mathbf{b}}^I$ can be described by inequalities whose normal vectors have entries of absolute value at most $n^{2n}\Delta^n$. Then M is taken to be the matrix whose rows are all of the non-zero integer vectors in the cone generated by the rows of A whose entries have absolute value at most $n^{2n}\Delta^n$. For example, the matrix

$$A = \begin{pmatrix} 1 & 2 \\ -2 & -3 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

has $\Delta = 3$ and $n^{2n}\Delta^n = 144$. Therefore, the rows of M would be all the non-zero

integer vectors in the box

$$\{(x, y) \in \mathbb{R}^2 : -144 \leq x, y \leq 144\}.$$

However, it suffices to use

$$M = \begin{pmatrix} 1 & 2 \\ -2 & -3 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \\ 0 & -1 \\ -1 & -2 \\ -1 & -1 \end{pmatrix}$$

and every row of M is actually a facet normal in some $Q_{\mathbf{b}}^I$. In fact, if $A \in \mathbb{Z}^{m \times 2}$ has rank two, we will show that there is always an M with at most $m\Delta$ rows (Corollary 5.18). Thus a sufficient M may be much smaller than the one constructed in [43, Theorem 17.4].

As in this example, we are interested in matrices M that are as economical as possible. We introduce a modified version of the Chvátal-Gomory procedure called *iterated basis normalization* (IBN) that in principle constructs M . The *small Chvátal rank* (SCR) of A is the least number of rounds of IBN necessary to generate an M that works for all $Q_{\mathbf{b}}^I$ as \mathbf{b} varies. For a fixed \mathbf{b} , the *small Chvátal rank* of $A\mathbf{x} \leq \mathbf{b}$ is the least number of iterations of IBN needed to obtain an M that describes that particular $Q_{\mathbf{b}}^I$. It can be shown that SCR of A (respectively $A\mathbf{x} \leq \mathbf{b}$) is at most its Chvátal rank and is hence finite. In contrast, IBN may not terminate when $n \geq 3$. These definitions and preliminaries are described in Section 5.2.

In Section 5.3 we show that when $n = 2$, SCR of A is at most one while it is known that Chvátal rank can be arbitrarily large. Such matrices are shown to exist for all $n \geq 2$. It is also shown that for a certain standard linear relaxation of the

co clique polytope of the complete graph K_n , SCR is either one or two depending on the parity of n while Chvátal rank is known to be $\mathcal{O}(\log n)$. These examples support our use of the adjective “small”.

In contrast, we show in Section 5.5 that in dimension greater than or equal to three, SCR may grow exponentially in the bit size of the matrix A , asymptotically just as fast as Chvátal rank. Thus for $n > 2$, SCR is not a function of m and n alone. Using a construction by Alon and Vü of 0/1 matrices with large determinants, we also show that SCR can be large even when $Q_{\mathbf{b}}$ is contained in the unit cube. The Chvátal rank of polytopes in the unit cube has been well studied. Our methods provide alternate ways to prove lower bounds on Chvátal rank.

In Section 5.4 we establish the relationship between SCR of A and *supernormality* of the vector configuration, \mathcal{A} , consisting of the rows of A [30]. If \mathcal{A} is supernormal then we prove that A also serves as M which implies that for every \mathbf{b} , $Q_{\mathbf{b}}^I$ has the form $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{d}\}$ for some integer \mathbf{d} . In this sense, supernormality is a generalization of unimodularity. We produce a family of vector configurations in growing dimension that are supernormal but not unimodular, answering a question in [30].

The *integrality gap* of a vector $\mathbf{c} \in \mathbb{Z}^n$ with respect to $Q_{\mathbf{b}}$ is

$$\max\{\mathbf{c}\mathbf{x} : \mathbf{x} \in Q_{\mathbf{b}}\} - \max\{\mathbf{c}\mathbf{x} : \mathbf{x} \in Q_{\mathbf{b}}^I\}.$$

A standard technique in estimating the Chvátal rank of $A\mathbf{x} \leq \mathbf{b}$ is to bound the integrality gap of a vector \mathbf{c} on $Q_{\mathbf{b}}$ and to use that to bound the number of iterations of the Chvátal-Gomory procedure on $Q_{\mathbf{b}}$. For the small Chvátal rank of A we only care about the normals of the inequalities describing $Q_{\mathbf{b}}^I$ and not about their right-hand-sides. Therefore, integrality gap does not appear to be a useful tool in the study of SCR. The problem here is to understand how deep in the IBN/Chvátal-Gomory procedure the last facet normal of an integer hull $Q_{\mathbf{b}}^I$ will be generated. Integrality gap and the tools needed for SCR appear to be complementary approaches

to understanding the difference between $Q_{\mathbf{b}}^I$ and $Q_{\mathbf{b}}$.

5.2 Main Definitions

Fix a matrix $A \in \mathbb{Z}^{m \times n}$ of rank n . Let $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_m\}$ be the vector configuration in \mathbb{Z}^n consisting of the rows of A . For each $\mathbf{b} = (b_1, \dots, b_m) \in \mathbb{Z}^m$, define the polyhedron

$$Q_{\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$$

and its *integer hull* (using conv for *convex hull*)

$$Q_{\mathbf{b}}^I := \text{conv}\{\mathbf{x} \in \mathbb{Z}^n : A\mathbf{x} \leq \mathbf{b}\}.$$

We may assume that the components of each \mathbf{a}_i are relatively prime since an inequality $\mathbf{a}_i \cdot \mathbf{x} \leq b_i$ in the description of $Q_{\mathbf{b}}$ is unaffected by division by a positive number. Since $\text{rank}(A) = n$, $m \geq n$. The Chvátal-Gomory procedure [11], [43, §23] for computing $Q_{\mathbf{b}}^I$ works as follows.

For each minimal face F of $Q_{\mathbf{b}}$, set

$$\mathcal{A}_F := \{\mathbf{a}_i : \mathbf{a}_i \cdot \mathbf{x} = b_i \ \forall \mathbf{x} \in F\}.$$

If \mathbf{h} is an integer vector in $\text{cone}(\mathcal{A}_F)$, the cone generated by \mathcal{A}_F , and $\mathbf{f} \in F$, then the inequality $\mathbf{h} \cdot \mathbf{x} \leq \mathbf{h} \cdot \mathbf{f}$ is valid on $Q_{\mathbf{b}}$, so the inequality $\mathbf{h} \cdot \mathbf{x} \leq \lfloor \mathbf{h} \cdot \mathbf{f} \rfloor$ is valid on $Q_{\mathbf{b}}^I$. Recall that a **Hilbert basis** of a rational polyhedral cone K is a set of integer vectors $\mathbf{h}_1, \dots, \mathbf{h}_t$ in K such that every integer vector $\mathbf{z} \in K$ can be written as a non-negative integer combination of $\mathbf{h}_1, \dots, \mathbf{h}_t$. The new polyhedron $Q_{\mathbf{b}}^{(1)}$ defined by the inequalities $\mathbf{h} \cdot \mathbf{x} \leq \lfloor \mathbf{h} \cdot \mathbf{f} \rfloor$ for every minimal face F of $Q_{\mathbf{b}}$ and every vector \mathbf{h} in a Hilbert basis of $\text{cone}(\mathcal{A}_F)$ satisfies

$$Q_{\mathbf{b}}^I \subseteq Q_{\mathbf{b}}^{(1)} \subseteq Q_{\mathbf{b}}.$$

Furthermore, every non-integral vertex of $Q_{\mathbf{b}}$ is cut off by this procedure (that is, lies outside $Q_{\mathbf{b}}^{(1)}$). For $i \geq 2$, inductively define $Q_{\mathbf{b}}^{(i)} := (Q_{\mathbf{b}}^{(i-1)})^{(1)}$. The **Chvátal rank** of

$A\mathbf{x} \leq \mathbf{b}$ is the smallest number t such that $Q_{\mathbf{b}}^{(t)} = Q_{\mathbf{b}}^I$. The **Chvátal rank** of A is the maximum over all $\mathbf{b} \in \mathbb{Z}^m$ of the Chvátal ranks of $A\mathbf{x} \leq \mathbf{b}$. The procedure and the ranks are finite [43, Chapter 23].

To study just the normals of inequalities needed for the integer hulls $Q_{\mathbf{b}}^I$ for all \mathbf{b} , we modify the Chvátal-Gomory procedure as follows. An n -subset $\tau \subseteq [m] := \{1, 2, \dots, m\}$ is called a *basis* if the submatrix A_{τ} consisting of the rows of A indexed by τ is non-singular. Let \mathcal{A}_{τ} be the set of rows of A_{τ} . Thus \mathcal{A}_{τ} is a basis of \mathbb{R}^n . We call $\text{cone}(\mathcal{A}_{\tau})$ a *basis cone*.

Algorithm 5.1. Iterated Basis Normalization (IBN)

Input: $A \in \mathbb{Z}^{m \times n}$ satisfying the assumptions above. Let \mathcal{A} be the set of rows of A .

1. Set $\mathcal{A}^{(0)} := \mathcal{A}$.
2. For $k \geq 1$, let $\mathcal{A}^{(k)}$ be the union of all the (unique) minimal Hilbert bases of all basis cones in $\mathcal{A}^{(k-1)}$.
3. If $\mathcal{A}^{(k)} = \mathcal{A}^{(k-1)}$, then stop.

Lemma 5.2. *Suppose σ is any subset of $[m]$ such that \mathcal{A}_{σ} linearly spans \mathbb{R}^n . Then the union of the minimal Hilbert bases of the basis cones $\text{cone}(\mathcal{A}_{\tau})$, as τ varies over the bases contained in σ , is a Hilbert basis for $\text{cone}(\mathcal{A}_{\sigma})$.*

Proof: Every integer point in $\text{cone}(\mathcal{A}_{\sigma})$ lies in $\text{cone}(\mathcal{A}_{\tau})$ for some basis $\tau \subseteq \sigma$ and hence can be written as a non-negative integer combination of the Hilbert basis elements of $\text{cone}(\mathcal{A}_{\tau})$. □

Proposition 5.3. *A sufficient set of normals needed by the Chvátal-Gomory procedure for obtaining $Q_{\mathbf{b}}^I$ from $Q_{\mathbf{b}}$, for all \mathbf{b} , is generated by IBN.*

Proof: If F is a minimal face of some intermediate polyhedron $Q^{(i)} = \{\mathbf{x} : U\mathbf{x} \leq \mathbf{u}\}$ in the Chvátal-Gomory procedure, then \mathcal{U}_F linearly spans \mathbb{R}^n and hence its index set

σ satisfies the hypothesis of Lemma 5.2. By induction, a Hilbert basis of $\text{cone}(\mathcal{U}_F)$ is produced within $(i + 1)$ iterations of IBN. \square

Let $A^{(k)}$ denote a matrix whose rows are the elements of the set $\mathcal{A}^{(k)}$ produced in the k th iteration of IBN.

Definition 5.4.

1. The **small Chvátal rank (SCR)** of the system of inequalities $A\mathbf{x} \leq \mathbf{b}$ defining $Q_{\mathbf{b}}$ is the smallest number k such that there is an integer vector \mathbf{b}' satisfying

$$Q_{\mathbf{b}}^I = \{\mathbf{x} \in \mathbb{R}^n : A^{(k)}\mathbf{x} \leq \mathbf{b}'\}.$$

2. The SCR of a matrix A is the maximum of the SCRs of all systems of the form $A\mathbf{x} \leq \mathbf{b}$ as \mathbf{b} varies in \mathbb{Z}^m .

Proposition 5.5. *For any $\mathbf{b} \in \mathbb{Z}^m$, the SCR of $A\mathbf{x} \leq \mathbf{b}$ is at most the Chvátal rank of the same system, and the SCR of $A \in \mathbb{Z}^{m \times n}$ is at most the Chvátal rank of A . In particular, SCR is always finite.*

Proof: By Lemma 5.2 and Proposition 5.3, if the Chvátal rank of $A\mathbf{x} \leq \mathbf{b}$ is t , then within t iterations, IBN generates the normals of all inequalities needed in the Chvátal-Gomory procedure on $A\mathbf{x} \leq \mathbf{b}$. \square

If IBN terminates after $k + 1$ iterations (since $\mathcal{A}^{(k+1)} = \mathcal{A}^{(k)}$) then k is an upper bound on the SCR of A and hence also on the SCR of $A\mathbf{x} \leq \mathbf{b}$ for all \mathbf{b} . So a natural first question to ask is whether IBN always terminates.

Lemma 5.6. *When $n = 2$, $\mathcal{A}^{(2)} = \mathcal{A}^{(1)}$.*

Proof: Pick $r, s \in \mathcal{A}^{(1)} \subset \mathbb{R}^2$ such that $\text{cone}(r, s)$ is a basis cone. Let $t_1 := r, t_2, \dots, t_{k-1}, t_k := s$ be the elements of $\mathcal{A}^{(1)}$ in $\text{cone}(r, s)$ in cyclic order from r to s . Then for each $i \in \{1, \dots, k - 1\}$, $\text{cone}(t_i, t_{i+1})$ is unimodular. (This is an artifact

of \mathbb{R}^2 . See [35, Corollary 3.11] for a proof.) Hence a Hilbert basis of $\text{cone}(r, s)$ is contained in $\{t_1, \dots, t_k\}$. Thus $A^{(2)} = \mathcal{A}^{(1)}$. \square

When $n > 2$, IBN need not terminate. An example appears in [30]; we independently discovered another as follows.

Example 5.7. Take

$$A = \begin{pmatrix} 0 & 3 & 1 \\ 1 & 1 & 1 \\ 2 & 5 & 5 \\ 1 & 4 & 3 \end{pmatrix}.$$

For each natural number j , set

$$\mathbf{u}_j := (j, 2j + 2, 2j + 1) \text{ and } \mathbf{v}_j := (j, 2j + 1, 2j).$$

Note that $\mathbf{u}_1 = (1, 4, 3)$ is a row of A . To show that IBN does not terminate on \mathcal{A} , one can check the following two assertions. We omit the details.

1. For each j , the unique minimal Hilbert basis of the cone C_j generated by $(0, 3, 1)$, $(1, 1, 1)$, and \mathbf{u}_j includes \mathbf{v}_j .
2. For each j , the unique minimal Hilbert basis of the cone D_j generated by $(0, 3, 1)$, $(2, 5, 5)$, and \mathbf{v}_j includes \mathbf{u}_{j+1} .

For a fixed \mathbf{b} , with $Q_{\mathbf{b}}$ full-dimensional, one can in principle compute SCR of $A\mathbf{x} \leq \mathbf{b}$ by first computing $Q_{\mathbf{b}}^I$ using the Chvátal-Gomory procedure and thus knowing the facet normals of $Q_{\mathbf{b}}^I$. However, just as we do not know a systematic way to compute the Chvátal rank of a matrix A , we also do not know an algorithm to compute SCR of A . Knowing a superset of the normals of inequalities needed to cut out all $Q_{\mathbf{b}}^I$ does not help in computing SCR since there may be vectors in this superset that are never going to be generated by IBN. There are several methods in the literature for

computing such supersets such as the method in [43, Theorem 17.4] or via *atomic fibers* [1]. Although these methods are not helpful in computing SCR, we will see that in many instances one can calculate or bound SCR.

5.3 Contrasting small Chvátal rank with Chvátal rank

In this section we justify our use of the adjective “small” by showing that SCR can be very small even for families of matrices whose Chvátal rank tend to infinity. Note that a lower bound on the SCR or Chvátal rank of a system $A\mathbf{x} \leq \mathbf{b}$ is also a lower bound on the corresponding rank of A . On the other hand, an upper bound on either rank of A is an upper bound on the corresponding rank of $A\mathbf{x} \leq \mathbf{b}$ for any \mathbf{b} .

Proposition 5.8. *If $A \in \mathbb{Z}^{m \times 2}$ then the SCR of A is at most one.*

Proof: This follows immediately from Lemma 5.6. □

Proposition 5.9. *There are systems $A\mathbf{x} \leq \mathbf{b}$ with $A \in \mathbb{Z}^{3 \times 2}$ whose SCRs are one but whose Chvátal ranks are arbitrarily large.*

Proof: Consider the family of inequality systems $A_j\mathbf{x} \leq \mathbf{b}_j$, $j = 1, 2, \dots$ where

$$A_j = \begin{pmatrix} -1 & 0 \\ 1 & 2j \\ 1 & -2j \end{pmatrix} \text{ and } \mathbf{b}_j = (0, 2j, 0)^t.$$

The polyhedron $Q_{\mathbf{b}_j}$ determined by the j th system is a triangle in \mathbb{R}^2 with vertices $(0, 0)$, $(0, 1)$ and $(j, 1/2)$, and its integer hull is the line segment with endpoints $(0, 0)$ and $(0, 1)$. It is noted in [43, §23.3] that the Chvátal rank of $A_j\mathbf{x} \leq \mathbf{b}_j$ is at least j . In contrast, Proposition 5.8 proves that the SCR of any matrix A with only two columns is at most one. For these systems, SCR is one. □

This family can be modified to produce families in higher dimensions.

Theorem 5.10. *For any $n \geq 2$ and $m \geq n + 1$, there are systems $A\mathbf{x} \leq \mathbf{b}$ with $A \in \mathbb{Z}^{m \times n}$ whose SCRs are one but whose Chvátal ranks are arbitrarily large.*

Proof: We prove this by induction using the systems $A_j\mathbf{x} \leq \mathbf{b}_j$ in Proposition 5.9 as base cases. It suffices to show that given a system $A\mathbf{x} \leq \mathbf{b}$, we can extend it in either of the following two ways without changing its Chvátal rank or SCR.

1. Adjoin one new inequality $\mathbf{a}'\mathbf{y} \leq b'$ to the system, thus increasing m by one while fixing n .
2. Adjoin one new row and one new column to A and one new entry to \mathbf{b} , thus increasing m and n by one each.

For the first construction, take $\mathbf{a}'\mathbf{y} \leq b'$ to be any inequality that is satisfied on an open set containing $Q_{\mathbf{b}}$. Such an inequality does not affect $Q_{\mathbf{b}}$, $Q_{\mathbf{b}}^I$, or the running of the Chvátal procedure, so if we had started with $A_j\mathbf{x} \leq \mathbf{b}_j$, the Chvátal rank of the new system stays unchanged. If A' is the matrix formed by adjoining the new vector \mathbf{a}' to A , then since \mathcal{A}' contains \mathcal{A} , it follows from the definition of IBN that $\mathcal{A}'^{(i)}$ contains $\mathcal{A}^{(i)}$ for all i . In particular this holds for $i = 1$, so $\mathcal{A}'^{(1)}$ contains all facet normals of the (unchanged) integer hull $Q_{\mathbf{b}}^I$, leaving SCR also unchanged.

For the second construction, set

$$A^* = \left(\begin{array}{c|c} A & \mathbf{0} \\ \hline \mathbf{0} & 1 \end{array} \right) \in \mathbb{Z}^{(m+1) \times (n+1)}$$

and form \mathbf{b}^* by adjoining any integer z to the end of \mathbf{b} . The polyhedron $Q_{\mathbf{b}^*}$ is just the product of $Q_{\mathbf{b}}$ with the ray $(-\infty, z]$ orthogonal to the hyperplane containing $Q_{\mathbf{b}}$. The same will hold after each iteration of the Chvátal procedure and for the integer hulls, so the Chvátal rank is left unchanged. Furthermore, the new vector $(\mathbf{0}, 1)$ in \mathcal{A}^* cannot contribute to any Hilbert basis constructed during IBN. This is because any linear combination of the vectors of \mathcal{A}^* with multipliers in $[0, 1)$ will have a fractional

last coordinate if the multiplier of $(\mathbf{0}, 1)$ is non-zero. Thus no Hilbert basis element created during IBN involves $(\mathbf{0}, 1)$ and IBN proceeds for \mathcal{A}^* exactly as it did for \mathcal{A} , so the SCR is left unchanged by (2). \square

Our next example also has SCR much less than Chvátal rank and comes from combinatorial optimization. Given a graph $G = (V(G), E(G))$ on n vertices, the *coclique polytope* of G , denoted by $P_{\text{COCL}}(G)$, is the convex hull in \mathbb{R}^n of all incidence vectors of cocliques (i.e., independent sets of vertices) of G . In general, an inequality description of $P_{\text{COCL}}(G)$ is unknown [43, §23.5]. Let $Q(G)$ be the polytope in \mathbb{R}^n given by the inequalities $x_i \geq 0$ for all $i \in V(G)$ and $x_i + x_j \leq 1$ for every edge $\{i, j\} \in E(G)$. Then $P_{\text{COCL}}(G) = Q(G)^I$.

When G is the complete graph K_n , its coclique polytope is just the simplex $\text{conv}\{0, \mathbf{e}_1, \dots, \mathbf{e}_n\}$. This polytope is defined by the inequalities $x_i \geq 0$ for all $i \in [n]$ and $\sum_{i=1}^n x_i \leq 1$. However, $Q(K_n)$ is far larger, and the Chvátal rank of the system defining $Q(K_n)$ is $\mathcal{O}(\log(n))$ [43, Example 23.2]. In contrast, the SCR of the system defining $Q(K_n)$ is small.

Theorem 5.11. *The small Chvátal rank of the system defining $Q(K_n)$ is one if n is odd and two if n is even.*

Proof: The matrix A used to define $Q(K_n)$ consists of the rows $\mathbf{e}_i + \mathbf{e}_j$ for all $1 \leq i < j \leq n$ and $-\mathbf{e}_i$ for all $i \in [n]$. The SCR of this system is the smallest k such that the all-ones vector is an element of $\mathcal{A}^{(k)}$.

If n is odd, form an $n \times n$ non-singular submatrix of A , using the $\mathbf{e}_i + \mathbf{e}_j$ rows, that is the incidence matrix of an n -cycle in G . The all-ones vector is half the sum of the rows of this submatrix and is the only nontrivial integer vector in the open parallelepiped they span, so it is in the Hilbert basis of the cone spanned by these rows. In particular, the all-ones vector is an element of $\mathcal{A}^{(1)}$, and thus the small Chvátal rank is exactly one.

If n is even, the above construction will not work since the all-ones vector is in the semigroup spanned by the rows of the incidence matrix of an n -cycle and hence is not in the Hilbert basis of the cone spanned by these rows. Instead, for each $1 \leq i \leq n$, consider the $n \times n$ non-singular submatrix of A consisting of the row $-\mathbf{e}_i$ and the incidence matrix of an $(n-1)$ -cycle that does not pass through vertex i . The Hilbert basis of the cone spanned by the rows of this submatrix contains the vector with all ones except for a zero in the i th position; weight the rows of the cycle by $1/2$ and $-\mathbf{e}_i$ by zero. Call this vector \mathbf{v}_i . Hence $\mathbf{v}_i \in \mathcal{A}^{(1)}$. The sub-configuration of $\mathcal{A}^{(1)}$ consisting of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ has the all-ones vector in its Hilbert basis, so the all-ones vector appears in $\mathcal{A}^{(2)}$.

It remains to show that when n is even, the all-ones vector does not appear in $\mathcal{A}^{(1)}$. Suppose on the contrary that it does appear. Then there is a basis $\mathcal{U} := \{\mathbf{u}_1, \dots, \mathbf{u}_n\} \subset \mathcal{A}$ and scalars $0 \leq c_1, \dots, c_n < 1$ such that

$$c_1 \mathbf{u}_1 + \dots + c_n \mathbf{u}_n = (1, \dots, 1).$$

But since no component of any vector in \mathcal{A} is more than one and every c_i is strictly less than one, for every i there must be at least two vectors in \mathcal{U} whose i th component is one. In other words, for each vertex, \mathcal{U} contains the incidence vectors of at least two edges covering that vertex. This implies that there are at least n edges contributing to \mathcal{U} and since $|\mathcal{U}| = n$, every $\mathbf{u} \in \mathcal{U}$ comes from an edge and these edges cover every vertex exactly twice. Thus these edges form a set of disjoint cycles covering all the vertices. If there is only one cycle, then since n is even, as we observed earlier, the all-ones vector is not in the Hilbert basis of $\text{cone}(\mathcal{U})$. If there is more than one cycle, then again the all-ones vector is not in the Hilbert basis of $\text{cone}(\mathcal{U})$ for the following reason. For each cycle C (now of size less than n), the all-ones vector with support in the vertices of C is in the cone spanned by the \mathbf{u} 's corresponding to the edges in C . The big all-ones vector with support $[n]$ is the sum of all these all-ones vectors with smaller support since the cycles are disjoint, and hence is not in the Hilbert basis of

$\text{cone}(\mathcal{U})$. □

5.4 Supernormality and small Chvátal rank

A vector configuration \mathcal{A} in \mathbb{Z}^n is **normal** if every integer point in $\text{cone}(\mathcal{A})$ is a non-negative integer combination of \mathcal{A} . Normality means exactly that \mathcal{A} is a Hilbert basis for $\text{cone}(\mathcal{A})$, so has immediate relevance to the Chvátal-Gomory procedure. Hoşten, Maclagan, and Sturmfels generalize normality to *supernormality*, which we show is intimately related to SCR.

Definition 5.12. [30] A configuration \mathcal{A} is **supernormal** if for every subset \mathcal{A}' of \mathcal{A} , every integer point in $\text{cone}(\mathcal{A}')$ is a non-negative integer combination of the set $\mathcal{A} \cap \text{cone}(\mathcal{A}')$.

Proposition 5.13. [30, Proposition 3.1] *For a configuration \mathcal{A} , the following are equivalent.*

1. \mathcal{A} is supernormal.
2. Every triangulation of \mathcal{A} that uses all vectors is unimodular.
3. Every regular triangulation of \mathcal{A} that uses all vectors is unimodular.

Theorem 5.14. *A configuration \mathcal{A} is supernormal if and only if $\mathcal{A} = \mathcal{A}^{(1)}$.*

Proof: The forward direction is immediate from the definitions. For the reverse direction, suppose $\mathcal{A} = \mathcal{A}^{(1)}$ and let T be a triangulation of \mathcal{A} using all vectors. Let σ be a maximal simplex of T . Then the sub-configuration \mathcal{A}_σ is a basis of \mathcal{A} . Since $\mathcal{A} = \mathcal{A}^{(1)}$, \mathcal{A} contains the minimal Hilbert basis of $\text{cone}(\mathcal{A}_\sigma)$. But since all vectors in \mathcal{A} are used in the triangulation T , none can lie inside or on the boundary of $\text{cone}(\mathcal{A}_\sigma)$ except those in \mathcal{A}_σ itself. Thus \mathcal{A}_σ is the Hilbert basis of its own cone. This implies that \mathcal{A}_σ is a lattice basis, so σ is a unimodular simplex. Since σ is arbitrary, it follows that T is a unimodular triangulation, so by Proposition 5.13, \mathcal{A} is supernormal. □

Corollary 5.15. *If $\mathcal{A}^{(k)}$ is supernormal then the SCR of A is at most k . In particular, if \mathcal{A} is supernormal then the SCR of A is zero.*

Thus if \mathcal{A} is supernormal then for each $\mathbf{b} \in \mathbb{Z}^m$, the integer hull $Q_{\mathbf{b}}^I$ equals $\{\mathbf{x} \in \mathbb{R}^m : A\mathbf{x} \leq \mathbf{b}'\}$ for some integer vector \mathbf{b}' . In other words, the integer hulls $Q_{\mathbf{b}}^I$ can be cut out by inequalities whose normals already appear among the rows of A . This is clearly possible if $Q_{\mathbf{b}}^I = Q_{\mathbf{b}}$ for each \mathbf{b} which is exactly when A is *unimodular* (that is, for every subset \mathcal{A}' of \mathcal{A} , every integer point in $\text{cone}(\mathcal{A}')$ is a non-negative integer combination of \mathcal{A}'). However, supernormality is more general than unimodularity. Hosten, Maclagan, and Sturmfels exhibit a four-dimensional configuration that is supernormal but not unimodular and observe in [30]:

It would be interesting to identify infinite families of configurations in higher dimensions which are supernormal but not unimodular. Such families might arise from graph theory or combinatorial optimization.

Proposition 5.16. *There exists an infinite family of configurations of increasing dimension which are supernormal but not unimodular.*

Proof: Let k be any positive integer and let \mathcal{A} be the rows of the matrix A of size $(2k + 1) \times (2k + 1)$:

$$A = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \dots & 1 & 1 \\ 1 & 0 & 0 & \dots & 0 & 1 \end{pmatrix}.$$

That is, A is the edge-vertex incidence matrix of an odd cycle. The determinant of A is two, so there is exactly one Hilbert basis element of $\text{cone}(\mathcal{A})$ that does not generate an extreme ray. In this case, this is the all-ones vector $\mathbf{1}$. So $\mathcal{A}^{(1)} = \mathcal{A} \cup \{\mathbf{1}\}$.

We claim that all maximal minors of $A^{(1)}$ except for $\det(A)$ are ± 1 . This implies that $\mathcal{A}^{(1)}$ equals $\mathcal{A}^{(2)}$, and hence by Theorem 5.14, $\mathcal{A}^{(1)}$ is supernormal. But since \mathcal{A} is not unimodular, neither is $\mathcal{A}^{(1)}$, proving the proposition.

To prove the claim, by symmetry it suffices to check a single minor of $A^{(1)}$ different from $\det(A)$, for instance the one obtained by removing the last row of A from $A^{(1)}$. By cofactor expansion on the last column, this minor equals $-\det(D_1) + \det(D_2)$ where D_1 and D_2 are the $2k \times 2k$ matrices

$$D_1 = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \dots & 1 & 1 \\ 1 & 1 & 1 & \dots & 1 & 1 \end{pmatrix}$$

and

$$D_2 = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \dots & 1 & 1 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{pmatrix}.$$

Now D_1 looks like A except with even size, so its last row is the sum of its 1st, 3rd, \dots , and $(2k - 1)$ st rows, hence $\det(D_1) = 0$. Further, D_2 is upper triangular with 1's on the main diagonal, so $\det(D_2) = 1$. \square

We finish this section by considering the $n = 2$ case.

Proposition 5.17. *If $A \in \mathbb{Z}^{m \times 2}$, then $\mathcal{A}^{(1)}$ is supernormal but not necessarily unimodular. Further, the vectors in $\mathcal{A}^{(1)}$ are exactly the facet normals of the full-dimensional integer polyhedra $Q_{\mathbf{b}}^I$ as \mathbf{b} varies in \mathbb{Z}^m .*

Proof: The first statement follows from Lemma 5.6. Every $Q_{\mathbf{b}}^I$ can be cut out

by inequalities whose normals are in $\mathcal{A}^{(1)}$ which in particular implies that all facet normals of all full-dimensional $Q_{\mathbf{b}}^I$ occur in $\mathcal{A}^{(1)}$.

To finish the proof we need to argue that each $\mathbf{h} \in \mathcal{A}^{(1)}$ is a facet normal of some $Q_{\mathbf{b}}^I$. Suppose \mathbf{h} is in the Hilbert basis of the basis cone $\text{cone}(\mathbf{a}_i, \mathbf{a}_j)$ where \mathbf{a}_i and \mathbf{a}_j are consecutive vectors in the plane. For a vector $\mathbf{p} \in \mathbb{Z}^2$, define $\mathbf{p}^\perp := (-p_2, p_1)$. Set $b_i = 0$ and $b_j = \mathbf{a}_j \cdot \mathbf{h}^\perp$ and all other b_k , $k \neq i, j$ such that the inequalities $\mathbf{a}_k \mathbf{x} \leq b_k$ are very far from the intersection \mathbf{v} of $\mathbf{a}_i \mathbf{x} = 0$ and $\mathbf{a}_j \mathbf{x} = -a_{j1}h_2 + a_{j2}h_1$. We will prove that the line segment with endpoints $(0, 0)$ and \mathbf{h}^\perp is an edge of this $Q_{\mathbf{b}}^I$ and hence that \mathbf{h} is a facet normal of $Q_{\mathbf{b}}^I$.

Note that by construction, both $(0, 0)$ and \mathbf{h}^\perp are on the boundary of $Q_{\mathbf{b}}^I$. Suppose the line segment joining them is not an edge of $Q_{\mathbf{b}}^I$. Then there is a sequence of consecutive lattice points $\mathbf{q}_0 := (0, 0), \mathbf{q}_1, \dots, \mathbf{q}_{t-1}, \mathbf{q}_t := \mathbf{h}^\perp$, with $t \geq 2$, on the boundary of $Q_{\mathbf{b}}^I$ and in the triangle with vertices $(0, 0)$, \mathbf{h}^\perp and \mathbf{v} . Let $\mathbf{p}_j := \mathbf{q}_j - \mathbf{q}_{j-1}$. Then $\mathbf{h}^\perp = (\mathbf{q}_1 - \mathbf{q}_0) + (\mathbf{q}_2 - \mathbf{q}_1) + \dots + (\mathbf{q}_t - \mathbf{q}_{t-1}) = \mathbf{p}_1 + \dots + \mathbf{p}_t$. This implies that $\mathbf{h} := \mathbf{p}_1^\perp + \dots + \mathbf{p}_t^\perp$. By construction, each \mathbf{p}_j^\perp lies in the cone spanned by \mathbf{a}_i and \mathbf{a}_j . Since $\mathbf{p}_j^\perp \in \mathbb{Z}^2$ and $t \geq 2$, we have shown that \mathbf{h} is not in the minimal Hilbert basis of $\text{cone}(\mathbf{a}_i, \mathbf{a}_j)$ which is a contradiction. Therefore, the line segment with endpoints $(0, 0)$ and \mathbf{h}^\perp is an edge of $Q_{\mathbf{b}}^I$. \square

In contrast, when $n = 3$ not every vector in $\mathcal{A}^{(1)}$ need be a facet normal of a $Q_{\mathbf{b}}^I$. See Example 5.24 for such an instance.

Corollary 5.18. *For $A \in \mathbb{Z}^{m \times 2}$, with Δ the maximum absolute value of a (2×2) -minor, there exists a matrix M with at most $m\Delta$ rows such that for each $\mathbf{b} \in \mathbb{Z}^m$, there exists $\mathbf{d} \in \mathbb{Z}^m$ such that $Q_{\mathbf{b}}^I = \{\mathbf{x} \in \mathbb{R}^n : M\mathbf{x} \leq \mathbf{d}\}$.*

Proof: The vector configuration \mathcal{A} , when drawn in the plane, creates a fan with m maximal cones if the fan is complete and $m - 1$ cones if the fan is not complete. The number of Hilbert basis elements of any of these maximal cones is at most $\Delta + 1$. The union of these Hilbert bases is $\mathcal{A}^{(1)}$ and its cardinality is therefore, at most $m\Delta$. \square

5.5 Lower Bounds on SCR

In contrast to the results of Section 5.3, we now exhibit systems of inequalities (and hence matrices) that establish lower bounds on SCR. Recall (Proposition 5.5) that for either an inequality system or the corresponding matrix, the Chvátal rank is an upper bound for the SCR.

Theorem 5.19. *For $m = n = 3$, the small Chvátal rank of $A\mathbf{x} \leq \mathbf{b}$ (and hence A) can be arbitrarily large and can grow exponentially in the size of the input.*

By adjoining inequalities that do not affect Chvátal rank or SCR, we immediately obtain the following.

Corollary 5.20. *The conclusion of Theorem 5.19 holds for any $m, n \in \mathbb{Z}$ satisfying $m \geq n \geq 3$.*

When n is fixed, Chvátal rank is known to grow no faster than exponentially in the size of the input [43, §23.3], so Theorem 5.19 shows that SCR is asymptotically as large as Chvátal rank in the worst case.

To prove Theorem 5.19, we consider the matrices

$$A_j = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & j & 2j - 1 \end{pmatrix}$$

for $j \geq 2$ an integer. We show that the SCR of A_j is $j - 1$ which is exponential in the bit size of A_j . To do this, we explicitly describe the configuration $\mathcal{A}_j^{(k)}$ for all k and prove that $\mathcal{A}_j^{(j-1)} = \mathcal{A}_j^{(j)}$, so the SCR of A_j is at most $j - 1$. Then we prove that the vector $(1, j, j)$ is a facet normal of the integer hull

$$Q_{(0,0,j-1)^t}^I = \text{conv}(\{\mathbf{x} \in \mathbb{Z}^3 : A_j\mathbf{x} \leq (0, 0, j - 1)^t\})$$

and is contained in $\mathcal{A}_j^{(j-1)}$ but not $\mathcal{A}_j^{(j-2)}$. Thus the SCR of A_j is exactly $j - 1$.




Figure 5.1: The polygon R_j^{j-1} (with $j = 6$) used to prove Theorem 5.19. Each integer point (a, b) in the polygon is labeled with the smallest number k such that $(1, a, b)$ appears in $\mathcal{A}_6^{(k)}$.

For $1 \leq k \leq j - 1$, define an integral polygon in \mathbb{R}^2 by

$$R_j^k := \text{conv}\{(0, 0), (k + 1, k + 1), (j, 2j - 1 - k), (j, 2j - 1)\} \subseteq \mathbb{R}_{\geq 0}^2.$$

For $k = j - 1$, the second and third points coincide; in all other cases the four points are distinct and in convex position. An inequality description of R_j^k is

$$R_j^k = \{(x, y) \in \mathbb{R}^2 : x \leq y, 2x \leq y + k + 1, x \leq j, y \leq (\frac{2j - 1}{j})x\}.$$

See Figure 5.1 for an illustration.

Lemma 5.21. [30, Proposition 5.1] *Let R be an integral polygon in \mathbb{R}^2 . The configuration in \mathbb{R}^3 of all vectors $(1, a, b)$ such that (a, b) is an integer point in R is supernormal.*

Lemma 5.22. For $1 \leq k \leq j - 1$, $\mathcal{A}_j^{(k)}$ consists of the vector $(0, 1, 0)$ along with all vectors of the form $(1, a, b)$ where (a, b) is an integer point in R_j^k .

Proof: Induct on k . For $k = 1$, we have

$$R_j^1 = \{(x, y) \in \mathbb{R}^2 : x \leq y, 2x \leq y + 2, x \leq j, y \leq (\frac{2j-1}{j})x\}.$$

Combining the second and fourth of these inequalities gives

$$2x - 2 \leq y \leq 2x - \frac{x}{j}.$$

From the third inequality and the fact that $R_j^1 \subseteq \mathbb{R}_{\geq 0}^2$ we see that $0 \leq \frac{x}{j} \leq 1$. Thus any integer point $(a, b) \in R_j^1$ must have $b = 2a - 2$, $b = 2a - 1$ or $b = 2a$. This implies that $R_j^1 \cap \mathbb{Z}^2$ is contained in

$$\{(i, 2i)\} \cup \{(i, 2i - 1)\} \cup \{(i, 2i - 2)\} \text{ where } 0 \leq i \leq j, i \text{ integer.}$$

By going through each of the three sets in the union, check that

$$R_j^1 = \{(0, 0)\} \cup \{(i, 2i - 1) : 1 \leq i \leq j\} \cup \{(i, 2i - 2) : 2 \leq i \leq j\}.$$

Observe that

$$(1, i, 2i - 1) = (\frac{2j-2i}{2j-1}, \frac{j-i}{2j-1}, \frac{2i-1}{2j-1})A_j$$

for $1 \leq i \leq j$ and that

$$(1, i, 2i - 2) = (\frac{2j-2i+1}{2j-1}, \frac{2j-i}{2j-1}, \frac{2i-2}{2j-1})A_j$$

for $2 \leq i \leq j$, so all the vectors in $\{1\} \times R_j^1 \cap \mathbb{Z}^2$ are in the fundamental parallelepiped of \mathcal{A}_j . Since all the first coordinates are one, no element of $\{1\} \times R_j^1 \cap \mathbb{Z}^2$ is a sum of others. Also, no two elements of $\{1\} \times R_j^1 \cap \mathbb{Z}^2$ differ by a multiple of $(0, 1, 0)$. These facts imply that $\{1\} \times R_j^1 \cap \mathbb{Z}^2$ is in the Hilbert basis of $\text{cone}(\mathcal{A}_j)$, and hence in $\mathcal{A}_j^{(1)}$. Thus $\mathcal{A}_j^{(1)}$ contains $(0, 1, 0)$ and $\{1\} \times R_j^1 \cap \mathbb{Z}^2$.

On the other hand, if $h = c_1(1, 0, 0) + c_2(0, 1, 0) + c_3(1, j, 2j - 1)$ is an integer point in the fundamental parallelepiped of \mathcal{A}_j (so $0 \leq c_1, c_2, c_3 < 1$), then $c_3 = \frac{p}{2j-1}$ for

some integer $1 \leq p \leq 2j - 2$ and c_1 and c_2 are uniquely determined by c_3 , so there is just one such h in the fundamental parallelepiped with a fixed z -coordinate: the vector with that z -coordinate already shown to be in the Hilbert basis.

For the induction step, first assume that $\mathcal{A}_j^{(k-1)}$ contains $\{1\} \times R_j^{k-1} \cap \mathbb{Z}^2$ for some $k \geq 2$. The difference between R_j^{k-1} and R_j^k is that the inequality $2x \leq y + k$ is replaced by $2x \leq y + k + 1$. Using this, we see that our task is to show that the set of new vectors in $\mathcal{A}_j^{(k)}$ includes the following.

$$\{(1, k + i, k + 2i - 1) : 1 \leq i \leq j - k\} \quad (5.1)$$

For each $1 \leq i \leq j - k$, the three vectors $(0, 1, 0)$, $(1, k + i - 1, k + 2i - 2)$, and $(1, k + i, k + 2i)$ appear in $\mathcal{A}_j^{(k-1)}$ by the induction hypothesis. The basis cone they span has normalized volume two and $(1, k + i, k + 2i - 1)$ (half the sum of the three vectors) is the unique integer point in the interior of the fundamental parallelepiped. Thus $(1, k + i, k + 2i - 1)$ appears in the Hilbert basis of the cone, and hence in $\mathcal{A}_j^{(k)}$, so $\mathcal{A}_j^{(k)}$ contains $\{1\} \times R_j^1 \cap \mathbb{Z}^2$.

We now must show that $\mathcal{A}_j^{(k)}$ does not contain any other vectors. Again, assume this is true for $k - 1$. By Lemma 5.21, the previous paragraph, and the induction hypothesis, $\mathcal{A}_j^{(k-1)} \setminus \{(0, 1, 0)\} = R_j^k \cap \mathbb{Z}^2$ is supernormal. Thus the only bases of $\mathcal{A}_j^{(k-1)}$ that might contribute new vectors to $\mathcal{A}_j^{(k)}$ are those that include $(0, 1, 0)$. Any new vector arising this way must be of the form $(1, a, b)$ where $(a - 1, b)$ is strictly in the interior of R_j^{k-1} and (a, b) is outside R_j^{k-1} . From the inequality description of R_j^{k-1} , this vector must indeed be of the form (5.1); see Figure 5.1. Thus no other vectors can occur in $\mathcal{A}_j^{(k)}$. \square

Lemma 5.23. *The configuration $\mathcal{A}_j^{(j-1)}$ is supernormal.*

Proof: Applying the same argument we used to complete the proof of the previous lemma, we conclude that any vector \mathbf{v} in $\mathcal{A}_j^{(j)} \setminus \mathcal{A}_j^{(j-1)}$ would have to be of the form $(1, a, b)$ where $(a - 1, b)$ is an integer point in the interior of R_j^{j-1} and (a, b) is outside

R_j^{j-1} . However, the inequality description of R_j^{j-1} shows it to be a triangle whose right boundary consists only of segments of the line $y = x$ and of the line $x \leq j$, so no such (a, b) exists; see Figure 5.1. Thus $\mathcal{A}_j^{(j)} = \mathcal{A}_j^{(j-1)}$. \square

Proof of Theorem 5.19: By Lemma 5.23 the SCR of \mathcal{A}_j is at most $j - 1$. By Lemma 5.22, the vector $(1, j, j)^t$ appears in $\mathcal{A}_j^{(j-1)}$ but not in $\mathcal{A}_j^{(j-2)}$. So it suffices to show there is a vector \mathbf{b} such that $(1, j, j)^t$ is a facet normal of the integer hull of $\{\mathbf{x} \in \mathbb{R}^3 : A_j \mathbf{x} \leq \mathbf{b}\}$.

Choose $\mathbf{b} = (0, 0, j - 1)^t$. We will prove that the integer hull

$$P_j := \{\mathbf{x} \in \mathbb{R}^3 : A_j \mathbf{x} \leq (0, 0, j - 1)^t\}^I$$

is cut out by the original inequalities $A_j \mathbf{x} \leq (0, 0, j - 1)^t$ and the single new inequality

$$(1, j, j) \mathbf{x} \leq 0 \tag{5.2}$$

and that this new inequality defines a facet of P_j .

Let $\mathbf{y} = (y_1, y_2, y_3)$ be any integer point in $Q_{(0,0,j-1)^t}$. We first show that \mathbf{y} satisfies (5.2). If $y_3 \leq 0$, then since we already know $y_1, y_2 \leq 0$, immediately \mathbf{y} satisfies (5.2). If $y_3 = 1$, then there are four cases to consider:

1. $y_1 = y_2 = 0$: Then $(0, 0, 1) \in Q_{(0,0,j-1)^t}$ and from the third inequality in $A_j \mathbf{x} \leq (0, 0, j - 1)^t$ we would have $j \leq 0$ which is not possible since $j \geq 2$ by assumption. So this case does not occur.
2. $y_1, y_2 < 0$: In this case \mathbf{y} satisfies (5.2).
3. $y_1 = 0, y_2 \leq -1$: Again, \mathbf{y} satisfies (5.2).
4. $y_2 = 0$: In this case, to satisfy the last inequality in $A_j \mathbf{x} \leq (0, 0, j - 1)^t$, $y_1 \leq -j$ and then \mathbf{y} satisfies (5.2).

Finally, suppose $y_3 \geq 2$. Rewrite $x_1 + jx_2 + (2j - 1)x_3 \leq j - 1$ as

$$x_1 + jx_2 \leq (j - 1) - x_3(2j - 1). \quad (5.3)$$

Then

$$\begin{aligned} (1, j, j) \mathbf{y} &= y_1 + jy_2 + jy_3 \\ &\leq (j - 1) - y_3(2j - 1) + jy_3 \\ &= j + y_3 - jy_3 - 1 \\ &= j + y_3(1 - j) - 1 \\ &\leq j + 2(1 - j) - 1 \\ &= 1 - j \\ &< 0 \end{aligned}$$

where the first inequality follows from (5.3), the second from $y_3 \geq 2$, and the last from $j \geq 2$. Thus the inequality (5.2) is valid on all integer points of $Q_{(0,0,j-1)^t}$ and hence is a valid inequality of P_j .

To finish the proof we need to argue that (5.2) is a facet inequality of P_j . This follows from the observation that the three affinely independent integer points $(0, -1, 1)^t$, $(0, 0, 0)^t$, and $(-j, 0, 1)^t$ in P_j satisfy (5.2) with equality. \square

We now use the above family to show that not every vector in $\mathcal{A}^{(k)}$ for $k \leq \text{SCR}(A)$ is needed to cut out the integer hulls $Q_{\mathbf{b}}^I$ as \mathbf{b} varies. In fact, we show that there may be superfluous vectors even in $\mathcal{A}^{(1)}$.

Example 5.24. Consider the matrix A_4 as above. By Lemma 5.22, the vector $(1, 3, 4)$ is an element of $\mathcal{A}_4^{(1)}$. Since A_4 is a non-singular square matrix, the lattice spanned by its columns has finite index in \mathbb{Z}^3 . Specifically, this index is seven, the determinant of A_4 . Thus there are only seven distinct $Q_{\mathbf{b}}$'s up to lattice translation, obtained by choosing one \mathbf{b} from each equivalence class of \mathbb{Z}^3 modulo the lattice. Using the

Chvátal-Gomory procedure one can compute the seven polyhedra $Q_{\mathbf{b}}^I$ explicitly and verify that the vector $(1, 3, 4)$ is not a facet normal of any of them. \square

Our last lower bound for SCR comes from polytopes in the unit cube. A *0/1 polytope* in \mathbb{R}^n is the convex hull of any subset of $\{0, 1\}^n$; that is, an integral polytope contained in the n -dimensional unit cube C_n . Many problems in combinatorial optimization can be phrased in terms of 0/1 polytopes. Accordingly, their Chvátal rank has been specifically studied, and bounds quite different from those that apply to the general case have been found.

Proposition 5.25. [18]

1. *The Chvátal rank of any polytope contained in C_n is at most $n^2(1 + \log n)$.*
2. *There exist polytopes contained in C_n whose Chvátal rank is at least $(1 + \epsilon)n$.*

Since SCR is always bounded above by Chvátal rank, the upper bound in Proposition 5.25 applies to SCR as well. We will derive a lower bound for SCR that is of the same order as that for Chvátal rank. Since our methods are quite different from those in [18], they provide a new way to prove a weaker version of their lower bound.

Theorem 5.26. *There are systems $A\mathbf{x} \leq \mathbf{b}$ that define polytopes contained in C_n whose small Chvátal ranks are at least $n/2 - o(n)$.*

Lemma 5.27. *If \mathbf{v} is in a minimal Hilbert basis of the cone spanned by n vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, then $\|\mathbf{v}\|_{\infty} \leq n \cdot \max_{1 \leq i \leq n} \|\mathbf{v}_i\|_{\infty}$.*

Proof: This is immediate because \mathbf{v} is in the fundamental parallelepiped spanned by $\mathbf{v}_1, \dots, \mathbf{v}_n$; that is, it is a non-negative combination of these n vectors where all coefficients are less than one. \square

Proof of Theorem 5.26: Given any 0/1 polytope Q , we can find a relaxation P contained in C_n and whose facet normals are 0/1/-1 vectors. To do this, set $\mathbf{e}_I = \sum_{i \in I} \mathbf{e}_i$

for each $I \subseteq \{1, \dots, n\}$ where \mathbf{e}_i is the i th standard unit vector in \mathbb{R}^n , and note that the inequality

$$\sum_{i \in I} x_i - \sum_{i \notin I} x_i \leq |I| - 1$$

is violated by \mathbf{e}_I but satisfied by every other vertex of C_n . Thus we can define P by taking such an inequality for every 0/1 vector not in Q , along with the inequalities $0 \leq x_i \leq 1$ that define C_n itself. The normals of all of these inequalities are 0/1/-1 vectors. Let \mathcal{A} be the configuration of all of these normals.

Using a construction by Alon and Vü [2] of 0/1 matrices with large determinant, Ziegler [54, Corollary 26] constructs an n -dimensional 0/1 polytope Q with a (relatively prime integer) facet normal \mathbf{v} whose ∞ -norm is at least $\frac{(n-1)^{(n-1)/2}}{2^{2n+o(n)}}$. Let P be the polytope we constructed above whose integer hull is Q , and let k be the SCR of the system $A\mathbf{x} \leq \mathbf{b}$ defining P . By definition, $\mathbf{v} \in \mathcal{A}^{(k)}$. Since \mathcal{A} consists entirely of 0/1/-1 vectors, we get by inductively applying Lemma 5.27 that

$$n^k > \frac{(n-1)^{(n-1)/2}}{2^{2n+o(n)}}.$$

Taking the logarithm of both sides, we see that

$$\begin{aligned} k \log n &> \left(\frac{n-1}{2}\right) \log(n-1) - (2n+o(n)) \log 2 \\ &= \frac{n}{2} \log(n-1) - \frac{1}{2} \log(n-1) - 2n \log 2 - o(n) \\ &= \frac{n}{2} \log n - o(n \log n) \\ &= \left(\frac{n}{2} - o(n)\right) \log n \end{aligned}$$

so $k > n/2 - o(n)$, as claimed. □

BIBLIOGRAPHY

- [1] W.W. Adams, S. Hoşten, P. Loustau, and L.J. Miller. Gröbner bases over principal ideal domains. *J. Symbolic Computation*, 27(1):31–47, 1999.
- [2] N. Alon and V. Vu. Anti-Hadamard matrices, coin weighing, threshold gates and indecomposable hypergraphs. *J. Combinatorial Theory, Series A*, 79(1):133–160, 1997.
- [3] F. Ardila and C. Klivans. The Bergman complex of a matroid and phylogenetic trees. *Journal of Combinatorial Theory, Series B*, 96:38–49, 2006.
- [4] M.F. Atiyah and I.G. McDonald. *Introduction to Commutative Algebra*. Westview Press, Boulder, CO, 1969.
- [5] G. Bergman. The logarithmic limit set of an algebraic variety. *Trans. Amer. Math. Soc.*, 157:459–469, 1971.
- [6] R. Bieri and J.R.J. Groves. The geometry of the set of characters induced by valuations. *J. Reine Angew. Mathematik*, 347:168–195, 1984.
- [7] L. J. Billera, P. Filliman, and B. Sturmfels. Constructions and complexity of secondary polytopes. *Advances in Mathematics*, 83:155–179, 1990.
- [8] T. Bogart, A.N. Jensen, D. Speyer, B. Sturmfels, and R.R. Thomas. Computing tropical varieties. *Journal of Symbolic Computation*, 42(1-2):54–73, 2007.
- [9] T. Bogart, A.N. Jensen, and R.R. Thomas. The ideal of circuits of a vector configuration. *Journal of Algebra*, 309(2):518–542, 2007.
- [10] B. Buchberger. *On Finding a Vector Space Basis of the Residue Class Ring Modulo a Zero Dimensional Polynomial Ideal*. PhD thesis, University of Innsbruck, Austria, 1965. (German).
- [11] V. Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4:305–337, 1973.
- [12] S. Collart, M. Kalkbrener, and D. Mall. Converting bases with Gröbner walks. *Journal of Symbolic Computation*, 24:465–469, 1997.

- [13] P. Conti and C. Traverso. Buchberger algorithm and integer programming. In H. F. Mattson, T. Mora, and T. R. N. Rao, editors, *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, pages 130–139. Lecture Notes in Computer Science 539, Springer-Verlag, 1991.
- [14] D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties and Algorithms*. Springer-Verlag, New York, 1992.
- [15] P. Diaconis, D. Eisenbud, and B. Sturmfels. Lattice walks and primary decomposition. In B. E. Sagan and R. P. Stanley, editors, *Mathematical Essays in Honor of Gian-Carlo Rota*, pages 173–193, Boston, 1998. Birkhäuser.
- [16] P. Diaconis and B. Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Annals of Statistics*, 26:363–397, 1998.
- [17] M. Einsiedler, M. Kapranov, and D. Lind. Non-Archimedean amoebas and tropical varieties. *Journal für die Reine und Angewandte Mathematik*, 601:139–157, 2006.
- [18] F. Eisenbrand and A.S. Schulz. Bounds on the Chvátal rank of polytopes in the 0/1 cube. *Combinatorica*, 23(2):245–261, 2003.
- [19] D. Eisenbud. *Commutative Algebra with a View Towards Algebraic Geometry*. Springer Graduate Texts in Mathematics, 1994.
- [20] D. Eisenbud and B. Sturmfels. Binomial ideals. *Duke Math. J.*, pages 1–45, 1996.
- [21] T. Granlund et al. Gnu multiple precision arithmetic library 4.1.4. Available from <http://swox.com/gmp/>, 2004.
- [22] G. Ewald. *Combinatorial Convexity and Algebraic Geometry*, volume 168 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1996.
- [23] K. Fukuda. cddlib reference manual, cddlib version 093b. Available from www.ifor.math.ethz.ch/~fukuda/cdd_home/cdd.htm, 2005.
- [24] K. Fukuda, A.N. Jensen, N. Lauritzen, and R.R. Thomas. The generic Gröbner walk. *Journal of Symbolic Computation*, 43(3):298–312, 2007.
- [25] W. Fulton. *Introduction to Toric Varieties*. Princeton University Press, 1993.

- [26] D. Gale. Irreducible convex sets. In *Proc. International Congress of Mathematicians*, volume 2, pages 217–218. Amsterdam, 1954.
- [27] D. Grayson and M. Stillman. Macaulay 2, a software system for research in algebraic geometry. Available from <http://www.math.uiuc.edu/Macaulay2>.
- [28] B. Grünbaum. *Convex Polytopes, Second Edition*. Springer GTM, New York, 2003.
- [29] R. Hartshorne. Connectedness of the Hilbert scheme. *Publ. Math. IHES*, 29:261–304, 1966.
- [30] S. Hoşten, D. Maclagan, and B. Sturmfels. Supernormal vector configurations. *J. Algebraic Combinatorics*, 19 no. 3:297–313, 2004.
- [31] S. Hoşten and B. Sturmfels. GRIN : An implementation of Gröbner bases for integer programming. In E. Balas and J. Clausen, editors, *Integer Programming and Combinatorial Optimization*, pages 267–276. Lecture Notes in Computer Science 920, Springer Verlag, 1995.
- [32] S. Hoşten and R.R. Thomas. The associated primes of initial ideals of lattice ideals. *Mathematical Research Letters*, 6:83–97, 1999.
- [33] S. Hoşten and R.R. Thomas. Gomory integer programs. *Math. Programming, Series B*, 96:271–292, 2003.
- [34] A.N. Jensen. Gfan, a software system for gröbner fans. Available from <http://home.imf.au.dk/ajensen/software/gfan/gfan.html>.
- [35] D. Maclagan and R.R. Thomas. The toric Hilbert scheme of a rank two lattice is smooth and irreducible. *Journal of Combinatorial Theory, Series A*, 2002. to appear, math.AG/0208031.
- [36] P. McMullen. Indecomposable convex polytopes. *Israel Journal of Mathematics*, 58(3):321–323, 1987.
- [37] G. Mikhalkin. Enumerative tropical algebraic geometry in \mathbb{R}^2 . *Journal of the American Mathematical Society*, 18:313–377, 2005.
- [38] T. Mora and L. Robbiano. The Gröbner fan of an ideal. *Journal of Symbolic Computation*, 6:183–208, 1988.

- [39] L. Pachter and B. Sturmfels. Tropical geometry of statistical models. *Proceedings of the National Academy of Sciences*, 101:16132 – 16137, 2004.
- [40] I. Peeva and M. Stillman. Toric Hilbert schemes. *Duke Math. Journal*, 111:419–449, 2002.
- [41] J. Richter-Gebert, B. Sturmfels, and T. Theobald. First steps in tropical geometry. In G.L. Litvinov and V.P. Maslov, editors, *Proc. Conference on Idempotent Mathematics and Mathematical Physics*, pages 289–317, Vienna, 2003. American Mathematical Society, Contemporary Mathematics Vol 377.
- [42] M. Saito, B. Sturmfels, and N. Takayama. *Gröbner Deformations of Hypergeometric Differential Equations*, volume 6. Algorithms and Computation in Mathematics, Springer-Verlag, New-York, 1999.
- [43] A. Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience Series in Discrete Mathematics and Optimization, New York, 1986.
- [44] G.C. Shephard. Decomposable convex polyhedra. *Mathematika*, 10:321–323, 1963.
- [45] D. Speyer and B. Sturmfels. Tropical mathematics. Clay Institute lecture at Park City Utah, 2004, math.CO/0408099.
- [46] D. Speyer and B. Sturmfels. The tropical Grassmannian. *Advances in Geometry*, 4:389–411, 2004.
- [47] B. Sturmfels. *Gröbner Bases and Convex Polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1996.
- [48] B. Sturmfels. *Solving Systems of Polynomial Equations*, volume 97. American Mathematical Society, Providence, RI, 2002.
- [49] B. Sturmfels and R. R. Thomas. Variation of cost functions in integer programming. *Mathematical Programming*, 77:357–387, 1997.
- [50] B. Sturmfels, N. Trung, and W. Vogel. Bounds on projective schemes. *Mathematische Annalen*, 302:417–432, 1995.
- [51] E. Tevelev. Compactifications of subvarieties of algebraic tori. *American Journal of Mathematics*, 2007. to appear, math.AG/0412329.

- [52] T. Theobald. On the frontier of polynomial computations in tropical geometry. *Journal of Symbolic Computation*, 41:1360–1375, 2006.
- [53] G.M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [54] Günter M. Ziegler. Lectures on 0/1-polytopes. In *Polytopes—combinatorics and computation (Oberwolfach, 1997)*, volume 29 of *DMV Sem.*, pages 1–41. Birkhäuser, Basel, 2000.

VITA

Tristram Charles Johann Bogart was born in Monterey, California on 2 February 1980. He grew up in Palo Alto, California and attended Oberlin College, where he graduated with a bachelor of arts degree in mathematics with a minor in history in 2001. He entered the graduate program in mathematics at the University of Washington that autumn and completed his doctorate in 2007 under the direction of Professor Rekha R. Thomas. He enjoys classical music, cycling, ultimate, reading, and card games.