J. Symbolic Computation (1999) **28**, 155–186
Article No. jsco.1998.0271
Available online at http://www.idealibrary.com on IDEAL®

Solving Degenerate Sparse Polynomial Systems Faster

J. MAURICE ROJAS

This paper is dedicated to my son, Victor Lorenzo

Department of Mathematics, City University of Hong Kong,
83 Tat Chee Avenue, Kowloon, Hong Kong

Consider a system F of n polynomial equations in n unknowns, over an algebraically closed field of arbitrary characteristic. We present a fast method to find a point in every irreducible component of the zero set $\mathcal{Z}$ of F. Our techniques allow us to sharpen and lower prior complexity bounds for this problem by fully taking into account the monomial term structure. As a corollary of our development we also obtain new explicit formulae for the exact number of isolated roots of F and the intersection multiplicity of the positive-dimensional part of $\mathcal{Z}$. Finally, we present a combinatorial construction of non-degenerate polynomial systems, with specified monomial term structure and maximally many isolated roots, which may be of independent interest.

1. Introduction

The rebirth of resultants, especially through the **toric**† resultant (Gel'fand *et al.*, 1994), has begun to provide a much needed alternative to Gröbner basis methods for solving polynomial systems. Continuing this philosophy, we will use toric geometry to derive significant speed-ups and extensions of resultant-based methods for solving polynomial systems with infinitely many roots.

The importance of dealing with degenerate polynomial systems has been observed in earlier work on quantifier elimination over algebraically closed fields (Chistov and Grigoriev, 1984; Canny, 1988; Renegar, 1989; Fitchas *et al.*, 1990): Many reasonable algorithms for polynomial system solving fail catastrophically when presented with a system F (of n polynomials in n unknowns) having a positive-dimensional zero set $\mathcal{Z}$. Even worse, this kind of failure can also occur when F has only finitely many roots, if F has infinitely many roots "at infinity". When such failures occur, it is of considerable benefit to the user to at least be given some sort of description of the zero-dimensional part of $\mathcal{Z}$.

We will present two new techniques for handling such degeneracies. The *twisted Chow form* (cf. Main Theorem 2.2) allows one to coordinatize quickly many (but not all) degenerate $\mathcal{Z}$, simply by injecting some extra combinatorics into the classical u-resultant. Our second technique builds on the twisted Chow form and works for **all** degenerate $\mathcal{Z}$: The *toric perturbation* (cf. Main Theorem 2.4) refines and generalizes an earlier algebraic perturbation trick used by Chistov and Grigoriev (1984), Renegar (1989), and Canny (1990).

†Other commonly used prefixes for this modern generalization of the classical resultant (van der Waerden, 1950) include: sparse, mixed, sparse mixed, $\mathcal{A}$-, $(\mathcal{A}_1, \ldots, \mathcal{A}_k)$-, and Newton.

0747-7171/99/070155 + 32 $30.00/0

Our refinement takes sparsity into account and allows one to replace the polynomial degrees present in earlier complexity bounds by more intrinsic geometric parameters (cf. Main Theorems 2.1 and 2.4). We will see in Sections 3.4 and 6 that our bounds are a definite improvement, sometimes even by a factor exponential in n. Our framework also allows us to work over any algebraically closed field (as opposed to some earlier restrictions to the complex numbers) and to isolate the zero-dimensional part of $\mathcal{Z}$.

We also derive four corollaries which may be of independent interest:

(1) An explicit method to compute field extensions involving the roots of F (Corollary 2.1).
(2) An explicit formula for the exact, as opposed to generic, number of isolated† roots of F (Corollaries 2.2 and 2.3).
(3) A combinatorial construction, within polynomial time for fixed n, of F with specified monomial term structure and no roots "at infinity" (Main Theorem 2.3).
(4) A lower bound (conjecturally an exact formula) for the intersection multiplicity of the positive-dimensional part of $\mathcal{Z}$ (Corollary 2.3).

Our main results are stated precisely in Section 2. We then give several simple examples of our main results in Section 3. There we also give an intuitive discussion of roots "at infinity" and show how our results include Canny's earlier **generalized characteristic polynomial (GCP)** as a special case. Section 4 then details our aforementioned combinatorial construction of "generic" F with specified monomial term structure. Our main results are then proved in Section 5, and we discuss the computational complexity of our techniques in Section 6.

2. Summary of Main Results

Before describing our results in detail, we will introduce some necessary notation: in what follows, we will let $\bar{F} := (f_1, \ldots, f_{n+1})$, where for all i, $f_i(x) = \sum_{a \in E_i} c_{i,a} x^a$, E_i is a nonempty finite subset of $(\mathbb{N} \cup \{0\})^n$, and x^a is understood to be the monomial term $x_1^{a_1} \cdots x_n^{a_n}$. Given the $c_{i,a}$, we will be solving for $x := (x_1, \ldots, x_n)$. So the $(n+1)$-tuple $\bar{E} := (E_1, \ldots, E_{n+1})$ thus controls which monomial terms are allowed to appear in our systems of equations. An accepted shorthand is to say that $\bar{F}$ **is an** $(n+1) \times n$ **polynomial system with support contained in** $\bar{E}$. (This generalizes in an obvious way to $k \times n$ systems.)

Of course, our given polynomial systems will usually be $n \times n$, so we will let $F := (f_1, \ldots, f_n)$ and $E := (E_1, \ldots, E_n)$. We also let $\mathrm{Conv}(B)$ denote the convex hull of (i.e., smallest convex set containing) a point set $B \subseteq \mathbb{R}^n$, and let $[k] := \{1, \ldots, k\}$ for any positive integer k. An important geometric invariant for $n \times n$ systems of equations is $\mathcal{M}(E)$ — the **mixed volume** (Burago and Zalgaller, 1988; Schneider, 1994; Gritzmann and Klee, 1993; Emiris and Canny, 1995; Ewald, 1996; Dyer *et al.*, 1998) of the convex hulls of the E_i. For $(n+1) \times n$ systems, we also have the following two important complexity-theoretic parameters: $R(\bar{E}) := \sum_{i=1}^{n+1} \mathcal{M}(E_1, \ldots, E_{i-1}, E_{i+1}, \ldots, E_{n+1})$ and $S(\bar{E}) = \mathcal{O}(\sqrt{n} e^n \mathcal{M}_{\bar{E}}^{\mathrm{ave}})$, where $\mathcal{M}_{\bar{E}}^{\mathrm{ave}}$ is the average value of $\mathcal{M}(\mathcal{E})$ as $\mathcal{E}$ ranges over all n-tuples $(\mathcal{E}_1, \ldots, \mathcal{E}_n)$ with $\mathcal{E}_j \in \{E_1, \ldots, E_{n+1}\}$ for all $j \in [n]$. The true definition of $S(\bar{E})$ depends on the efficiency of a particular class of algorithms described later in Sections 3.2, 5.1, and 6.

†By an isolated root, we will simply mean a root not lying in a positive-dimensional component of $\mathcal{Z}$.

We will usually take all polynomial coefficients to be constants in a fixed algebraically closed field $\mathbb{K}$ or polynomials in $\mathbb{K}[s]$ for some new parameter s. Also, we let $\mathbb{K}^* := \mathbb{K}\setminus\{0\}$ and $\Delta := \mathrm{Conv}(\{\mathbf{O}, \hat{e}_1, \ldots, \hat{e}_n\})$, where $\mathbf{O} \in \mathbb{R}^n$ denotes the origin and $\hat{e}_i \in \mathbb{R}^n$ is the ith standard basis vector. Finally, using $\#$ for set cardinality, let $\varphi_A : (\mathbb{K}^*)^n \longrightarrow \mathbb{P}_{\mathbb{K}}^{\#A-1}$ be the rational map defined by $x \mapsto [x^a | a \in A]$. On occasion, we will extend the domain of φ_A to a suitable toric variety (cf. Section 5).

2.1. FINDING POINTS IN ALL COMPONENTS IN INTRINSIC POLYNOMIAL TIME

Our first main result allows us to use efficiently exact arithmetic to find a point in every irreducible component of $\mathcal{Z}$. In what follows, $\mathcal{O}^*(T)$ means $\mathcal{O}(T\log^r T)$ for some constant $r > 0$.

MAIN THEOREM 2.1. *Let F be an $n\times n$ polynomial system with support contained in E, assume $\mathcal{M}(E) > 0$, and set $E_{n+1} = A = \Delta\cap\mathbb{Z}^n$. Also let $\varphi_A(\mathcal{Z})$ be the zero set† of F in $\mathbb{P}^n_{\mathbb{K}}$. Then we can find univariate polynomials $h, h_1, \ldots, h_n$ with the following properties:*

(0) *The degrees of h and $h_1, \ldots, h_n$ are all bounded above by $\mathcal{M}(E)$.*
(1) *For any root θ of h, define $\gamma(\theta) := (h_1(\theta), \ldots, h_n(\theta))$. Then $\gamma(\theta) \in (\mathbb{K}^*)^n \Longrightarrow \gamma(\theta)$ is a root of F.*
(2) *There is at least one $\gamma(\theta)$ in every irreducible component of $\varphi_A(\mathcal{Z}) \cap (\mathbb{K}^*)^n$. In particular, the set of points $\{\gamma(\theta)\}_{h(\theta)=0}$ is finite and contains all the isolated roots of F in $(\mathbb{K}^*)^n$.*
(3) *Let K be $\mathbb{Q}(c_{i,a} | i \in [n], a \in E_i)$ or $(\mathbb{Z}/p\mathbb{Z})(c_{i,a} | i \in [n], a \in E_i)$, according to whether $\mathrm{char}\,\mathbb{K}$ is zero or a prime p. Then all the coefficients of $h, h_1, \ldots, h_n$ (and all intermediate calculations thereof) are in K, or a degree $\lceil 2\log_p((n+1)\mathcal{M}(E))\rceil$ algebraic extension of K, according as $\mathrm{char}\,\mathbb{K}$ is zero or p.*

Furthermore, we can find $h, h_1, \ldots, h_n$ deterministically within $\mathcal{O}^*(n^4\mathcal{M}(E)^3 R(\bar{E})^2 S(\bar{E})^{2.376})$ arithmetic steps and $\mathcal{O}(nS(\bar{E})^2)$ space. Finally, at the expense of replacing E by $\mathbf{O}\cup E := (\{\mathbf{O}\}\cup E_1, \ldots, \{\mathbf{O}\}\cup E_n)$, we can ensure that $\{\gamma(\theta)\}_{h(\theta)=0}$ includes all the isolated roots of F in $\mathbb{K}^n$ as well.

REMARK 2.1. The above time bound can be reinterpreted as "near-heptic in the number of roots of a system closely related to F" and is clearly polynomial-time for fixed n. Also, depending on the combinatorial data E and the algebraic data $\mathrm{char}\,\mathbb{K}$, the above complexity bounds can be lowered considerably, especially if randomization is allowed. These improvements are detailed further in Section 6. In particular, Main Theorem 2.1 already improves an earlier intrinsic complexity bound due to Giusti *et al.* (1995)‡.

REMARK 2.2. The assumption that $\mathcal{M}(E) > 0$ can actually be checked in polynomial time, via Lemma 4.1 of Section 4. Furthermore, if $\mathcal{M}(E) = 0$, then we can simply add $\leq n$ appropriately chosen points to E (within the same asymptotic time bound) to make $\mathcal{M}(E)$ positive. In particular, one can also use Main Theorem 2.1 to solve $k\times n$ polynomial systems and this is detailed further in Rojas (1999b).

†Zero sets in projective space (and more general toric varieties) are defined in Section 5.

‡It should be noted that Giusti *et al.* (1995) also deals with the more general problem of complexity bounds for polynomial system solving in terms of arithmetic networks and straight-line programs.

As fast algorithms for univariate factoring over algebraically closed fields are already available (Kaltofen, 1992, 1995), our univariate reduction from Main Theorem 2.1 thus yields a fast and general way to find a point in every components of $\mathcal{Z}$. (See also Malajovich-Muñoz and Zubelli (1998) for a fast and numerically-stable univariate factoring method for $\mathbb{K} = \mathbb{C}$.) Our first main theorem thus removes a final geometric/complexity-theoretic bottleneck from solving polynomial systems: earlier algorithms had larger complexity bounds or failed to be general enough.

For example, a fast algorithm for finding approximations within $\varepsilon > 0$ of all the roots of F in $(\mathbb{C}^*)^n$ (within time $\mathcal{O}^*(12^n \mathcal{M}(E)^2 \log\log \frac{1}{\varepsilon})$, neglecting some preprocessing) has recently been announced by Mourrain and Pan (1998). However, their algorithm assumes that $\mathcal{Z}$ is zero-dimensional and $\mathbb{K} = \mathbb{C}$. On the other hand, while the results of Canny (1988, 1990) and Canny *et al.* (1989) yield an algorithm which can handle[†] positive-dimensional $\mathcal{Z}$, one is forced to assume $\mathbb{K} = \mathbb{C}$ in order to obtain a Las Vegas complexity bound of $\mathcal{O}^*\left(nD_\Pi \binom{D_\Sigma + 1}{n}^3\right)$. (We use respectively D_Π and D_Σ for the product and sum of the total degrees of the f_i.) We will see in Sections 3.4 and 6 that our algorithm above is at least this fast, and is in fact frequently much faster. We also point out that when $\mathcal{Z}$ is positive-dimensional, Gröbner basis techniques for solving F suffer from a worst-case arithmetic complexity doubly exponential in n (Mayr and Meyer, 1982).

Main Theorem 2.1 is also useful for certain rationality questions via the following corollary, proved in Section 5.2.

COROLLARY 2.1. *Following the notation of Main Theorem 2.1, suppose now that char* $\mathbb{K} = 0$ *and* F *has only finitely many roots in* $(\mathbb{K}^*)^n$. *Let* g *be the greatest common divisor of* h *and* $\prod_{i=1}^n h_i$. *Then* $K(\zeta_i | (\zeta_1, \ldots, \zeta_n) \in (\mathbb{K}^*)^n$ is a root of $F)$ *is exactly the splitting field of* g.

By combining this corollary with a result of Landau and Miller (1985), it then follows that deciding whether F can be solved in terms of radicals can be done within time (roughly) polynomial in the number of roots of F in $\mathbb{K}^n$. The proof makes use of sparse height bounds (Rojas, 1999b) (analogous to our sparse complexity bounds) and will be pursued in another paper.

To make Main Theorem 2.1 more precise, we now outline its underlying toric geometric techniques.

2.2. MAIN GEOMETRIC RESULTS

First recall that there is a natural addition of point sets in $\mathbb{R}^n$ defined by $B + B' := \{b + b' | b \in B, b' \in B'\}$. In the notation of Rojas (1997a, b), we can associate to any $(n+1)$-tuple of point sets in $\mathbb{Z}^n$, $\bar{E}$, a **toric resultant** $\mathrm{Res}_{\bar{E}}(\bar{F})$. This important operator is amply detailed in Sturmfels (1993, 1994, 1997), Gel'fand *et al.* (1994) and Emiris and Canny (1995), so let us state our first geometric construction.

DEFINITION 2.1. Let $P := \sum_{i=1}^n \mathrm{Conv}(E_i)$ and $\bar{P} := P + \mathrm{Conv}(E_{n+1})$. Also let $A \subset \mathbb{Z}^n$ be any finite subset with at least two points and define $f_{n+1}(x) := \sum_{a \in A} u_a x^a$ and $u := (u_a | a \in A)$, where the u_a are new parameters. We then call $\mathrm{Chow}_A(u) :=$

[†]That is, construct $h, h_1, \ldots, h_n$ as in Main Theorem 1.

$\mathrm{Res}_{(E_\cdot,A)}(F, f_{n+1})$ a **twisted Chow form** of F. (Frequently, we will set $E_{n+1} = A$ and thus $\bar{P} = P + \mathrm{Conv}(A)$ as well.)

Note that $\mathrm{Chow}_A(u)$ will be a polynomial in the parameters u_a, encoding (in a manner to be described below) the roots of F. Twisted Chow forms are a generalization of the classical u-**resultant** (van der Waerden, 1950) as the latter simply corresponds to the case where we use the classical "dense" resultant and let $A = \Delta \cap \mathbb{Z}^n$. For convenience, we will frequently respectively write u_0 and u_i in place of $u_{\mathbf{0}}$ and $u_{\hat{e}_i}$.

EXAMPLE 2.1. Suppose we take $\mathrm{char}\,\mathbb{K} \notin \{2,3\}$, $n = 2$, $E_1 = E_2 = 2\Delta \cap \mathbb{Z}^2$, $A = \Delta \cap \mathbb{Z}^2$, and $F = (1+2y-x^2+y^2, 1+2x+x^2-4y^2)$. Then Chow_A is simply the u-resultant, and this polynomial in u_0, u_1, u_2 factors (modulo a nonzero constant multiple) as $(u_0+\frac{1}{3}u_1-\frac{2}{3}u_2)\times (u_0 + 3u_1 + 2u_2)(u_0 - u_1)^2$. It is also not hard to see that F has exactly three roots: $(\frac{1}{3}, -\frac{2}{3})$, $(3,2)$, and $(-1,0)$; the last occurring with multiplicity 2. Better still, we can read this off directly from our u-resultant by computing $(\frac{\text{coefficient of } u_1}{\text{coefficient of } u_0}, \frac{\text{coefficient of } u_2}{\text{coefficient of } u_0})$ for each linear factor (with u_0 appearing) of the u-resultant. (See Main Theorem 2.2 below.)

Our next main theorem tells us exactly how and when we can use a twisted Chow form to compute monomials in the roots of F. Recall that to any n-dimensional rational polytope $Q \subset \mathbb{R}^n$ one can associate its corresponding **toric variety** (over $\mathbb{K}$) $\mathcal{T}(Q)$ (Kempf *et al.*, 1973; Danilov, 1978; Kapranov *et al.*, 1992; Fulton, 1993; Gel'fand *et al.*, 1994; Rojas, 1999a), and this $\mathcal{T}(Q)$ always has† a naturally embedded copy of $(\mathbb{K}^*)^n$. To state our results fully, we will require some toric variety terminology, but the underlying idea is simple: by working in compactifications more general than the projective spaces $\{\mathbb{P}^n_{\mathbb{K}}\}_{n=1}^{\infty}$, we can make better use of the monomial term structure of our polynomial systems.

MAIN THEOREM 2.2. *Following the notation of Definition 2.1, set $E_{n+1} = A$ and let $\mathcal{Z}$ denote the zero set of F in $\mathcal{T}(\bar{P})$. Then $\mathrm{Chow}_A(u)$ is a homogeneous polynomial, either identically zero or of degree $\mathcal{M}(E)$, with the following properties:*

(1) *The polynomial Chow_A is indentically zero $\iff$ $\varphi_A(\mathcal{Z})$ is positive-dimensional.*
(2) *If $\zeta \in \mathcal{T}(\bar{P})$ is a root of F then Chow_A is divisible by $\sum_{a\in A} \gamma_a u_a$, where $[\gamma_a | a \in A] = \varphi_A(\zeta)$.*
(3) *The polynomial $\mathrm{Chow}_A(u)$ splits completely (over $\mathbb{K}$) into linear factors. In particular, if $\mathrm{Chow}_A \not\equiv 0$ and a nonzero linear form $\sum_{a\in A} \gamma_a u_a$ divides Chow_A, then $[\gamma_a | a \in A] = \varphi_A(\zeta)$ for some root $\zeta \in \mathcal{T}(\bar{P})$ of F.*

The zero set of F in a toric variety is formalized in Section 5. Note in particular that assertions (2) and (3) tell us that calculating $\mathrm{Chow}_A(u)$ allows us to reduce the computation of the projective coordinates $[\zeta^a | a \in A]$, for any root $\zeta \in \mathcal{T}(\bar{P})$ of F, to a multivariate factorization problem. Of course, this reduction only works if $\mathrm{Chow}_A(u)$ is not indentically zero, and assertion (1) tells us exactly when this happens.

We also obtain the following almost immediate corollary.

†It is not always the case that $\mathcal{T}(Q)$ also has a naturally embedded copy of $\mathbb{K}^n$. However, with some extra work, one can modify Q so that this is true.

COROLLARY 2.2. *Following the notation of Main Theorem 2.2, we may check if* Chow_A *is identically zero (and thus whether* $\dim \varphi_A(\mathcal{Z}) > 0$*) within* $\mathcal{O}^*(n^2\mathcal{M}(E)R(\bar{E})S(\bar{E})^{2.376})$ *arithmetic steps and* $\mathcal{O}(nS(\bar{E})^2)$ *space.*[‡] *Furthermore, if* $\mathrm{Chow}_A(u)$ *does not vanish identically, then we can compute the exact number of roots of* F *in* $(\mathbb{K}^*)^n$*, counting multiplicities, within* $\mathcal{O}^*(n^4\mathcal{M}(E)^3R(\bar{E})S(\bar{E})^{2.376})$ *arithmetic steps and* $\mathcal{O}(nS(\bar{E})^2)$ *space.*[7]

Even better, by combining with Corollary 5.1 of Section 5, we can also see how many roots lie at various parts of "toric infinity". Corollary 2.2 thus generalizes Bernshtein's famous mixed volume bound (Bernshtein, 1975) to exact root counting over an algebraically closed field.

However, there is still another improvement to be made: it is actually possible for F to have infinitely many roots in $\mathcal{T}(\bar{P})$ but only finitely many roots in $(\mathbb{K}^*)^n$. In such cases, sometimes the right A will permit an exact count of the roots of F in $(\mathbb{K}^*)^n$ via Corollary 2.2. For example, it is easy to construct F, A, and A' where Chow_A vanishes identically but $\mathrm{Chow}_{A'}$ does not (cf. Section 3.3). On the other hand, those F with infinitely many roots in $(\mathbb{K}^*)^n$ will never have a nontrivial twisted Chow form.

Our next construction works for **all** F and A, and begins as follows:

DEFINITION 2.2. Following the notation of Main Theorem 2.2, assume further that $\mathcal{M}(E) > 0$. Let F^* be any $n \times n$ system with constant coefficients and support contained in E, such that F^* has only finitely many roots in $\mathcal{T}(P)$. We then say that $\mathcal{H}(u;s) := \mathrm{Res}_{(E,A)}(F - sF^*, f_{n+1})$ (where s is a new indeterminate) is a **toric generalized characteristic polynomial for** (F, A). Furthermore, we define $\mathrm{Pert}_{A,F^*}(u) \in K[u_a | a \in A]$ to be the coefficient of the term of $\mathcal{H}(u;s)$ of lowest degree in s. We call Pert_{A,F^*} a **toric perturbation of** (F, A) and, when no confusion is possible, we will sometimes write Pert_A instead.

The polynomial Pert_A is what we can use in place of Chow_A when Chow_A vanishes identically. We will describe this shortly, but first we digress momentarily to describe how to construct the necessary "generic" F^* above: if we simply fix the support of F^* to be E, and pick random numbers for the coefficients (using any probability distribution on $\mathbb{K}^{\#\text{monomial terms}}$ yielding probability 1 avoidance of algebraic hypersurfaces), Lemma 5.3 of Section 5 tells us that F^* will satisfy the above hypothesis with probability 1. Alternatively, a deterministic method for constructing suitable F^* is the following.

DEFINITION 2.3. (ROJAS, 1994; ROJAS AND WANG, 1996) *Given* n*-tuples* $D := (D_1, \ldots, D_n)$ *and* $E := (E_1, \ldots, E_n)$ *of nonempty compact subsets of* $\mathbb{R}^n$*, we say that* D *fills* E *(or* D *is a fill of* E*) iff (0)* $D_i \subseteq E_i$ *for all* $i \in [n]$ *and (1)* $\mathcal{M}(D) = \mathcal{M}(E)$*. We then call* D *irreducible iff the removal of any point of* D *causes* $\mathcal{M}(D)$ *to decrease.*

MAIN THEOREM 2.3. *Following the notation of Definition 2.3, suppose* $E_i \subset \mathbb{Z}^n$ *for all* i*,* $\mathcal{M}(E) > 0$*, and* D *is an irreducible fill of* E*. Then, for any choice of nonzero* $c_{i,a} \in \mathbb{K}^*$*, the polynomial system* $(\sum_{a \in D_1} c_{1,a}x^a, \ldots, \sum_{a \in D_n} c_{n,a}x^a)$ *has exactly* $\mathcal{M}(E)$ *roots, counting multiplicities, in* $(\mathbb{K}^*)^n$ *and no roots in* $\mathcal{T}(P) \setminus (\mathbb{K}^*)^n$*. Furthermore, letting*

[‡]Just as in Main Theorem 1, these complexity bounds can be significantly lowered under certain reasonable assumptions. Also, unless otherwise stated, arithmetic steps will always be counted over the finite extension of K described in Main Theorem 2.1.

roots, counting multiplicities, in $(\mathbb{K}^*)^n$ *and no roots in* $\mathcal{T}(P)\backslash(\mathbb{K}^*)^n$. *Furthermore, letting* $m := \sum_{i=1}^n \#E_i$, *an irreducible fill of* E *can be found within* $\mathcal{O}(n^{2.616}m^{2n+2})$ *arithmetic steps over* $\mathbb{Q}$.

Some simple examples of fills appear in Section 3.1 and we present further background on filling in Section 4. We emphasize that while it is much more practical to pick a generic F^* via randomization, the cost of derandomizing via fills can sometimes be amortized when one solves many F with similar monomial term structure. In particular, the selection of an F^* need only be done once for a given n-tuple E, regardless of the coefficients of F.

Toric perturbations improve on twisted Chow forms as follows:

MAIN THEOREM 2.4. *Following the notation of definition 2,* $\mathrm{Pert}_A(u)$ *is a nonzero homogeneous polynomial of degree* $\mathcal{M}(E)$ *with the following properties:*

(1) $\mathrm{Chow}_A \not\equiv 0 \iff \mathcal{H}(s)$ *has a nonzero constant term. Also, when the latter holds,* $\mathrm{Chow}_A = \mathrm{Pert}_A$.
(2) *If* $\zeta \in \mathcal{T}(\bar{P})$ *is an isolated root of* F *then* Pert_A *is divisible by* $\sum_{a\in A}\gamma_a u_a$, *where* $[\gamma_a | a \in A] = \varphi_A(\zeta)$.
(3) *The polynomial* $\mathrm{Pert}_A(u)$ *splits completely (over* $\mathbb{K}$*) into linear factors. In particular, extending the correspondence of assertion (2), for every irreducible positive-dimensional component* W *of* $\mathcal{Z}$, *there is at least one factor of* Pert_A *corresponding to a root* $\zeta \in W$.

Furthermore, we may evaluate Pert_A *at any point in* $\mathbb{K}^{\#A}$ *within* $\mathcal{O}^*(nR(\bar{E})^2S(\bar{E})^{2.376})$ *arithmetic steps over* $\mathbb{K}$ *and* $\mathcal{O}(nS(\bar{E})^2)$ *space.*†

We emphasize that the main advantage of Pert_A is that we can pick any A we prefer and still obtain a useful analogue of Chow_A. For instance, even if the u-resultant unluckily vanishes identically, we can always simply set $A = \Delta \cap \mathbb{Z}^n$ and directly read off the coordinates of the isolated roots of F from the factors of $\mathrm{Pert}_A(u)$ (assuming one can do multivariate factoring over $\mathbb{K}$). Indeed, $\mathrm{Pert}_{\Delta\cap\mathbb{Z}^n}$ and assertion (3) are central to our construction of points in every irreducible component‡ of $\mathcal{Z}$, not to mention the proof of Main Theorem 2.1.

Better still, we can sometimes (conjecturally always) distinguish which roots of F are isolated.

COROLLARY 2.3. *Following the notation above, let* $\mathcal{Z}_0$ *and* $\mathcal{Z}_\infty$ *respectively denote the zero-dimensional and positive-dimensional parts of* $\mathcal{Z}$. *Then* $\mathcal{Z}_\infty \cap (\mathbb{K}^*)^n = \emptyset \Longrightarrow$ *we can count the number of points in* $\mathcal{Z}_0 \cap (\mathbb{K}^*)^n$, *with or without multiplicity, within the same asymptotic complexity bounds as stated in Main Theorem 2.1. More generally, there is a randomized algorithm which computes upper bounds on the cycle class degrees* $\deg \mathcal{Z}_0$ *and* $\deg \mathcal{Z}_0 \cap (\mathbb{K}^*)^n$, *and a lower bound on* $\deg \mathcal{Z}_\infty$, *within the same complexity bounds. Conjecturally, these bounds are all actually exact formulae with probability* 1.

† Just as in Corollary 2.2 and Main Theorem 1, these complexity bounds can also be significantly lowered under certain reasonable assumptions.

‡ The analogue of assertion (3) had been conjectured for Canny's GCP. We have thus proved this conjecture and generalized it to the toric GCP.

A simple example of this final main result (and Main Theorem 2.4) also appears in Section 3.2. So in summary, as the zero set of F in $\mathcal{T}(\bar{P})$ becomes more and more degenerate, we can successively use Corollaries 2.2 and 2.3 to count roots in $(\mathbb{K}^*)^n$ with complete generality. We also point out that a special case of Corollary 2.3 was used in Rojas (1998a, b, c) in connection with a fast general algorithm for exact multivariate root counting in $(\mathbb{R}^*)^n$.

We can also construct the corresponding analogues of h and the h_i to describe $\mathcal{Z}_0$ explicitly, but this becomes more technical (cf. Section 5.7). The same can be said for the analogous results in $\mathbb{K}^n$, and this is covered in greater depth in Rojas (1997a, b) and Rojas (1998a, b, c). We thus obtain a first step toward an algorithmic foundation for excess intersections. (See Fulton (1984) for a brief historical description of this problem.) In particular, Corollary 2.3 gives a toric geometric algorithm further strengthening Shub's extension (Shub, 1993) of Bézout's theorem over $\mathbb{C}$ (see also Lemma 5.4 of Section 5.6).

We now illustrate our results and theory.

3. Examples

We begin with two small examples of filling. We will then see applications of the toric GCP and twisted Chow form to some degenerate 2×2 and 3×3 polynomial systems. Finally, we will see a brief comparison of the toric GCP to the original GCP. In what follows, we will sometimes respectively write x, y, and z in place of x_1, x_2, and x_3.

3.1. FILLING SQUARES AND CUBES

For our first example, consider the pair of rectangles $\mathcal{P} := ([0,a] \times [0,b], [0,c] \times [0,d])$ where a, b, c, and d are positive integers. Then it is easily verified (via Theorem 4.1 of Section 4) that the pair $D = (\{(0,0),(a,b)\}, \{(0,d),(c,0)\})$ fills $\mathcal{P}$. In this case, the mixed area of both pairs is easily checked to be $ad+bc$. Note also that D is a pair of oppositely slanting diagonals of our initial pair of rectangles (modulo taking convex hulls). Finally, it is easily checked that D is indeed irreducible, since the removal of any point of D results in a mixed area of 0.

By Main Theorem 2.3, we thus obtain that for **any** $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathbb{K}^*$, the bivariate polynomial system $(\alpha_1 + \alpha_2 x^a y^b, \beta_1 x^a + \beta_2 y^b)$ will have exactly $ad + bc$ roots, counting multiplicities, in $(\mathbb{K}^*)^2$.

For our second example, let $\mathcal{P}$ instead be a triple of standard cubes (so that the vertex set of each cube is simply $\{0,1\}^3$). Then, using the criterion from Theorem 4.1 once again, it is easily verified that the triple $D = (\{(1,0,0),(0,1,0),(0,0,1)\}, \{(1,1,0),(1,0,1),(0,1,1)\}, \{(0,0,0),\ (1,1,1)\})$ fills $\mathcal{P}$. (This is depicted in Figure 1.) Also, it is easily checked that the mixed volume of both triples is 6. Finally, note that this D is irreducible as well by Theorem 4.1. Alternatively, one can easily check this by brute force, using any one of the publically web-accessible software packages for mixed volume computation by Emiris, Gao, Huber, or Verschelde.

By Main Theorem 2.3, we thus obtain that for any $\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \gamma_1, \gamma_2 \in \mathbb{K}^*$, the trivariate polynomial system $(\alpha_1 x + \alpha_2 y + \alpha_3 z, \beta_1 xy + \beta_2 xz + \beta_y yz, \gamma_1 + \gamma_2 xyz)$ will have exactly 6 roots, counting multiplicities, in $(\mathbb{K}^*)^3$.

In summary, Theorem 4.1 of Section 4 gives a necessary and sufficient criterion for D to fill a given n-tuple E, and Main Theorem 2.3 tells us that we can construct some irreducible fill for E within time singly exponential in n (and within polynomial time for fixed n).

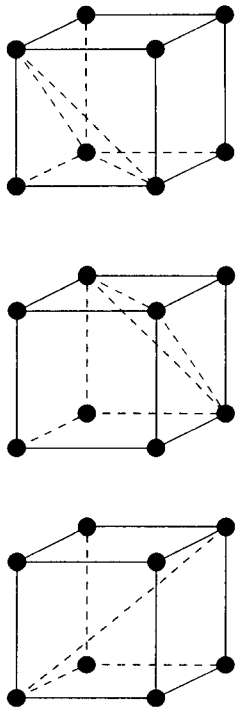

Figure 1. An irreducible fill of three 3-cubes.

3.2. Pert_A APPLIED TO A DEGENERATE 2×2 SYSTEM

Consider the bivariate polynomial system

$$F = (1 + 2x - 2x^2y - 5xy + x^2 + 3x^3y, 2 + 6x - 6x^2y - 11xy + 4x^2 + 5x^3y)$$

over any field of characteristic not equal to 2, 3, or 7. Letting E be the support of F, the reader can easily verify[†] that $\mathcal{M}(E) = 4$, and that the only roots of F are the points $\{(1,1), (\frac{1}{7}, \frac{7}{4})\}$ and the line $\{-1\} \times \mathbb{K}$. So it would appear that the u-resultant (and even $\mathrm{Chow}_{\Delta \cap \mathbb{Z}^2}$) will vanish identically and not give us any useful information about any of these roots. Let us see how we can use Pert_A (with $A = \Delta \cap \mathbb{Z}^2$) to recover everything we need to know about the roots of F.

First, via combinatorial means (Sturmfels, 1993; Emiris and Canny, 1995), we construct a **toric resultant matrix**, $M_{\bar{E}}$. This matrix has the property that its determinant is a multiple of the toric resultant defining the toric GCP (the precursor to Pert_A). With the assistance of a `Matlab` program, `res2.m` (publically available from the author's webpage), we can obtain the following 17×17 matrix:

[†]For $n = 2$, there is the simple formula $\mathcal{M}(E) = \mathrm{Area}(\mathrm{Conv}(E_1 + E_2)) - \mathrm{Area}(\mathrm{Conv}(E_1)) - \mathrm{Area}(\mathrm{Conv}(E_2))$. Also, both polynomials are divisible by $x + 1$. Furthermore, when $\mathrm{char}\,\mathbb{K} = 2$, the second isolated root becomes an isolated root lying on the x-axis.

$$M_{\bar{E}} = \begin{bmatrix}
u_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & u_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & u_0 \\
u_0 & u_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & u_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & u_2 & 0 & 0 & 0 & 0 & 0 & u_1 & 0 & 0 & 0 & 0 & 0 & 0 & u_0 & 0 \\
0 & 0 & 0 & u_2 & 0 & 0 & 0 & 0 & u_0 & u_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & b_3 & b_4 & b_5 & 0 & 0 & 0 & b_2 & 0 & 0 & 0 & 0 & 0 & b_0 & b_1 & 0 \\
0 & 0 & 0 & b_3 & b_4 & b_5 & 0 & 0 & b_1 & b_2 & 0 & 0 & 0 & 0 & 0 & b_0 & 0 \\
0 & 0 & 0 & 0 & b_3 & b_4 & b_5 & 0 & b_0 & b_1 & 0 & 0 & b_2 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & b_3 & b_4 & b_5 & 0 & b_0 & 0 & 0 & b_1 & b_2 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & a_3 & a_4 & a_5 & 0 & a_0 & a_1 & 0 & 0 & a_2 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & a_3 & a_4 & a_5 & 0 & a_0 & 0 & 0 & a_1 & a_2 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & a_4 & a_5 & a_0 & a_1 & 0 & 0 & 0 & a_3 & a_2 \\
a_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & a_3 & a_4 & 0 & a_0 & a_5 & 0 & 0 & 0 & a_1 \\
b_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & b_3 & b_4 & 0 & b_0 & b_5 & 0 & 0 & 0 & b_1 \\
b_1 & b_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & b_3 & 0 & 0 & b_4 & b_5 & 0 & 0 & b_0 \\
0 & 0 & a_3 & a_4 & a_5 & 0 & 0 & 0 & a_2 & 0 & 0 & 0 & 0 & 0 & a_0 & a_1 & 0 \\
0 & 0 & 0 & a_3 & a_4 & a_5 & 0 & 0 & a_1 & a_2 & 0 & 0 & 0 & 0 & 0 & a_0 & 0 \\
a_1 & a_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & a_3 & 0 & 0 & a_4 & a_5 & 0 & 0 & a_0
\end{bmatrix}$$

where the a_i (resp. b_i) are indeterminates correponding to the coefficients of f_1 (resp. f_2). Note in particular that $R(\bar{E}) = 4+4+4 = 12$. As for the other complexity parameter $S(\bar{E})$, its true definition is the size of any available toric resultant matrix. So $S(\bar{E}) = 17$ in the case at hand.

Now note that by Theorem 4.1, $D := (\{\mathbf{O}, (3,1)\}, \{(1,1), (2,0)\})$ is an irreducible fill of E. So by Main Theorem 2.3, we can take $F^* = (1 + x^3y, xy + x^2)$ and apply Main Theorem 2.4 to construct the toric GCP, $\mathcal{H}(u; s)$. By setting $(a_0, \ldots, a_5) = (1 - s, 2, -2, -5, 1, 3-s)$, $(b_0, \ldots, b_5) = (2, 6, -6, -11-s, 4-s, 5)$, and taking the determinant of $M_{\bar{E}}$, we then obtain a nonzero constant multiple of $\mathcal{H}(u_0, u_1, u_2; s)$.

However, multivariate symbolic expansions are typically slow and memory-intensive. So to "solve" efficiently F — that is, to find quickly a point in every irreducible component of its zero set — we will instead compute the univariate polynomials h, h_1, h_2 of Main Theorem 1 via interpolation. The polynomial h is derived simply by specializing Pert_A at some suitable value of (u_1, u_2) and then interpolating through $1+\mathcal{M}(E)$ values of u_0. The derivation of h_1 and h_2 is essentially the same but involves an additional intermediate step described in Section 5.1. As Pert_A is in turn a coefficient of $\mathcal{H}(u; s)$, there is also another level of interpolation through $1 + S(\bar{E}) - \mathcal{M}(E)$ values of s.

For example, setting $(u_1, u_2) = (\frac{1}{2}, 1)$ (and setting u_0 equal to a parameter t), we easily obtain via `Maple` that

$$h(t) = -153 + 120t + 1540t^2 + 1600t^3 + 448t^4$$

$$h_1(t) = -\frac{11762}{7511} + \frac{19150}{22533}t + \frac{114736}{22533}t^2 + \frac{7264}{3219}t^3$$

$$h_2(t) = -\frac{5881}{7511} + \frac{32108}{22533}t + \frac{57368}{22533}t^2 + \frac{3632}{3219}t^3.$$

As $h(t)$ factors as $(2t+3)(28t+51)(2t+1)(4t-1)$, we thus immediately obtain from Corollary 2.1 (and the fact that u_1 and u_2 were chosen within K) that the zero-dimensional part of $\mathcal{Z} \cap (\mathbb{K}^*)^2$ actually lies in $(K^*)^2$, where K is the quotient field generated by the canonical image of $\mathbb{Z}$ in $\mathbb{K}$. Furthermore, by Main Theorem 1 (and the fact that $\mathbf{O} \in E_1 \cap E_2$), we can at last recover a set of points lying in $\mathcal{Z}$ (including all the isolated roots of F in $\mathbb{K}^2$) by substituting $\{-\frac{3}{2}, -\frac{51}{28}, -\frac{1}{2}, \frac{1}{4}\}$ into the pair $(h_1(t), h_2(t))$. In particular, we obtain the set $\{(1,1), (\frac{1}{7}, \frac{7}{4}), (-1, 1), (-1, \frac{1}{4})\}$.

For the curious, we can easily compute via **Maple** that the full expansion of $\mathcal{H}(u;s)$ is, up to a nonzero constant multiple,

$$
\begin{aligned}
&(u_2^4 - u_0^4 + u_1^4 + 6u_1^2u_2^2 - 4u_1u_2^3 - 4u_1^3u_2)\mathbf{s^8}\\
&\quad +(36u_1^2u_2^2 - 20u_2u_0^3 - 20u_2^3u_0 - 4u_1u_0^3 - 19u_0^4 - 24u_2^4 + 6u_0^2u_1u_2\\
&\quad +36u_1u_2^3 + 36u_1^4 - 12u_0u_1^2u_2 - 9u_1^2u_0^2 + 3u_2^2u_0^2 + 36u_0u_1u_2^2 - 4u_0u_1^3 - 84u_1^3u_2)\mathbf{s^7}\\
&\quad +(220u_2^4 - 170u_2u_0^3 - 394u_1^3u_2 - 98u_1u_0^3 - 98u_0^2u_1u_2 - 20u_0^4 + 370u_2^3u_0\\
&\quad +14u_0u_1u_2^2 - 110u_0u_1^3 - 226u_1^2u_0^2 - 354u_1^2u_2^2 + 454u_1^4 - 274u_0u_1^2u_2 + 74u_1u_2^3)\mathbf{s^6}\\
&\quad +(1008u_2u_0^3 - 1612u_0^2u_1u_2 + 903u_0^4 - 624u_1u_0^3 - 2632u_2^3u_0 - 2104u_0u_1^2u_2 - 970u_2^4\\
&\quad -1010u_1u_2^3 + 418u_1^3u_2 - 2104u_0u_1u_2^2 - 642u_1^2u_2^2 - 1547u_1^2u_0^2 - 936u_0u_1^3 - 1557u_2^2u_0^2 + 2204u_1^4)\mathbf{s^5}\\
&\quad +(538u_0^2u_1u_2 + 1271u_0^4 + 12253u_2^2u_0^2 + 6972u_2u_0^3 + 1929u_1^4 - 3075u_1^2u_2^2 + 654u_0u_1u_2^2\\
&\quad +50u_1u_0^3 + 2156u_2^4 - 960u_1^2u_0^2 - 2290u_0u_1^3 + 132u_1u_2^3 - 5344u_0u_1^2u_2 - 1142u_1^3u_2 + 8708u_2^3u_0)\mathbf{s^4}\\
&\quad +(4384u_1u_0^3 - 24988u_2^2u_0^2 - 1582u_1^3u_2 - 6756u_0^4 + 10884u_0u_1u_2^2 + 3802u_1u_2^3 + 15438u_0u_1^2u_2\\
&\quad +1024u_0u_1^3 + 8324u_1^2u_0^2 - 12826u_2^3u_0 + 11270u_0^2u_1u_2 - 6976u_1^4\\
&\quad +7164u_1^2u_2^2 - 21326u_2u_0^3 - 2408u_2^4)\mathbf{s^3}\\
&\quad +(3436u_1^3u_2 + 3800u_0u_1^3 + 7756u_2^3u_0 - 3886u_1u_2^3 + 1225u_2^4 + 17059u_2^2u_0^2 - 5984u_1^2u_0^2\\
&\quad +15708u_2u_0^3 - 12232u_0u_1u_2^2 + 5180u_0^4 - 2091u_1^2u_2^2 - 6828u_0u_1^2u_2\\
&\quad +1316u_1^4 - 12700u_0^2u_1u_2 - 4312u_1u_0^3)\mathbf{s^2}\\
&\quad +(384u_0^3u_1 - 1792u_0^3u_2 + 512u_0^2u_1^2 + 1536u_0^2u_1u_2 + 1920u_0u_1u_2^2 - 1288u_0u_2^3 - 768u_1^3u_2\\
&\quad -448u_0^4 - 2436u_0^2u_2^2 - 384u_0u_1^3 + 1024u_0u_1^2u_2 - 64u_1^4 + 260u_1^2u_2^2 + 768u_1u_2^3 - 196u_2^4)\mathbf{s}.
\end{aligned}
$$

So our toric perturbation Pert_{A,F^*} is just the coefficient of s or s^2 in this polynomial, as $\mathrm{char}\,\mathbb{K} \neq 2$ or $\mathrm{char}\,\mathbb{K} = 2$. Let us now examine Pert_{A,F^*} itself in detail: factoring with **Maple**, we obtain that Pert_{A,F^*} splits as follows:

$$-4(u_0 + u_1 + u_2)(28u_0 + 4u_1 + 49u_2)(u_0 - u_1 + u_2)(4u_0 - 4u_1 + u_2).$$

In particular, given any factor above, the ratio of the coefficients of u_i and u_0 is precisely the ith coordinate of some corresponding root of F. Thus the first two factors correspond precisely to the two isolated roots we already know. As for the last two factors, note that they both give isolated points lying on the aforementioned line $\{-1\} \times \mathbb{K}$. We can then guess that this line should be assigned an excess intersection multiplicity of 2. Of course, we might not know at the outset which of these roots is isolated, i.e., a zero-dimensional component of $\mathcal{Z}$. However, as the constant term of $\mathcal{H}(s)$ vanishes, assertions (1) of Main Theorems 2.2 and 2.4 at least tell us that $\mathcal{Z}$ is indeed positive-dimensional.

To distinguish the isolated roots, let us employ an algorithm from the proof of Corollary 2.3: apply Main Theorem 2.3 once more to pick $F^{**} = (1 + x^3y, xy + 2x^2)$. Noting that (due to their second equations) F^* and F^{**} will have no roots in common in $(\mathbb{K}^*)^2$, let us then define the **double toric perturbation**, Pert_A^{**}, to be the greatest common divisor of Pert_{A,F^*} and $\mathrm{Pert}_{A,F^{**}}$.

Repeating the same calculation we used for h, h_1, h_2, but with Pert_A^{**} instead, we obtain new polynomials $h^{**}, h_1^{**}, h_2^{**}$. Let us compute the gcd, g^{**}, of h^{**} and $h_1^{**}h_2^{**}$. It then turns out that the number of isolated roots of F is at most $\deg h^{**} - \deg g^{**}$ (cf. Section 5.7).

More explicitly, via **Maple** again, we easily see that $h^{**}(t) = (2t + 1)(2t + 3)$ and $g^{**}(t) = 1$. So the number of isolated roots in $(\mathbb{K}^*)^2$ is at most 2, and the positive-dimensional part of $\mathcal{Z}$ (the line $\{-1\} \times \mathbb{K}$) should be assigned an intersection multiplicity of at least $\mathcal{M}(E) - 2 = 2$. Fortuitously (conjecturally always), our lower bound is actually an equality.

For completeness, we now reveal $\mathrm{Pert}_{A,F^{**}}$ (up to a constant multiple):

$$(u_0+u_1+u_2)(28u_0+4u_1+49u_2)\left(u_0-u_1+\frac{\frac{1}{\sqrt{-3}}-1}{4}u_2\right)\left(u_0-u_1-\frac{\frac{1}{\sqrt{-3}}+1}{4}u_2\right).$$

(In particular, $\mathrm{Pert}_{A,F^{**}}$ is again the coefficient of s in $\mathcal{H}(s)$.) Note also that the last two factors of this toric perturbation again correspond to roots lying on the line $\{-1\}\times\mathbb{K}$. We thus see that varying the coefficients of our perturbation of F has moved two of our points lying in the positive-dimensional part of $\mathcal{Z}$.

Note (via `Maple` again) that the original GCP could have been used above, but would have resulted in a variant of Pert_A of degree 16 (the product of the degrees of f_1 and f_2) — four times larger than the degree of our Pert_A. Also, the old GCP is significantly larger, having 672 terms, compared with 110 for our above toric GCP $\mathcal{H}(u;s)$.

3.3. WHICH COMPACTIFICATION FOR Chow_A?

Here we show how the twisted Chow form Chow_A can vanish identically for the wrong A, thus giving no information about the roots of F. Along the way, we will also obtain a more precise visualization of the toric compacta $\mathbb{P}^3_{\mathbb{K}}$, $\mathcal{T}(P)$, and $\mathcal{T}(\bar{P})$. We also point out that while it is sometimes customary to consider the roots of F in $\mathcal{T}(P)$ (as in Fulton (1993), Gel'fand *et al.* (1994) and Rojas (1999a)), the construction of Chow_A and Pert_A necessitate the consideration of roots in $\mathcal{T}(\bar{P})$ as well.

To define our next example, set $n=3$, $A=\Delta\cap\mathbb{Z}^3$, and consider the 3×3 system $F=(a_1yz+a_2xz+a_3xy+a_4xyz, b_1yz+b_2xz+b_3xy+b_4xyz, c_1yz+c_2xz+c_3xy+c_4xyz)$. Note that the mixed volume bound for this system is 1. Furthermore, it is clear that $\frac{1}{xyz}F$ is a linear system in $\{\frac{1}{x},\frac{1}{y},\frac{1}{z}\}$. So by Cramer's rule, we can express x, y, and z as ratios of 3×3 determinants in the coefficients.

Combining this with the product formula for toric resultants (Pedersen and Sturmfels, 1993) (and clearing denominators) we obtain that Chow_A is precisely† $[423][143][124]u_0+[123][143][124]u_1+[123][423][124]u_2+[123][423][143]u_3$ where the **bracket** $[ijk]$ (Dalbec and Sturmfels, 1995) is the 3×3 subdeterminant

$$\det\begin{bmatrix} a_i & a_j & a_k\\ b_i & b_j & b_k\\ c_i & c_j & c_k\end{bmatrix}$$

of the coefficient matrix of F. This compactly expressed resultant can be thought of as a **semi-mixed** Chow form — a toric resultant of a system of $n+1$ polynomials with $k\leq n$ distinct supports.

Now consider the specialization of F to $(yz+xz+2xy+3xyz, yz+xz+4xy+9xyz, yz+xz+8xy+27xyz)$. It is then easily verified that F has no roots in $(\mathbb{K}^*)^3$, but F does have exactly one root † in $\mathcal{T}(P)$. Also, in our particular example, $\mathcal{T}(P)\cong\mathbb{P}^3_{\mathbb{K}}$ and, locally (within $(\mathbb{K}^*)^n$), the isomorphism is given by $(x,y,z)\mapsto[\frac{1}{x}:\frac{1}{y}:\frac{1}{z}:1]$. In particular, using the latter set of coordinates, our one root of F in $\mathcal{T}(P)$ is exactly the point $[1:-1:0:0]$. More to the point, $\mathrm{Chow}_A\equiv 0$ for this specialization of the coefficients of F.

†We also need the fact that the Pedersen–Sturmfels formula, originally stated only over $\mathbb{C}$, remains true over a general algebraically closed field (cf. Section 5.3).

†If $\mathrm{char}\,\mathbb{K}\in\{2,3\}$ then F will actually have infinitely many roots in $\mathcal{T}(P)$. So let us assume henceforth that $\mathrm{char}\,\mathbb{K}\notin\{2,3\}$. (It is easy to construct similar examples when $\mathrm{char}\,\mathbb{K}\in\{2,3\}$ as well.)

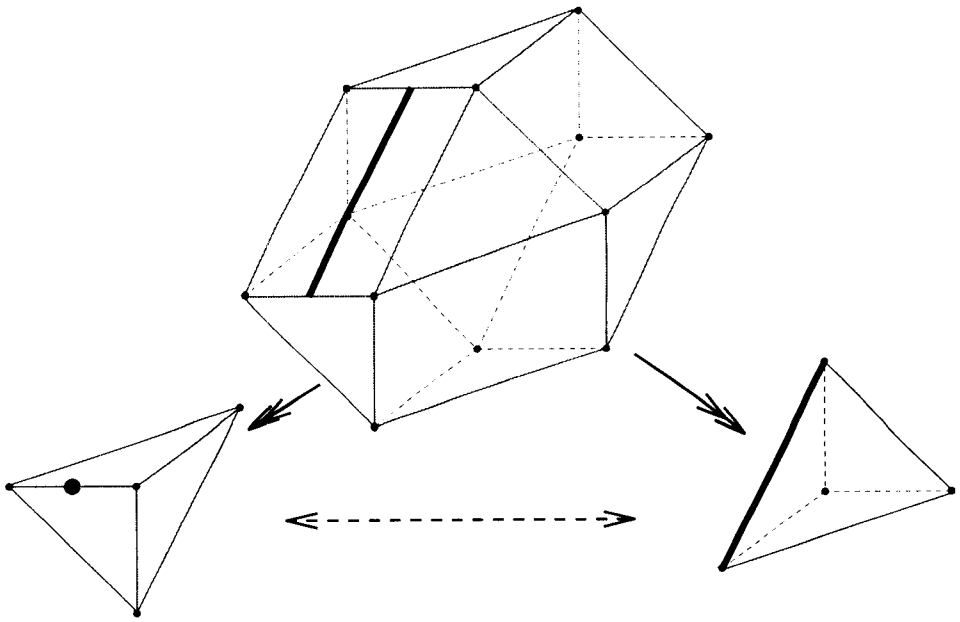

Figure 2. One root in the lower left toric compactification ($\mathcal{T}(P)$) becomes infinitely many roots in the other two compactifications ($\mathcal{T}(\bar{P})$ and $\mathcal{T}(A)$).

A simple geometric explanation for this behavior of $\text{Chow}_\star(\cdot)$ is that the choice of A defines a toric variety $\mathcal{T}(A)$ into which the roots of F in $\mathcal{T}(\bar{P})$ are projected. (The variety $\mathcal{T}(A)$ is the toric variety corresponding to a point set (Gel'fand *et al.*, 1994), and is simply the image of $\mathcal{T}(\bar{P})$ under the morphism φ_A.) So depending on our choice of A, the roots of F in $\mathcal{T}(P)$ may or may not correspond to roots of F in $\mathcal{T}(A)$ in a well-defined way. For instance, in our example, F actually has infinitely many roots in $\mathcal{T}(A)$, so Main Theorem 2.2 tells us that Chow_A must vanish.

So it is more useful to work within $\mathcal{T}(\bar{P})$, as the roots of F in $\mathcal{T}(P)$ and $\mathcal{T}(A)$ are actually images of the roots of F in $\mathcal{T}(\bar{P})$. In particular, the underlying algebraic maps induce projections of certain faces of $\bar{P}$ (corresponding to certain parts of $\mathcal{T}(\bar{P}) \setminus (\mathbb{K}^*)^n$) onto certain faces of P and $\text{Conv}(A)$. Figure 2 above illustrates this, along with where the root $[1 : -1 : 0 : 0] \in \mathcal{T}(P)$ of F "goes" within these various compacta. For instance, note that $\bar{P}$ is a cuboctahedron, and φ_A is constant on the portions of $\mathcal{T}(\bar{P}) \setminus (\mathbb{K}^*)^n$ corresponding to the triangular faces with inner normals $-\hat{e}_1$, $-\hat{e}_2$, $-\hat{e}_3$, and $(1,1,1)$.

Algebraically, we have the following maps:

$$\begin{array}{ccccc} & & \mathcal{T}(\bar{P}) & & \\ & \overset{\pi}{\swarrow} & & \overset{\varphi_A}{\searrow} & \\ \mathbb{P}^3_{\mathbb{K}} \cong \mathcal{T}(P) & & \overset{\phi}{\dashleftarrow\!\!\!\dashrightarrow} & & \mathcal{T}(A) \hookrightarrow \mathbb{P}^3_{\mathbb{K}} \end{array}$$

where π is the natural projection between compatible toric compacta (cf. Section 5), and ϕ is the rational map (defined just on $(\mathbb{K}^*)^n$) from $\mathcal{T}(P)$ to $\mathcal{T}(A)$ obtained from

$x \mapsto [x^a | a \in A]$. In the case at hand, the latter map is simply the identity map between the two corresponding naturally embedded copies of $(\mathbb{K}^*)^3$.

To remedy the preceding trivial Chow_A, we can instead use $\mathrm{Chow}_{A'}(u)$ with $A' := \{(0,1,1),(1,0,1),(1,1,0),(1,1,1)\}$. (This choice is motivated by trying to pick an A' which is compatible with P (cf. Section 5).) In particular, when the coefficients of F are unspecialized,

$$\mathrm{Chow}_{A'}(u) = \det \begin{bmatrix} a_1 & a_2 & a_3 & a_4 \\ b_1 & b_2 & b_3 & b_4 \\ c_1 & c_2 & c_3 & c_4 \\ u_{(0,1,1)} & u_{(1,0,1)} & u_{(1,1,0)} & u_{(1,1,1)} \end{bmatrix}.$$

So under our last specialization, this becomes $12u_{(1,0,1)} - 12u_{(0,1,1)}$. Note that we now recover our root $[1 : -1 : 0 : 0]$ from the coordinates of our new twisted Chow form. For example, the ratio of the x-coordinate to the y-coordinate is just $\frac{x}{y} = \frac{x^1y^0z^1}{x^0y^1z^1} = \frac{12}{-12} = -1$.

Alternatively, we can simply use Pert_A and forget about cleverly chosen A'. For example, by Main Theorem 3 (and Theorem 4.1), we can simply take $F^* = (yz + xyz, xz + xyz, xy + xyz)$. After an application of `Maple`, we then obtain that $\mathrm{Pert}_{\Delta \cap \mathbb{Z}^3}$ is exactly $5u_1 + 21u_2$. In particular, while the point $[0 : 5 : 21 : 0] \in \mathcal{T}(A)$ does not correspond (in any obvious way) to a root of F in $\mathcal{T}(P)$, it is the image of a bona fide root of F in $\mathcal{T}(\bar{P})$ under the morphism φ_A.

In closing, we emphasize that in practice we would never actually compute the full monomial expansions of $\mathrm{Chow}_A(u)$, $\mathrm{Pert}_A(u)$, or $\mathcal{H}(u;s)$ — we would instead recover the roots of F (or evaluate monomials thereof) via rapid and sophisticated interpolation techniques, e.g. Canny (1988, 1990), Canny *et al.* (1989) and Díaz and Kaltofen (1995). In particular, this is the approach of Main Theorem 1, and our calculations can be sped up tremendously with suitably optimized code.

3.4. THE "DENSE" CASE

Our last example illustrates a simple fundamental case.

Suppose E is the n-tuple $(d_1\Delta \cap \mathbb{Z}^n, \ldots, d_n\Delta \cap \mathbb{Z}^n)$ where $d_i \in \mathbb{N}$ for all i. (So we are now considering the family of all $n \times n$ polynomial systems where f_i has total degree $\leq d_i$ for all i.) This is usually referred to as the dense case. It is then easily verified that the system $F^* = (x_1^{d_1}, \ldots, x_n^{d_n})$ (with support contained in E) has only finitely many roots in $\mathcal{T}(P)$. Indeed, in this case, $\mathcal{T}(P) \cong \mathbb{P}^n_{\mathbb{K}}$ and there is exactly one root (of multiplicity $\prod d_i$) at the origin $\mathbf{O}$. Note also that our current setting is sufficiently simple that we could find a suitable F^* with just n terms, without the need for an irreducible fill.

REMARK 3.1. Letting $D_\Pi := \prod_{i=1}^n d_i$ and $D_\Sigma := \sum_{i=1}^n d_i$ it is easily checked by the basic properties of the mixed volume that for general E we have

$$\mathcal{M}(E) \leq D_\Pi, R(\bar{E}) \leq D_\Pi(1 + \sum_{i=1}^n \frac{1}{d_i}), \text{ and } S(\bar{E}) \leq \binom{D_\Sigma + 1}{n},$$

where d_i is the degree of f_i for all i. (The last inequality follows from Macaulay's 19th century construction of the multivariate resultant (Canny, 1987).) Furthermore, equality occurs for all three bounds in the dense case. When these upper limits on $\mathcal{M}$, R, and S are reached, our complexity bound from Main Theorem 2.1 then specializes to the

best bounds from Canny (1988, 1990) and Canny *et al.* (1989), once $\operatorname{char}\mathbb{K} = 0$ and randomization is allowed (cf. Corollary 6.1 of Section 6).

Letting $A = \Delta \cap \mathbb{Z}^n$, we then see that our polynomial $\mathcal{H}(u; s)$ is simply the original GCP (Canny, 1990), but extended to a general algebraically closed field. In particular, our $F - sF^*$ is the polynomial system $(f_1 - sx^{d_1}, \ldots, f_n - sx^{d_n})$. (Note also that if we set $d_1 = \cdots = d_n = 1$ then $\mathcal{H}(a_{n+1,0} - \lambda, a_{n+1,1}, \ldots, a_{n+1,n}; \lambda)$ is just the usual characteristic polynomial of a matrix.) Finally, note that $\mathcal{T}(A) \cong \mathcal{T}(\bar{P}) \cong \mathcal{T}(P) \cong \mathbb{P}^n_{\mathbb{K}}$ and the map φ_A is the identity. So by considering the zero set of F in $\mathcal{T}(\bar{P})$, in the dense case, we are just considering the zero set of F in $\mathbb{P}^n_{\mathbb{K}}$ in the usual way via homogenizations. Thus by Main Theorem 2.4, Canny's original GCP indeed finds a point in every irreducible component of $\mathcal{Z}$ in $\mathbb{P}^n_{\mathbb{K}}$, as conjectured in 1990. Of course, the advantage of the toric GCP is that we can do the same with greater efficiency for sparse systems with small $\mathcal{M}(E)$.

4. Filling

Here we briefly recount filling and some related concepts. Some of the material below is covered at greater length in Rojas (1994). The results below form the basis for our combinatorial approach to perturbing degenerate polynomial systems.

Let $\mathcal{S}^{n-1} \subset \mathbb{R}^n$ denote the unit $(n-1)$-sphere centered at the origin. For any compact $B \subset \mathbb{R}^n$ and any $w \in \mathbb{R}^n$, define B^w to be the set of $x \in B$ where the inner-product $x \cdot w$ is minimized. (Thus B^w is the intersection of B with its supporting hyperplane in the direction w.) We then define $E^w := (E_1^w, \ldots, E_n^w)$ and $D \cap E^w := (D_1 \cap E_1^w, \ldots, D_n \cap E_n^w)$.

Recall that the dimension of any $B \subseteq \mathbb{R}^n$, $\dim B$, is the dimension of the smallest subspace of $\mathbb{R}^n$ containing a translate of B. The following definition is fundamental to our development.

DEFINITION 4.1. Suppose $C := (C_1, \ldots, C_n)$ is an n-tuple of polytopes in $\mathbb{R}^n$ or an n-tuple of finite subsets of $\mathbb{R}^n$. We will allow any C_i to be empty and say that a nonempty subset $J \subseteq [n]$ is essential for C (or C has essential subset J) $\iff$ (0) $C_i \neq \emptyset$ for all $i \in J$, (1) $\dim(\sum_{j \in J} C_j) = \#J - 1$, and (2) $\dim(\sum_{j \in J'} C_j) \geq \#J'$ for all nonempty proper $J' \subsetneq J$.

Equivalently, J is essential for $C \iff$ the $\#J$-dimensional mixed volume of $(C_j | j \in J)$ is 0 and no smaller subset of J has this property. Figure 3 below shows some simple examples of essential subsets for C, for various C in the case $n = 2$.

A basic fact about mixed volumes is that $\mathcal{M}(E) = 0 \iff E$ has an essential subset, whenever $\operatorname{Supp}(E) = [n]$. However, there is an even deeper connection between filling and essentiality:

THEOREM 4.1. (ROJAS 1994, SECTION 2.5) Suppose D and E are n-tuples of finite subsets of $\mathbb{Z}^n$ such that $\mathcal{M}(E) > 0$. Then D fills $E \iff$ for all $w \in \mathcal{S}^{n-1}$, $\operatorname{Supp}(D \cap E^w)$ contains a subset essential for E^w.

REMARK 4.1. One certainly need not check infinitely many w. In fact, we need only check one w (just pick any inner normal) for each face of the polytope $P = \sum_{i=1}^n \operatorname{Conv}(E_i)$.

We also present the following important observation.

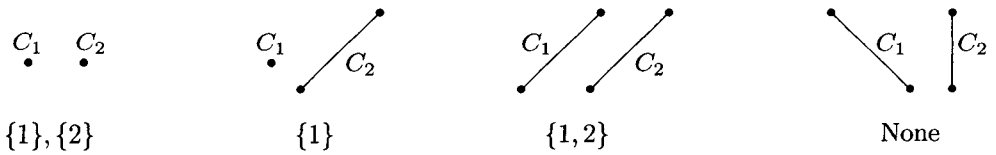

Figure 3. The essential subsets for four different pairs of plane polygons. (The segments in the third pair are meant to be parallel.)

LEMMA 4.1. *Let* $m := \sum_{i=1}^{n} \#E_i$. *Then we can decide* $\mathcal{M}(E) \stackrel{?}{>} 0$ *within* $\mathcal{O}(mn^{1.616})$ *arithmetic steps over* $\mathbb{Q}$. *Furthermore, if* $\mathcal{M}(E) = 0$, *then we can find points* $p_1, \ldots, p_n \in (\mathbb{N} \cup \{0\})^n$, *within the same asymptotic complexity bound, such that* $\mathcal{M}(\{p_1\} \cup E_1, \ldots, \{p_n\} \cup E_n) > 0$.

The first portion was stated in terms of a non-explicit polynomial-time bound in Dyer *et al.* (1998, Theorem 8). To the best of our knowledge, Lemma 4.1 gives the first precise complexity bound for the above geometric problem, so we will supply a proof.

PROOF OF LEMMA 4.1. By the translation invariance of the mixed volume, we may assume that $\mathbf{O} \in E_i$ for all i. Let $\mathcal{I}$ be the set of all pairs (i,j) with $i \in [n]$ and $j \in \{0, \ldots, \#E_i - 1\}$. Also let $\mathcal{B}$ be the $n \times m$ matrix whose columns are the elements of all the E_i. We will let $\mathcal{I}$ index the columns of $\mathcal{B}$ in the most obvious way. Finally, let $\mathcal{G}$ be the bipartite graph with vertex set $\mathcal{I}$ where (i,j) and (i',j') are connected by an edge iff $i = i'$ and either $j = 0$ or $j' = 0$.

There are two natural matroid structures on $\mathcal{I}$: the linear matroid and the partition matroid (Grötschel *et al.*, 1993, Section 7.5). The independent sets of the first (resp. second) matroid are exactly the index sets defining linearly independent multisets of columns of $\mathcal{B}$ (resp. matchings in $\mathcal{G}$). It is a simple corollary of Dyer *et al.* (1998, Proposition 2) that the mixed volume is nonzero iff we can find linearly independent vectors $a_1 \in E_1$, $\ldots$, $a_n \in E_n$. So it suffices to know if there is a set $I \subseteq \mathcal{I}$ which is a **basis** for both matroids. In other words, we have reduced our vanishing volume problem to determining whether there is an I which simultaneously defines n linearly independent columns of $\mathcal{B}$ and a matching in $\mathcal{G}$. This is an instance of the unweighted intersection problem for linear matroids, and an $\mathcal{O}(mn^{1+1/(4-\omega)} \log n)$ algorithm for this problem appears in Gabow and Xu (1996). (We use ω (<2.376) for the famous matrix multiplication complexity exponent (Coppersmith and Winograd, 1990).) So we have proved the first portion of our lemma.

As for the second portion of the lemma, there are many ways to find $p_1, \ldots, p_n$: One naive but valid way is simply to set $p_i := \hat{e}_i$ for all i. This takes only $\mathcal{O}(n)$ operations. □

REMARK 4.2. In the semi-mixed case — that is, when there are only $n' < n$ distinct E_i — we can easily alter the above argument to replace m by m', where $m' := \sum_{j=1}^{n'} \#E_{i_j}$ and $\{E_i | i \in [n]\} = \{E_{i_j} | j \in [n']\}$.

Oddly enough, filling seems to have originated from an algebraic problem: genericity conditions for counting the roots of sparse polynomial systems. This aspect is explored

much further in Rojas (1994, 1999a), Rojas and Wang (1996). We also emphasize that constructing a fill need only be done once for a given family of problems, provided E remains fixed. The situation where the monomial term structure of a polynomial system remains fixed once and for all, and the coefficients may vary many thousands of times, actually occurs frequently in many practical contexts such as robot control or computational geometry.

To conclude our background, we will need the following lemma characterizing irreducible fills.

LEMMA 4.2. *Following the preceding notation, assume $\mathcal{M}(D) > 0$. Then D is irreducible $\Longleftrightarrow$ for any v lying in some D_i, there exists a $w \in \mathbb{Q}^n \setminus \{\mathbf{O}\}$ such that $D_i^w = \{v\}$ and $\mathcal{M}(D_1^w, \ldots, D_{i-1}^w, D_{i+1}^w, \ldots, D_n^w) > 0$.*

PROOF. First note that the mixed volume condition above is equivalent to $\{i\}$ being the unique essential subset of D^w. This follows immediately from definition 4.1 and, say, the development of Burago and Zalgaller (1988).

The "$\Longleftarrow$" direction then follows almost immediately from Theorem 4.1: if the mixed volume condition holds, then the removal of any point from D would indeed violate the filling condition from Theorem 4.1. So the removal of any point from D would make $\mathcal{M}(D)$ decrease. The converse implication follows almost as easily.

Suppose, to derive a contradiction, that D is irrreducible but there is some v in some D_i satisfying the following property: for all $w \in \mathbb{Q}^n \setminus \{\mathbf{O}\}$, $\#D_i^w \geq 2$ or $\mathcal{M}(D_1^w, \ldots, D_{i-1}^w, D_{i+1}^w, \ldots, D_n^w) = 0$. Let us then consider the n-tuple $D' := (D_1, \ldots, D_{i-1}, D_i \setminus \{v\}, D_{i+1}, \ldots, D_n)$. Then by Theorem 4.1 once again, D' fills D. But this contradicts the irreducibility of D, so we are done. □

5. Toric Geometry and the Proofs of Main Theorems

Our notation is a slight variation of that used in Fulton (1993), and is described at greater length in Rojas (1999a). However, we will briefly review a few important facts and definitions.

The (inner) normal fan of a polytope $Q \subset \mathbb{R}^n$, $\mathrm{Fan}(Q)$, is simply the collection of cones of inner normals of faces of Q (Gel'fand *et al.*, 1994). (For instance, the inner normal fan of the standard unit square in the plane consists of nine cones: the four quadrants, the four nonnegative coordinate rays, and the origin.) We will assume the reader to be familiar with the construction of a toric variety from a fan, a polytope, or a finite point set (Fulton, 1993; Gel'fand *et al.*, 1994).

EXAMPLE 5.1. When $A = \Delta \cap \mathbb{Z}^n$, it is easy to derive from scratch that $\mathcal{T}(A)$ is just the projective space $\mathbb{P}^n_{\mathbb{K}}$. More generally, if $\mathrm{Conv}(A)$ is a product of simplices, then $\mathcal{T}(A)$ is a product of twisted projective spaces (Fulton, 1993) — hence our appelation for Chow_A. Note also that the coefficients of Chow_A are multisymmetric functions of $\{\varphi_A(\zeta)\}_\zeta$ as ζ ranges over the roots of F in $\mathcal{T}(\bar{P})$.

Recall that any n-dimensional toric variety $\mathcal{T}$ over $\mathbb{K}$ has a $(\mathbb{K}^*)^n$-action extending the natural action of $(\mathbb{K}^*)^n$ on itself. Let us now list our cast of main characters:

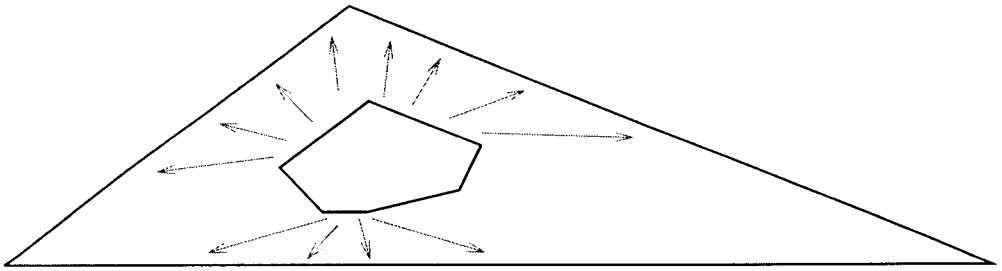

Figure 4. The inner polytope is compatible with the outer polytope. Also, the corresponding "outer" toric variety can be obtained as a deformation (or image under a proper morphism) of the "inner" toric variety.

DEFINITION 5.1. (FULTON, 1993; ROJAS, 1999A) Given any $w \in \mathbb{R}^n$, we will use the following notation:

$$\begin{aligned}
T &= \textit{The algebraic torus } (\mathbb{K}^*)^n.\\
Q^w &= \textit{The face of } Q \textit{ with inner normal } w.\\
\sigma_w &= \textit{The closure of the cone generated by the inner normals of } Q^w.\\
\sigma_w^\vee &= \textit{The dual (or angle) cone } \{w' \in \mathbb{R}^n | w' \cdot y \geq 0 \textit{ for all } y \in \sigma_w\}.\\
U_w &= \textit{The affine chart of } \mathcal{T}(Q) \textit{ corresponding to all semigroup homorphisms}^\dagger\ \sigma_w^\vee \cap \mathbb{Z}^n \longrightarrow \mathbb{K}.\\
O_w &= \textit{The } T\textit{-orbit corresponding to the relative interior of } Q^w.\\
\mathcal{E}_Q(Q') &= \textit{The } T\textit{-invariant Weil divisor of } \mathcal{T}(Q) \textit{ corresponding to a polytope } Q'.\\
\mathrm{Div}(f) &= \textit{The Weil divisor of } \mathcal{T}(Q) \textit{ defined by a rational function } f \textit{ on } (\mathbb{K}^*)^n.\\
\mathcal{D}_Q(f,Q') &= \mathrm{Div}(f) + \mathcal{E}_Q(Q') = \textit{The toric effective divisor of } \mathcal{T}(Q) \textit{ corresponding to } (f,Q').\\
\mathcal{D}_Q(F,\mathcal{P}) &= \textit{The (nonnegative) cycle in the Chow ring of } \mathcal{T}(Q) \textit{ defined by } \textstyle\bigcap_{i=1}^k \mathcal{D}_Q(f_i,P_i), \textit{ whenever } \mathcal{P} = (P_1,\ldots,P_k).
\end{aligned}$$

We say that P is compatible with Q iff every cone of $\mathrm{Fan}(Q)$ is a union of cones of $\mathrm{Fan}(P)$ (Khovanskii, 1977; Fulton, 1993; Rojas, 1999a). (So P compatible with $Q \Longrightarrow P$ has at least as many facets as Q.)

Finally, whenever F is a $k \times n$ polynomial system with support contained in E, we will define the zero set[‡] of F in $\mathcal{T}(Q)$ to be the toric cycle $\mathcal{D}_Q(F,\mathcal{P})$, where $\mathcal{P} := (\mathrm{Conv}(E_1),\ldots,\mathrm{Conv}(E_k))$. **Toric infinity** is then defined relative to Q: it is simply the set $\mathcal{T}(Q) \setminus (\mathbb{K}^*)^n$.

EXAMPLE 5.2. (ZERO SETS IN $\mathbb{P}^n_{\mathbb{K}}$) Suppose $Q = \alpha + \beta\Delta$, for any $\alpha \in \mathbb{Q}^n$ and any rational $\beta > 0$. Then $\mathcal{T}(Q) \cong \mathbb{P}^n_{\mathbb{K}}$ canonically. As for explicitly defining the zero set of F in $\mathcal{T}(Q)$, we can do the following: (1) Define vectors $p_1,\ldots,p_n \in \mathbb{Z}^n$ such that for all i, $x^{p_i} f_i \in \mathbb{K}[x]$ is not divisible by any x_j, (2) define $\tilde{f}_i(x) := x_\infty^{d_i} x^{p_i} f(\frac{x_1}{x_\infty},\ldots,\frac{x_n}{x_\infty})$ for all i, where d_i is the total degree of $x^{p_i} f_i$. Then $[z_1 : \cdots : z_n : z_\infty] \in \mathbb{P}^n_{\mathbb{K}}$ is a root of F iff $\tilde{f}_1(z) = \cdots = \tilde{f}_n(z) = 0$. In particular, note that this toric definition differs from the

[†]Note that the domain and range spaces are respectively semigroups under the natural operations of vector addition and field multiplication.

[‡]When necessary, we will also use the underlying scheme structure.

classical definition of "zero set of F in $\mathbb{P}^n_{\mathbb{K}}$", due to the extra step (1). For instance, our toric definition might omit some affine roots, for certain E and F. However, note that step (1) is unnecessary when $\mathbf{O} \in E_i$ for all i.

REMARK 5.1. By Rojas (1999a, Section 6.1), the zero scheme of F in $\mathbb{K}^n$ embeds naturally in $\mathcal{D}_P(F,\mathcal{P})$ (and $\mathcal{D}_{\bar{P}}(F,\mathcal{P})$) when we replace $\bar{E}$ by $\mathbf{O} \cup \bar{E}$. Hence the introduction of $\mathbf{O} \cup E$ in (and $A = \Delta \cap \mathbb{Z}^n$ in the proof of) Main Theorem 1.

The following result will provide some necessary geometric intuition for specializing resultants. The lemma immediately following then gives a more explicit algebraic analogy between the faces of Q and the affine charts of $\mathcal{T}(Q)$.

VANISHING THEOREM FOR RESULTANTS. (ROJAS, 1998B) *Suppose $\bar{F}$ is an $(n+1) \times n$ polynomial system (over $\mathbb{K}$) with support contained in $\bar{E}$. Then, provided $\mathcal{M}(E_1, \ldots, E_{i-1}, E_{i+1}, \ldots, E_{n+1}) > 0$ for some $i \in [n+1]$, $\mathrm{Res}_{\bar{E}}(\bar{F}) = 0 \iff \mathcal{D}_{\bar{P}}(F,\mathcal{P}) \neq \emptyset$, where $\mathcal{P} := (\mathrm{Conv}(E_1), \ldots, \mathrm{Conv}(E_{n+1}))$ and $\bar{P} = \sum_{i=1}^{n+1} \mathrm{Conv}(E_i)$.*

LEMMA 5.1. (ROJAS, 1999A, SECTION 4.2–5.1) *Suppose F is a $k \times n$ polynomial system over $\mathbb{K}$ with support contained in a k-tuple of integral polytopes $\mathcal{P} := (P_1, \ldots, P_k)$ in $\mathbb{R}^n$. Assume further that Q is a rational polytope in $\mathbb{R}^n$. Then the defining ideal in $\mathbb{K}[x^a | a \in \sigma_w^\vee \cap \mathbb{Z}^n]$ of $U_w \cap \mathcal{D}_Q(F,\mathcal{P})$ is $\langle x^{b_i} f_i |$ for all $i \in [k]$ and $b_i \in \mathbb{Z}^n$ such that $b_i + P_i \subseteq \sigma_w^\vee \rangle$.* □

Lifting (or projecting) from one toric variety to another is an important fundamental idea we will also use. The following lemma follows directly from the development of Fulton (1993).

LEMMA 5.2. *Suppose $Q \subset \mathbb{R}^n$ is an n-dimensional rational polytope, and B is either a nonempty finite subset of $\mathbb{Z}^n$ or a rational polytope in $\mathbb{R}^n$. Assume further that Q is compatible with $\mathrm{Conv}(B)$. Then there is a natural (surjective) proper morphism $\pi : \mathcal{T}(Q) \twoheadrightarrow \mathcal{T}(B)$. In particular, $\pi(\mathcal{D}_Q(F,\mathcal{P})) = \mathcal{D}_B(F,\mathcal{P})$, where the latter cycle is the image of $\mathcal{D}_{\mathrm{Conv}(B)}(F,\mathcal{P})$ under the natural proper morphism from $\mathcal{T}(\mathrm{Conv}(B))$ to $\mathcal{T}(B)$. Furthermore, $\pi(O_w) = O_w$, where the corresponding T-orbits are considered in their respective domains, and $\pi|_{(\mathbb{K}^*)^n} = \mathrm{id}$.*

REMARK 5.2. Following the notation of Main Theorem 1, it easily follows that if $A = \Delta \cap \mathbb{Z}^n$ then the multiplicity of any root of F in $(\mathbb{K}^*)^n$ is preserved under the map φ_A. If $A = dQ \cap \mathbb{Z}^n$ for some rational polytope compatible with P, and $d \in \mathbb{N}$ is sufficiently large, then the same will be true of any root of F in $\mathcal{T}(\bar{P})$ (Fulton, 1993). In general, thanks to the functoriality of Chow forms (Dalbec and Sturmfels, 1995), Chow_A is precisely the Chow form of the subscheme $\varphi_A(\mathcal{Z})$ of $\mathbb{P}^{\#A-1}_{\mathbb{K}}$.

Another immediate corollary of our last lemma is the following result on the meaning of the projective coordinates $[\zeta^a | a \in A]$.

COROLLARY 5.1. *Following the notation of Main Theorem 4, let $\zeta \in \mathcal{T}(\bar{P})$ be an isolated root of F and fix a vertex of $\mathrm{v} \in \mathrm{Conv}(A)$ with inner normal w. Then $\varphi_A(\zeta)$ lies in the*

affine chart U_w *of* $\mathcal{T}(A) \iff$ *the coefficient of* $u_{\mathbf{v}}$ *in the corresponding factor of* Pert_A *is nonzero.*

EXAMPLE 5.3. Suppose we take $A = \Delta \cap \mathbb{Z}^n$ as usual. Then $\mathcal{T}(A) \cong \mathbb{P}^n_{\mathbb{K}}$ canonically, and there are exactly $n+1$ affine charts of $\mathcal{T}(A)$ corresponding to vertices. These charts are respectively isomorphic to $\mathbb{P}^n_{\mathbb{K}}$ minus the hyperplane at infinity, and $\mathbb{P}^n_{\mathbb{K}} \setminus \{x_i = 0\}$ as i runs through $[n]$. For example, given a factor of Pert_A such as $u_0 + u_3$, we know that it corresponds to a root image $\varphi_A(\zeta)$ which lies in $n-2$ of these affine charts and outside of 2 others, i.e., $\varphi_A(\zeta) = [0:0:1:0:\cdots:0:1]$ lies on the x_3-axis. Similarly, if all the coordinates of $\varphi_A(\zeta)$ are nonzero, then $\zeta, \varphi_A(\zeta) \in (\mathbb{K}^*)^n$.

Finally, we will need a version of the fundamental fact that F generically has exactly $\mathcal{M}(E)$ roots in $(\mathbb{K}^*)^n$. The case $\mathbb{K} = \mathbb{C}$ first appeared in Bernshtein (1975), and the general case is an immediate corollary of Rojas (1999a, Main Theorems 1 and 2).

LEMMA 5.3. *Let* C_E *be the vector of coefficients of* F *and define* $\#E := \sum_{i=1}^n \#E_i$. *Then there is an algebraic hypersurface* $\Sigma_E \subset \mathbb{K}^{\#E}$ *such that* $C \in \mathbb{K}^{\#E} \setminus \Sigma_E \Longrightarrow F$ *has no roots in* $\mathcal{T}(P) \setminus (\mathbb{K}^*)^n$. *Moreoever, the latter assertion implies that* F *has exactly* $\mathcal{M}(E)$ *roots, counting multiplicities, in* $(\mathbb{K}^*)^n$.

With all our technical background complete, we can now prove our main theorems.

5.1. POLYNOMIAL ALGEBRA AND THE FIRST HALF OF MAIN THEOREM 1

Our proof of assertions (0)–(2) of Main Theorem 1 will rely on two main constructions: the toric perturbation $\mathrm{Pert}_{\Delta\cap\mathbb{Z}^n}$ and an extension of Canny's constructive version (Canny, 1988) of the primitive element theorem. We thus emphasize that while Chow_A and Pert_A permit one to reduce polynomial system solving to multivariate factorization, we will not use factoring to build h and $h_1, \ldots, h_n$.

Algebraically, the idea is as follows: our techniques allow us to find a set of points $Z' \subset (\mathbb{K}^*)^n$ intersecting every irreducible component of the zero set of F in $(\mathbb{K}^*)^n$. Consider the field extension $L := K(Z')$, obtained by adjoining all the coordinates of all the points of Z'. Then L is a finite extension of K, and by the primitive element theorem (van der Waerden, 1950), $L = K(\theta)$ for some $\theta \in L$. Furthermore, by the same theorem, we should be able to recover the coordinates of every point in Z' in terms of rational functions (with coefficients in K) of θ. As $K(\theta) \cong K[t]/h(t)$ when h is the minimal polynomial of θ over K, we can further simplify the preceding rational representation to one in terms of polynomials in θ with coefficients in K. Our algorithm for Main Theorem 2.1 will explicitly construct this encoding for us.

To describe our algorithm, we will first need a bit of subresultant theory: for any univariate polynomials $f(t) = \alpha_0 + \alpha_1 t + \cdots + \alpha_{d_1} t^{d_1}$ and $g(t) = \beta_0 + \beta_1 t + \cdots + \beta_{d_2} t^{d_2}$,

consider the following $(d_1+d_2-2)\times(d_1+d_2-1)$ matrix

$$\begin{bmatrix}
\beta_0 & \cdots & \beta_{d_2} & 0 & \cdots & 0 & 0 \\
0 & \beta_0 & \cdots & \beta_{d_2} & 0 & \cdots & 0 \\
\vdots & \ddots & \ddots & & \ddots & \ddots & \vdots \\
0 & \cdots & 0 & \beta_0 & \cdots & \beta_{d_2} & 0 \\
0 & 0 & \cdots & 0 & \beta_0 & \cdots & \beta_{d_2} \\
\alpha_0 & \cdots & \alpha_{d_1} & 0 & \cdots & 0 & 0 \\
0 & \alpha_0 & \cdots & \alpha_{d_1} & 0 & \cdots & 0 \\
\vdots & \ddots & \ddots & & \ddots & \ddots & \vdots \\
0 & \cdots & 0 & \alpha_0 & \cdots & \alpha_{d_1} & 0 \\
0 & 0 & \cdots & 0 & \alpha_0 & \cdots & \alpha_{d_1}
\end{bmatrix}$$

with d_1-1 "β rows" and d_2-1 "α rows". Let M_1^1 (resp. M_0^1) be the submatrix obtained by deleting the last (resp. second to last) column, and let $\mathcal{R}_i(f,g) := \det(M_i^1)$ for $i \in \{0,1\}$. Finally, define the first subresultant of f and g to be $\mathcal{R}_0(f,g) + \mathcal{R}_1(f,g)t$. It is then a classical fact that if $\gcd(f,g) = a+bt$ with $b \neq 0$, then $\frac{a}{b} = \frac{\mathcal{R}_1(f,g)}{\mathcal{R}_0(f,g)}$ (González-Vega, 1991). We will make heavy use of this fact in our proof.

Recall also the following algorithmic facts about polynomials over any field (Bini and Pan, 1994):

(a) Given the values of a univariate polynomial of degree d at $d+1$ distinct points, the coefficients of the polynomial can be recovered within $\mathcal{O}^*(d)$ field operations.
(b) The gcd of two univariate polynomials of degree $\mathcal{O}(d)$ can be found within $\mathcal{O}^*(d)$ field operations.
(c) The coefficients of the square-free part of a univariate polynomial (of degree d) can be found within $\mathcal{O}^*(d)$ field operations.
(d) The subresultant of two univariate polynomials of degree $\mathcal{O}(d)$ can be computed within $\mathcal{O}^*(d)$ additions and multiplications.

We now proceed with our proof of the first half of Main Theorem 2.1.

PROOF OF ASSERTIONS (0)–(2). To simplify matters slightly, we will first derive a Las Vegas version of our algorithm for Main Theorem 2.1. The announced time bound will then follow from a simple derandomization. The construction of $h, h_1, \ldots, h_n$ will follow from evaluating $\mathrm{Pert}_A(u)$ at various specializations of u, thus reducing to $\mathcal{O}(n)$ univariate polynomial interpolation and gcd problems. In particular, our algorithm can be outlined as follows:

Step 0 Set $A = \Delta \cap \mathbb{Z}^n$ and fix generic values in $\mathbb{K}$ for $u_1, \ldots, u_n$.
Step 1 Define $h \in \mathbb{K}[t]$ to be $\mathrm{Pert}_A(t, u_1, \ldots, u_n)$.
Step 2 If $n = 1$, set $h_1(\theta) := \theta$ and stop. Otherwise, for all $i \in [n]$, let $q_i^-(t)$ be the square-free part of $\mathrm{Pert}_A(t, u_1, \ldots, u_{i-1}, u_i - 1, u_{i+1}, \ldots, u_n)$.
Step 3 Let α satisfy either $\alpha = 1$ or $\alpha(\alpha+1) = 1$ according as $\mathrm{char}\,\mathbb{K} \neq 2$ or $\mathrm{char}\,\mathbb{K} = 2$. Then define $q_i^*(t)$ to be the square-free part of $\mathrm{Pert}_A(t, u_1, \ldots, u_{i-1}, u_i + \alpha, u_{i+1}, \ldots, u_n)$ for all $i \in [n]$.
Step 4 For all $i \in [n]$ and $j \in \{0,1\}$, let $r_{i,j}(\theta)$ be the reduction of $\mathcal{R}_j(q_i^-(t), q_i^*((\alpha+1)\theta - \alpha t))$ modulo $h(\theta)$.
Step 5 For all $i \in [n]$, define $h_i(\theta)$ to be the reduction of $-\theta - \frac{r_{i,1}(\theta)}{r_{i,0}(\theta)}$ modulo $h(\theta)$.

Note that assertion (0) thus follows immediately from Steps 1 and 4, thanks to the beginning of Main Theorem 2.4. Let us now verify the correctness of our algorithm, clarifying the genericity assumption of Step 0 along the way.

Using Main Theorem 4 once more, we know that the factors of Pert_A define for us a set of points $Z = \{\zeta^{(j)}\}_{j\in[N]}$, with $N \leq \mathcal{M}(E)$, such that Z intersects every irreducible component of $\varphi_A(\mathcal{Z})$. In particular, we see that the roots of h are exactly $\{\theta^{(j)}\}_{j\in N'}$, where $\theta^{(j)} := -\sum_{i=1}^n \zeta_i^{(j)} u_i$, $Z' := \{\zeta^{(j)}\}_{j\in N'} = Z \cap \mathbb{K}^n$, and $\zeta^{(j)} = (\zeta_1^{(j)}, \ldots, \zeta_n^{(j)})$ for all $j \in N'$. Furthermore, it is easy to check that for all but finitely many $[u_1 : \cdots : u_n]$, $j \neq j' \Longrightarrow \theta^{(j)} \neq \theta^{(j')}$. (In which case, via Remark 5.2 from Section 5, the multiplicity of any isolated root $\zeta^{(j)} \in (\mathbb{K}^*)^n$ of F is exactly the multiplicity of the root $\theta^{(j)}$ of h.) Similarly, for any $i \in [n]$, $j \neq j' \Longrightarrow \theta^{(j)} + \zeta_i^{(j)} \neq \theta^{(j')} + \zeta_i^{(j')}$ and $\theta^{(j)} - \alpha\zeta_i^{(j)} \neq \theta^{(j')} - \alpha\zeta_i^{(j')}$, for all but finitely many $[u_1 : \cdots : u_n]$. The avoidance of these $1 + 2n$ finite sets of $[u_1 : \cdots : u_n]$ is precisely our genericity condition for Step 0. Furthermore, by checking square-free parts, we can check our genericity condition with negligible overhead (via fact (c)).

Now note that if $\theta = \theta^{(j)}$ for some j, then for all $i \in [n]$, $q_i^-(t) = q_i^\star((\alpha+1)\theta - \alpha t) = 0 \Longleftrightarrow t = \theta^{(j)} + \zeta_i^{(j)}$. Furthermore, by construction, this common root has multiplicity 1 for both q_i^- and $q_i^\star$. It is then easily checked that $h_i(\theta^{(j)}) = \zeta_i^{(j)}$.

Recalling that the zero scheme of F in $(\mathbb{K}^*)^n$ is exactly $\mathcal{D}_A(F,\mathcal{P}) \cap (\mathbb{K}^*)^n$ (Rojas, 1999a, Section 5.1), we at last obtain assertions (1) and (2) of Main Theorem 2.1 by an application of Lemma 5.2. (In the case at hand, $\pi = \varphi_A$.) □

REMARK 5.3. The probability of failure in our Las Vegas algorithm above is 0, assuming any probability distribution on the coefficients of F yielding probability 1 avoidance of algebraic hypersurfaces in $\mathbb{K}^{\#\text{monomial terms}}$.

5.2. CONCLUDING THE PROOFS OF MAIN THEOREM 1 AND COROLLARY 2.1

We begin by checking the complexity of our Las Vegas algorithm from the preceding section. First note that by Main Theorem 2.4, each evaluation of Pert_A (for constant $u_0, u_1, \ldots, u_n$) takes $\mathcal{O}^*(nR(\bar{E})^2 S(\bar{E})^{2.376})$ arithmetic steps over $\mathbb{K}$. So by observation (a) above (and assertion (0)), we can find h via interpolation within time $\mathcal{O}^*(n\mathcal{M}(E)R(\bar{E})^2 S(\bar{E})^{2.376})$. Similarly, by (a), (b), and (c), we can find each q_i^- and $q_i^\star$ within the same time bound. So the construction of all these polynomials thus takes a total of $\mathcal{O}^*(n^2\mathcal{M}(E)R(\bar{E})^2 S(\bar{E})^{2.376})$ arithmetic steps over $\mathbb{K}$.

Finding the coefficients of $q_i^\star((\alpha+1)\theta - \alpha t)$ takes time $\mathcal{O}^*(\mathcal{M}(E)^2)$ via another simple interpolation step. So by (d), we can then find $h_1, \ldots, h_n$ still within the latter asymptotic time bound. As for space, we only need to keep track of $\mathcal{O}(n\mathcal{M}(E)^2)$ coefficient values, and this falls well within the $\mathcal{O}(nS(\bar{E})^2)$ space requirement of Main Theorem 2.4.

To conclude, we need only derandomize our algorithm. This can be done as follows: replace the generic selection of $u_1, \ldots, u_n$ above by $u_i = \varepsilon^i$ for $i \in [n]$. We then obtain that at our genericity condition is violated iff the point $(1, \varepsilon, \ldots, \varepsilon^n) \in \mathbb{K}^{n+1}$ lies in at least one of $(2n+1)\binom{\mathcal{M}(E)}{2}$ hyperplanes depending on the input F. From the box principle, and the well-known properties of the van der Monde matrix (Bini and Pan, 1994), this can happen to at most $n(2n+1)\binom{\mathcal{M}}{2}(E)$ 2 distinct values of ε. So we can derandomize

by repeatedly running steps (1)–(3) with new ε at most $n(2n+1)\binom{\mathcal{M}(E)}{2}$ times, thus finally accounting for our aforementioned deterministic time bound.

Moving on, we must now further refine our algorithm so that our arithmetic is over K (or a small algebraic extension thereof) instead of $\mathbb{K}$. This can be done as follows: If $\operatorname{char}\mathbb{K} = 0$, then there are enough choices for ε in K to derandomize our algorithm (since K will be infinite). Otherwise, we simply choose ε in an algebraic extension of K of degree $\lceil \log_p((n+1)^2\mathcal{M}(E)^2)\rceil$, so that we have more than enough ε to choose from. Assertion (3) is now proved.

To conclude, Remark 5.1 tells us that the zero scheme of F in $\mathbb{K}^n$ embeds naturally in $\mathcal{D}_P(F,\mathcal{P})$ (and $\mathcal{D}_{\bar{P}}(F,\mathcal{P})$) if we replace E by $\mathbf{O}\cup E$. So this introduction of extra points into our supports indeed guarantees that Z' includes all the affine roots of F. □

REMARK 5.4. We have thus improved the complexity of finding all the affine roots (roughly) from polynomial in $\prod d_i$ to polynomial in $\mathcal{M}(\mathbf{O}\cup E)$. However, one can improve this even further to polynomial in $\mathcal{SM}(E)$ — the stable mixed volume (Huber and Sturmfels, 1997; Rojas, 1998a, b, c) of E. (In particular, $\mathcal{SM}(E) \leq \mathcal{M}(\mathbf{O}\cup E) \leq \prod d_i$ and the gaps between can be quite large (cf. Example 6.1).) To make this final improvement, it is necessary to use a more refined resultant operator — the affine toric resultant, denoted $\operatorname{AffRes}_{\bar{E}}(\bar{F})$ (Rojas, 1997a, b). This is covered at greater length in (Rojas 1997a, b, 1998a, b, c), and this new operator also allows us to extend Corollaries 2.2 and 2.3 to $\mathbb{K}^n$ minus an arbitrary union of coordinate hyperplanes.

PROOF OF COROLLARY 2.1. It follows immediately from our proof of Main Theorem 1 that the fields $K[\zeta_i \mid (\zeta_1,\ldots,\zeta_n)\in(\mathbb{K}^*)^n$ is a root of $F]$ and $K[\theta \mid h(\theta)=0, \prod h_i \neq 0]$ are identical when $u_1,\ldots,u_n$ are chosen from K. (So the assumption that $\operatorname{char}\mathbb{K}=0$ is actually stronger than necessary.) As the latter field is exactly the splitting field of g, we are done. □

5.3. THE PROOF OF MAIN THEOREM 2

We first note that the well-known results on the degree of $\operatorname{Res}_{\bar{E}}(f_1,\ldots,f_{n+1})$ with respect to the coefficients of various f_i (Sturmfels, 1994) remain true over any algebraically closed field. This follows easily from the formulation of the resultant for a collection of invertible sheafs on a projective variety (Gel'fand *et al.*, 1994). In particular, Chow_A should indeed be either be identically zero or a homogeneous polynomial (in the u_a) of degree $\mathcal{M}(E)$.

To prove assertions (1)–(3), we can then simply invoke the Vanishing Theorem for Resultants and Lemma 5.2 (since $\bar{P}$ is compatible with $\operatorname{Conv}(A)$). For instance, we obtain that $\varphi_A(\mathcal{Z})$ is positive-dimensional iff Chow_A has infinitely many distinct divisors of the form $\sum_{a\in A}\gamma_a u_a$. So assertion (1) follows immediately. Assertions (2) and (3) follow similarly. □

5.4. THE PROOF OF COROLLARY 2.2

Let ω (<2.376) denote the famous matrix multiplication complexity exponent (Coppersmith and Winograd, 1990) and set $E_{n+1} = A$. It then follows immediately from Coppersmith and Winograd (1990) and Emiris and Canny (1995, The Division

Method) that for any choice of constant coefficients in $\mathbb{K}$, $\mathrm{Res}_{\bar{E}}(\bar{F})$ can be evaluated within $\mathcal{O}^*(nR(\bar{E})S(\bar{E})^\omega)$ arithmetic operations over $\mathbb{K}$, using $\mathcal{O}(nS(\bar{E})^2)$ space.

The first part of Corollary 2.2 then follows immediately from a van der Monde type argument, as in the proof of Main Theorem 2.1. In particular, via interpolation, it suffices to evaluate $\mathrm{Chow}_A(u)$ at exactly $1 + n\mathcal{M}(E)$ distinct points of the form $(1, \varepsilon, \ldots, \varepsilon^n)$ to see if Chow_A is identically zero.

To then count the roots of F in $(\mathbb{K}^*)^n$ when Chow_A is not identically zero, we can begin with a variant of the algorithm from Main Theorem 2.1 where we evaluate $\mathrm{Chow}_{\Delta\cap\mathbb{Z}^n}$ instead of $\mathrm{Pert}_{\Delta\cap\mathbb{Z}^n}$. From our previous observations, we can thus construct h and $h_1, \ldots, h_n$ within time $\mathcal{O}^*(n^4\mathcal{M}(E)^3R(\bar{E})S(\bar{E})^\omega)$ and space $\mathcal{O}(nS(\bar{E})^2)$.

We then use the following trick: compute the gcd, g, of h and $\prod_{i=1}^n h_i$. By Remark 5.2 of Section 5, we immediately obtain that $\deg h - \deg g$ is exactly the number of roots of F in $(\mathbb{K}^*)^n$ counting multiplicities. (In fact, the roots of g tell us precisely which $\zeta^{(j)}$ lie **out** of $(\mathbb{K}^*)^n$.) By the same argument, we can also count the number of distinct roots simply by replacing h with its square-free part. By facts (b) and (c) of Section 5.1, and because the degree of $\prod_{i=1}^n h_i$ is at most $n\mathcal{M}(E)$, these computations cause a negligible growth in our asymptotic complexity bounds. So we are done. □

5.5. FACET SEARCHES AND THE PROOF OF MAIN THEOREM 3

The first portion of this result follows immediately from Lemma 4.2 and Rojas (1994, Corollary 3). The second portion is a consequence of the following algorithm:

Step 1 Compute the facet normals of P and the vertices of all the $\mathrm{Conv}(E_i)$.
Step 2 Find a vertex v of some E_i such that for any facet normal w of P, $v \in E_i^w \Longrightarrow [\#E_i^w \geq 2$ or $\mathcal{M}(E_1, \ldots, E_{i-1}, E_{i+1}, \ldots, E_n) = 0]$. If no such v exists, stop. Otherwise, delete v from E_i^w and go back to step 1.

By Lemma 4.2, the above algorithm will eventually stop with an irreducible fill of E. As for its complexity, note that the number of facets of P is $\mathcal{O}(m^{2n})$, and we can find the normals to these facets within that many arithmetic steps over $\mathbb{Q}$ (Gritzmann and Sturmfels, 1993), given the convex hulls of the E_i. Furthermore, this asymptotic bound dominates the complexity of finding the convex hulls of all the E_i (Preparata and Shamos, 1985; Chan, 1996). So the complexity of Step (1) is $\mathcal{O}(m^{2n})$. Step (2) thus amounts to $n\mathcal{O}(m^{2n})$ checks for zero mixed volume per vertex. So by Lemma 4.1, this takes $\mathcal{O}(n^{2.616}m^{2n+1})$ arithmetic steps over $\mathbb{Q}$. These steps will be executed at most m times, so we are done.

5.6. ALGEBRAIC HOMOTOPIES AND THE PROOF OF MAIN THEOREM 4

Main Theorem 2.4 is the cornerstone of our approach to solving degenerate systems of equations, so we will precede its proof by illustrating one of its underlying constructions: explicit algebraic deformation of degenerate zero sets.

More precisely, following the notation of Main Theorem 2.4, we will construct a family of curves C, fibered over the projective line, whose fiber over a particular point is a zero-dimensional variety $Z \subseteq \mathcal{Z}$ encoding the multiplicities of all the irreducible components of $\mathcal{Z}$. To do this, we begin with the following lemma, which follows easily from the development of Rojas (1999a, Section 5.1) and Fulton (1984, Section 11.3).

LEMMA 5.4. *Following the notation of Definition 2.2 and Main Theorem 2.4, let $\mathcal{Z}_0$ be the zero-dimensional part of $\mathcal{Z}$. Also let $\mathcal{Z}^\times$ be the zero scheme of $F - sF^*$ in $\mathcal{T}(\bar{P}) \times \mathbb{P}^1_{\mathbb{K}}$. Then $\mathcal{Z} = \mathcal{Z}^\times \cap (\mathcal{T}(\bar{P}) \times \{0\})$. Finally, let C be the algebraic curve (possibly reducible) defined by the union of all one-dimensional components of $\mathcal{Z}^\times$ with surjective projection onto the second factor of $\mathcal{T}(\bar{P}) \times \mathbb{P}^1_{\mathbb{K}}$. Then C has the following properties:*

(1) $\mathcal{Z}^\times \cap (\mathcal{T}(\bar{P}) \times \{s_0\}) = C \cap (\mathcal{T}(\bar{P}) \times \{s_0\})$ *for almost all* $s_0 \in \mathbb{P}^1_{\mathbb{K}}$.
(2) $Z := C \cap (\mathcal{T}(\bar{P}) \times \{0\})$ *is a subscheme of $\mathcal{Z}$ consisting of exactly $\mathcal{M}(E)$ points (counting multiplicities). Furthermore, $\mathcal{Z}_0$ is a subscheme of Z.*
(3) *Let W be any irreducible component of $\mathcal{Z}$. Then Z has at least one point in W and, for a generic choice of F^*, the number of points of Z in W (counting multiplicities) is exactly the cycle class degree of W.*

We can now begin our most important proof.

PROOF OF MAIN THEOREM 2.4. Similar to the beginning of the proof of Main Theorem 2.2, the results of Sturmfels (1994) (generalized to arbitrary algebraically closed $\mathbb{K}$) immediately imply that the degree of $\mathcal{H}$ as a polynomial in s should be $\sum_{i=1}^n \mathcal{M}(E_1, \ldots, E_{i-1}, E_{i+1}, \ldots, E_n, A) \leq R(\bar{E})$. Also each coefficient of $\mathcal{H}(s)$ should be a homogeneous polynomial (in the u_a) of degree $\mathcal{M}(E)$. These two assertions of course include the opening statement of Main Theorem 2.4 (on the degree and homogeneity of Pert_A), but they will follow only upon showing that $\mathcal{H}$ is not identically zero.

To see this, note that Lemma 5.1 and the Vanishing Theorem for Resultants readily imply that the coefficient of the highest power of s in $\mathcal{H}$ is precisely $\mathrm{Res}_{(E,A)}(F^*, f_{n+1})$. (Simply check the zero set of $F - sF^*$ in $\mathcal{T}(\bar{P})$ at $s = \infty$, via the homogenization $s'F - sF^*$.) By Definition 2.2, and the Vanishing Theorem once more, we see that this polynomial in the u_a is not identically zero. So $\mathcal{H} \not\equiv 0$ and we have finished the simplest part of our proof.

Part (1) of Main Theorem 2.4 follows similarly: one need only consider the unspecialized resultant polynomial $\mathrm{Res}_{(E,A)}(F, f_{n+1})$ and observe the terms of degree 0 in s as we specialize coefficients to obtain $F - sF^*$. In particular, $\mathrm{Chow}_A(u)$ is precisely $\mathcal{H}(u; 0)$. Note then that (2) and (3) also follow almost immediately, provided Chow_A is not identically zero.

To handle the cases of (2) and (3) properly, where we are actually working with a non-trivial toric perturbation, we now invoke Lemmata 5.2 and 5.4 to establish a precise correspondence between the factors of Pert_A and the points of Z.

Letting $\mathcal{Z}^\times_{A,+}$ be the zero set of $\mathcal{H}(u; s)$ in $\mathbb{P}^{\#A-1}_{\mathbb{K}} \times \mathbb{P}^1_{\mathbb{K}}$, note that if j is the least exponent of s in $\mathcal{H}$, then $\mathcal{Z}^\times_{A,+}$ and the zero set of $\frac{\mathcal{H}}{s^j}$ in $\mathbb{P}^{\#A-1}_{\mathbb{K}} \times \mathbb{P}^1_{\mathbb{K}}$ differ only by the presence of the hyperplane $\mathbb{P}^{\#A-1}_{\mathbb{K}} \times \{0\}$. The second zero set does not contain this hyperplane, so let us call the second zero set $\mathcal{Z}^\times_A$. By Lemmata 5.1 and 5.2, and the Vanishing Theorem for Resultants, we then derive that $\dim[\mathcal{Z}^\times \cap (\mathcal{T}(\bar{P}) \times \{s_0\})] = 0$ implies the following equivalence: $\mathcal{H}(H_{\varphi_A(\zeta)}; s_0) = 0 \iff \zeta \in \mathcal{Z}^\times \cap (\mathcal{T}(\bar{P}) \times \{s_0\})$, where H_p is the hyperplane dual to the point p.† By assertion (1) of Lemma 5.4, $\dim[\mathcal{Z}^\times \cap (\mathcal{T}(\bar{P}) \times \{s_0\})] = 0$ for almost all $s_0 \in \mathbb{P}^1_{\mathbb{K}}$. So $C^\vee$ is an open subset of $\mathcal{Z}^\times_A$, where we define $C^\vee := \{(y, s_0) | y \in H_{\varphi_A(\zeta)}\ ;\ \zeta \in C \cap (\mathcal{T}(\bar{P}) \times \{s_0\})\ ;\ s_0 \in \mathbb{P}^1_{\mathbb{K}}\}$. Therefore, as φ_A is a proper map, $\frac{\mathcal{H}}{s^j}$ must

†So if $p := [p_a | a \in A] \in \mathbb{P}^{\#A-1}_{\mathbb{K}}$ then $H_p := \{[y_a | a \in A] \in \mathbb{P}^{\#A-1}_{\mathbb{K}} | \sum_{a \in A} p_a y_a = 0\}$.

vanish on all of $C^\vee$. In particular, via Remark 5.2 of Section 5,

$$\mathrm{Pert}_A(u) = \alpha \cdot \prod_{\zeta \in C \cap (\mathcal{T}(\bar{P}) \times \{0\})} \left(\sum_{a \in A} \gamma_{\zeta,a} u_a \right),$$

where $\alpha \in \mathbb{K}^*$, $[\gamma_{\zeta,a} | a \in A] := \varphi_A(\zeta)$, and the product counts intersection multiplicities.

Continuing our main proof, assertions (2) and (3) follow immediately from our last formula and our preceding observations. As for the complexity bounds, these follow immediately from our earlier fact (a) and the Division Method (Emiris and Canny, 1995) to compute $\mathrm{Res}_\star(\cdot)$: to evaluate $\mathrm{Pert}_A(u)$, we simply find the coefficients of $\mathcal{H}(u;s)$ by evaluating $\mathcal{H}(u;s)$ at $R(\bar{E})+1$ distinct values of s and then interpolating. □

Note that our algebraic proof avoids the use of limiting arguments that were present in Canny (1990). Thus our result holds for any algebraically closed $\mathbb{K}$, instead of just $\mathbb{C}$.

5.7. DOUBLE PERTURBATIONS AND THE PROOF OF COROLLARY 2.3

The first portion of our final corollary follows immediately (thanks to Main Theorem 2.4) by simply replacing Chow_A with Pert_A in the algorithm from the proof of Corollary 2.2. In particular, we obtain that the exact number of roots of F in $(\mathbb{K}^*)^n$ (counting multiplicities) is exactly $\deg h^* - \deg g^*$, where h^* (resp. g^*) is the corresponding variant of h (resp. g), using the notation of the proof of Corollary 2.2. The number of distinct roots can of course be recovered by using square-free parts (as before), thanks to Remark 5.2 of Section 5. Also, by Main Theorem 2.4, the complexity of this algorithm is just the complexity estimate from Corollary 2.2 multiplied by $R(\bar{E})$.

As for the second portion of our corollary, we make a slightly more sophisticated variant of the preceding replacement of Chow_A.

DEFINITION 5.2. Let F^* and F^{**} be $n \times n$ polynomial systems with support contained in E such that (1) F^* and F^{**} each have only finitely many roots in $\mathcal{T}(P)$, and (2) F^* and F^{**} share no common roots. Following the notation of Main Theorem 2.4, define a double toric perturbation of F, Pert_A^{**}, to be the greatest common divisor of Pert_{A,F^*} and $\mathrm{Pert}_{A,F^{**}}$.

It is then clear (via Main Theorem 2.4 once again) that using Pert_A^{**} in place of Pert_A in our preceding algorithm will lead to a new estimate, $\deg h^{**} - \deg g^{**}$, for $\deg(\mathcal{Z}_0 \cap (\mathbb{K}^*)^n)$. Furthermore, by the above definition, it is clear that $\deg h^{**} - \deg g^{**} \leq \deg h^* - \deg g^*$.

As for estimating $\deg \mathcal{Z}_0$ and $\deg \mathcal{Z}_\infty$, our preceding theory tells us that we can simply respectively use $\deg h^{**}$ and $\mathcal{M}(E) - \deg h^{**}$.

REMARK 5.5. Our algorithm thus requires a generic choice of F^* and F^{**}. Just as in the construction of Pert_A, we can derandomize via combinatorial means: We simply use an irreducible fill (as in Main Theorem 2.3) to construct F^*, and then simply perturb a single coefficient of F^* to construct F^{**}. This is the trick used in our earlier example in Section 3.2.

REMARK 5.6. The basic idea behind the double perturbation is that the points in $Z := \{\gamma(\theta)\}_{h(\theta)=0}$ lying in positive-dimensional components of $\varphi_A(\mathcal{Z})$ will move as we vary F^*.

Thus, assuming that F^{**} is such that the new Z overlaps the old Z only on the isolated roots of F, we should be able to pick out these isolated roots simply by computing the gcd of Pert_{A,F^*} and $\mathrm{Pert}_{A,F^{**}}$. We hope to address this "motion of points within a deformation" in future work.

6. Computing Toric Resultants and the Complexity of the Sparse Encoding

Let us first recall some important facts on the computation of toric resultants.

As of 1998, the main method for computing $\mathrm{Res}_{\bar{E}}(\bar{F})$ is to first construct an $S(\bar{E}) \times S(\bar{E})$ **toric resultant matrix**, $M_{\bar{E}}$, whose nonzero entries are certain coefficients of $\bar{F}$. This matrix is specifically built so that $\det(M_{\bar{E}})$ is, for generic choices of the coefficients $c_{i,a}$, a nonzero multiple of $\mathrm{Res}_{\bar{E}}(\bar{F})$.

REMARK 6.1. So $S(\bar{E})$ is actually a parameter depending on which algorithm we use for constructing $M_{\bar{E}}$ — hence our earlier use of an asymptotic bound, instead of an explicit formula, for $S(\bar{E})$. The aforementioned bound is actually a simple estimate on the number of lattice points in the interior of the shifted Minkowski sum $\delta + \sum_{i=1}^{n+1} \mathrm{Conv}(E_i)$, where $\delta \in \mathbb{Q}^n$ is chosen generically. The derivation follows easily from Stirling's estimate for the Γ-function, the n-dimensional identity $\mathcal{M}(P,\ldots,P) = n!\,\mathrm{Vol}(P)$, and the multilinearity of the mixed volume.

Via some clever interpolation tricks (Canny and Emiris, 1995; Emiris and Canny, 1995; Emiris and Pan, 1997), one can recover the exact value of $\mathrm{Res}_{\bar{E}}(\bar{F})$ after interpolating $\det(M_{\bar{E}})$ through several-many specializations of the coefficients of $\bar{F}$. One such fundamental technique, which uses $n+1$ versions of $M_{\bar{E}}$, is known as the Division Method (Canny, 1987; Emiris and Canny, 1995). In general, the matrix $M_{\bar{E}}$ is highly structured (it is quasi-Toeplitz (Emiris and Pan, 1997)) and, when $\mathrm{char}\,\mathbb{K} = 0$, this permits $\mathrm{Res}_{\bar{E}}(\bar{F})$ to be computed much faster than would be expected.

In practice, the cost of building $M_{\bar{E}}$ (or several versions thereof) can be amortized when one works with many $\bar{F}$ with support contained in the same $\bar{E}$. (In fact, via the Cayley trick (Gel'fand *et al.*, 1994), it reasonably follows from standard results on triangulations (Preparata and Shamos, 1985; Chan, 1996) and lattice point enumeration (Barvinok, 1994) that the complexity of constructing $h, h_1, \ldots, h_n$ dominates the preprocessing complexity.) Furthermore, when randomization is allowed, the results of Canny and Emiris (1995), Emiris and Canny (1995) and Emiris and Pan (1997) tell us that this preprocessing is actually negligible.

As for the complexity of computing $\mathrm{Res}_{\bar{E}}(\bar{F})$ itself, we state the following additional facts:

I (Emiris and Canny, 1995, The GCD Method) When $\mathrm{char}\,\mathbb{K} = 0$, we can compute $\mathrm{Res}_{\bar{E}}(\bar{F})$ (for any choice of constant coefficients in $\mathbb{K}$ for $\bar{F}$) within $\mathcal{O}^*(S(\bar{E})^{1+\omega})$ arithmetic steps and $\mathcal{O}(S(\bar{E})^2)$ space.† However, we have the added benefit that we can also compute $\mathcal{H}(u;s)$ (for any constant $u \in \mathbb{K}^{\#A}$) within the same complexity bound.

†The restriction on $\mathrm{char}\,\mathbb{K}$ is due to a use of effective Hilbert irreducibility, which actually fails in positive characteristic (Lang, 1983).

II (Emiris and Pan, 1997) If we assume $\operatorname{char}\mathbb{K} = 0$ and allow randomization, then we can accelerate the Division Method (resp. GCD Method) to obtain a Las Vegas time bound of $\mathcal{O}^*(n^2R(\bar{E})S(\bar{E})^2)$ (resp. $\mathcal{O}^*(nS(\bar{E})^3)$). Furthermore, either of these improvements requires only $\mathcal{O}^*(nS(\bar{E}))$ space.

III If $\operatorname{Res}_{\bar{E}}(\bar{F}) = \det(M_{\bar{E}})$, then $S(\bar{E}) \leq R(\bar{E})$ and we can reduce the deterministic time bounds of the Division and GCD methods to $\mathcal{O}(R(\bar{E})^\omega)$, regardless of $\operatorname{char}\mathbb{K}$. Furthermore, if we also allow randomization and assume $\operatorname{char}\mathbb{K} = 0$, then we can further improve the time bounds of (I) and (II) to $\mathcal{O}^*(R(\bar{E})^2)$. However, characterizing when $\operatorname{Res}_{\bar{E}}(\bar{F})$ can be expressed as a "small" determinant is an open problem. (See Weyman and Zelevinsky (1994) for some interesting partial results, including some cases where the Newton polytopes are products of scaled standard simplices.)

The last fact is actually a simple corollary of the development of Emiris and Pan (1997). In particular, in the situation of (III), we can skip an interpolation procedure that would have multiplied our time bound by $\mathcal{O}^*(R(\bar{E}))$.

Let us now state and prove the best current speed-ups for all our preceding algorithmic results.

COROLLARY 6.1. *Suppose* $\operatorname{char}\mathbb{K} = 0$ *and we allow randomization in our algorithms. Then our main algorithmic results can be sped up as follows:*

	Sequential (Las Vegas) Time Bound $= \mathcal{O}^*(\cdots)$
Main Theorem 1	$n^3\mathcal{M}(E)R(\bar{E})^2S(\bar{E})^2$ *or* $n^2\mathcal{M}(E)S(\bar{E})^3$
Corollary 2.2 (First Bound)	$n^2\mathcal{M}(E)R(\bar{E})S(\bar{E})^2$ *or* $n\mathcal{M}(E)S(\bar{E})^3$
Corollary 2.2 (Second Bound)	$n^3\mathcal{M}(E)R(\bar{E})S(\bar{E})^2$ *or* $n^2\mathcal{M}(E)S(\bar{E})^3$
Main Theorem 4	$n^2R(\bar{E})^2S(\bar{E})^2$ *or* $nS(\bar{E})^3$

Furthermore, the space bound for each of the above algorithms is $\mathcal{O}^*(nS(\bar{E}))$. *Finally, if we also have that* $\operatorname{Res}_{\bar{E}}(\bar{F}) = \det(M_{\bar{E}})$, *then the four pairs of entries in the right-hand column (from top to bottom) can be replaced by the following sequence:* $n\mathcal{M}(E)R(\bar{E})^3$, $\mathcal{M}(E)R(\bar{E})^2$, $n\mathcal{M}(E)R(\bar{E})^2$, $\mathcal{M}(E)R(\bar{E})^3$.

REMARK 6.2. As before, the probability of failure in all our Las Vegas algorithms above is 0, assuming any probability distribution on the coefficients of F yielding probability 1 avoidance of algebraic hypersurfaces in $\mathbb{K}^{\#\text{monomial terms}}$. The total number of random choices of elements in K (or a small algebraic extension thereof) needed is $n+R(\bar{E})$. (This is just the number choices needed to construct h and a variant (Emiris and Pan, 1997) of $M_{\bar{E}}$.)

REMARK 6.3. Our algorithms are also well-parallelizable. In particular, all of the problems considered in Corollary 6.1 can be shown to lie in the complexity class **PSPACE** (and in **NC** for fixed n). While this was known for the first problem (e.g. Ben-Or *et al.* (1986), Canny (1988) and Fitchas *et al.* (1990)), our techniques enable us to derive a much sharper deterministic parallel time bound of $\mathcal{O}^*(\log S(\bar{E}))$ (using a number of processors polynomial in $S(\bar{E})$) for all of these problems (Rojas, 1999b).

So in summary, we can solve any $n \times n$ system, over an algebraically closed field of characteristic zero, in Las Vegas time near-quartic in the number of roots of a closely

related system.† Furthermore, our algorithms are well-parallelizable, and we can go even faster when we have a sufficiently compact toric resultant matrix. Before proving the above corollary, we will briefly explain what we mean by a "closely related system".

First recall that $\mathcal{M}(E)$ is precisely the cycle class degree of the toric divisor $\mathcal{D}_P(F,\mathcal{P})$ (Fulton, 1993; Rojas, 1999a). Put more simply, if we simply perturb the coefficients of F, we can expect F to have exactly $\mathcal{M}(E)$ roots in $\mathcal{T}(P)$. Thus, the quantity $\mathcal{M}_{\bar{E}}^{\text{ave}}$ defined earlier can be reinterpreted as follows: it is the average number of roots of an $n\times n$ system of equations with support contained in $(\mathcal{E}_1,\ldots,\mathcal{E}_n)$, as we let the $\mathcal{E}_i$ independently range over $\{E_1,\ldots,E_{n+1}\}$, and we assume generically chosen coefficients. So the quantity $S(\bar{E})$ can also be interpreted as a weighted average of a set of cycle class degrees. Similarly, note that the generic number of roots of the $(n+1)\times(n+1)$ system $(F-sF^*, s-s_0)$ is exactly $\mathcal{M}(E_1\times\{0,1\},\ldots,E_n\times\{0,1\},\{0,\hat{e}_{n+1}\})$. So by the multilinearity of the mixed volume, the last mixed volume is exactly $R(\bar{E})$.

Let us now prove our above corollary.

PROOF OF COROLLARY 6.1. The key bounds to begin with are those of Corollary 2.2. In particular, the first bound of Corollary 2.2 is the complexity of determining whether $\text{Chow}_A(u)$ vanishes identically. As this can be accomplished by evaluating $\text{Res}_{\bar{E}}(\bar{F})$ at $\mathcal{M}(E)+1$ random points, facts (I)–(III) above immediately imply our asserted bounds.

As for the second bound of Corollary 2.2, this is the complexity of running a variant of the algorithm of Main Theorem 1, where Pert_A is replaced by Chow_A. As Chow_A is just a specialized resulant, and because this algorithm boils down to evaluating Chow_A at $\mathcal{O}(n\mathcal{M}(E))$ distinct points, facts (I)–(III) immediately imply these bounds as well.

From the proof of Main Theorem 4, we know that the bound from Main Theorem 4 is simply the complexity of evaluating $\text{Res}_{\bar{E}}(\bar{F})$ at $R(\bar{E})+1$ different specializations of s. (Remember that s occurs only in the coefficients of $f_1,\ldots,f_n$, and all other parameters are assumed to be constants.) So this bound follows easily from facts (I)–(III) as well.

To conclude, the bound from Main Theorem 1 is simply the complexity of evaluating Pert_A at $\mathcal{O}(n\mathcal{M}(E))$ distinct points. From the bound of Main Theorem 4, we are done. □

Are the above sequential complexity bounds the best one can expect for solving polynomial systems specified in the sparse encoding?‡ Neglecting the precise values of the exponents (which we have seen range somewhere between 4 and 7.376, if not better), the answer is "yes". This is due to the fact that a generic F will have exactly $\mathcal{M}(E)$ distinct roots in $\mathbb{K}$, regardless of the number of terms present. Thus, it is really $\mathcal{M}(E)$, not the number of terms, which governs the complexity of global polynomial system solving over an algebraically closed field. So the quantities in the "base" of our bounds can not be any smaller (asymptotically) than $\mathcal{M}(E)$. As for the exponent, we so far only have the obvious worst case lower bound of 1.

However, the question of whether the number of terms more strongly governs the complexity of solving over a non-algebraically closed field, or solving for a single root, is quite open. For example, while Khovanskii has shown that the number of real roots of a sparse system of equations is singly exponential in the number of terms (Khovanskii, 1991), the complexity of real solving is not yet known to fall within such a bound, even when $n=1$. Similarly, while a recent algorithm of Ye (1994) for ε-approximating a

†We conjecture that this can be done in positive characteristic as well. The main current obstruction is the use (in current fast algorithms) of algebraic identities for recovering elementary symmetric functions from power sums, which fail for small positive characteristic.

‡That is, when we specify polynomial systems as a list of exponents and coefficients.

single dth root of $\alpha \in \mathbb{R}$ has sequential arithmetic complexity $\mathcal{O}((\log d) \log\log \frac{|\alpha|}{\epsilon})$, the complexity of finding a single root of F in $\mathbb{K}^n$ is quite open. It is also interesting to note that in spite of recent successes for fewnomials over a number field (Lenstra, 1997a, 19XX), it is still unknown whether a single real root of a degree d univariate trinomial can be ϵ-approximated within time polynomial in $\log(d)$ and $\log(\frac{1}{\epsilon})$. We hope to address these finer points of sparse algebraic complexity in future work.

We now close with a brief example of how $\mathcal{M}(E)$ can be smaller than D_Π (the product of the total degrees of $f_1, \ldots, f_n$) by an exponential factor.

EXAMPLE 6.1. (WELL DIRECTED SPIKES) Consider the system of equations F defined by

$$\begin{aligned} a_{1,1} + a_{1,2}x_1 + \cdots + a_{1,n}x_{n-1} + c_{1,1}(x_1 \cdots x_n) + \cdots + c_{1,d}(x_1 \cdots x_n)^d &= 0 \\ &\vdots \\ a_{n,1} + a_{n,2}x_1 + \cdots + a_{n,n}x_{n-1} + c_{n,1}(x_1 \cdots x_n) + \cdots + c_{n,d}(x_1 \cdots x_n)^d &= 0. \end{aligned}$$

In this case, the Newton polytopes are all equal to a single "spike", and this spike is equivalent (via an integer linear map with determinant 1) to a standard n-simplex scaled by d in the x_1-direction. So it is easy to check that $\mathcal{M}(E) = d$. However, the product of the total degrees of F is clearly $n^n d^n$. (It is also not hard to see that the best multigraded Bézout bound (Wampler, 1992) is $n!d^n$.) Generating infinite families of such examples is easy, simply by picking Newton polytopes which are n-dimensional, but "long" in a suitable fixed direction.

REMARK 6.4. The construction of toric resultant matrices is an area of active research and it can be reasonably expected that our earlier asymptotic estimate on $S(\bar{E})$ will be significantly improved in the near future. In particular, a significant first step would be to find an algorithm which always constructs a toric resultant matrix of size $\mathcal{O}(R(\bar{E}))$. Looking even further ahead, there is also hope for general algorithms which construct even smaller matrices, via the use of entries which are nonlinear polynomials in the coefficients of $\bar{F}$.

Acknowledgements

This research was partially funded by a Hong Kong CERG grant and a US National Science Foundation Mathematical Sciences Postdoctoral Fellowship.

The author would like to thank an anonymous referee for extensive comments on clarifying the exposition and development of this paper. The author also expresses his deep gratitude to Gregorio Malajovich-Muñoz for his assistance in computing the matrix from Section 3.2. Special thanks also go to Pino Italiano and Nini Wong for their help in obtaining a copy of Canny (1988).

References

Barvinok, A. I. (1994). Computing the Erhart polynomial of a convex lattice polytope. *Discrete Comput. Geom.*, **12**, 35–48.

Ben-Or, M., Kozen, D., Reif, J. (1986). The complexity of elementary algebra and geometry. *J. Comput. Syst. Sci.*, **32**, 251–264 (1986).

Bernshtein, D. N. (1975). The number of roots of a system of equations. *Functional Analysis and its Applications* (translated from Russian), **9**, 183–185.

Bini, D., Pan, V. Y. (1994). *Polynomial and Matrix Computations, Volume 1: Fundamental Algorithms,* Progress in Theoretical Computer Science. Basel, Birkhäuser.

Burago, Yu. D., Zalgaller, V. A. (1988). *Geometric Inequalities,* Grundlehren der mathematischen Wissenschaften 285. Berlin, Springer-Verlag.

Canny, J. F. (1987). The complexity of robot motion planning problems. *ACM Doctoral Dissertation Award Series.* New York, ACM Press.

Canny, J. F. (1988). Some Algebraic and Geometric Computations in PSPACE. In *Proc. 20th ACM Symp. on the Theory of Computing, Chicago.* ACM Press, U.S.A.

Canny, J. F. (1990). Generalised characteristic polynomials. *J. Symb. Comput.*, **9**, 241–250.

Canny, J. F., Emiris, I. Z. (1995). A subdivision-based algorithm for the sparse mixed resultant, preprint, INRIA .

Canny, J. F., Kaltofen, E., Lakshman, Y. (1989). Solving systems of non-linear polynomial equations faster. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pp. 121–128. ACM Press, U.S.A.

Chan, T. M. (1996). Output-sensitive results on convex hulls, extreme points, and related problems. *Discrete Comput. Geom.*, **16**, 369–387.

Chistov, A. L., Grigoriev, D. Y. (1984). In *Complexity of Quantifier Elimination in the Theory of Algebraically Closed Fields*, LNCS **176**. Berlin, Springer-Verlag.

Coppersmith, D., Winograd, S. (1990). Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, **9**, 251–280.

Dalbec, J., Sturmfels, B. (1995). *Introduction to Chow Forms, Invariant Methods in Discrete and Computational Geometry (Curaçao, 1994)*, pp. 37–58. Dordrecht, Kluwer.

Danilov, V. I. (1978). The geometry of toric varieties. *Russ. Math. Sur.*, **33**, 97–154.

Díaz, A., Kaltofen, E. (1995). On computing greatest common divisors with polynomials given by black boxes for their evaulations. In *Proceedings of ISSAC '95, Montreal Canada*, pp. 232–239. New York, ACM Press.

Dyer, M., Gritzmann, P., Hufnagel, A. (1998). On the complexity of computing mixed volumes. *SIAM J. Comput.*, **27**, 356–400.

Emiris, I. Z., Canny, J. F. (1995). Efficient incremental algorithms for the sparse resultant and the mixed volume. *J. Symb. Comput.*, **20**, 117–149.

Emiris, I. Z., Pan, V. Y. (1997). The structure of sparse resultant matrices. In *Proc. Int. Symp. on Symbolic and Algebraic Computation (ISSAC).* New York, ACM Press.

Ewald, G. (1996). *Combinatorial convexity and algebraic geometry*, Graduate Texts in Mathematics **168**. New York, Springer-Verlag.

Fitchas, N., Galligo, A., Morgenstern, J. (1990). Precise sequential and parallel complexity bounds for quantifier elimination over algebraically closed fields. *J. Pure Appl. Algebra*, **67**, 1–14.

Fulton, W. (1984). *Intersection Theory.* Berlin, Springer-Verlag.

Fulton, W. (1993). *Introduction to Toric Varieties*, Annals of Mathematics Studies **131**. Princeton, NJ, Princeton University Press.

Gabow, H. N., Ying, X. (1996). Efficient theoretic and practical algorithms for linear matroid intersection problems. *J. Comput. Syst. Sci.*, **53**, 129–147.

Gel'fand, I. M., Kapranov, M. M., Zelevinsky, A. V. (1994). *Discriminants, Resultants and Multidimensional Determinants.* Boston, MA, Birkhauser.

Giusti, M., Heintz, J., Morais, J. E., Pardo, L. M. (1995). When polynomial equation systems can be 'solved' fast? In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes* (Paris, 1995), LNCS **948**, pp. 205–231. Berlin, Springer.

González-Vega, Laureano, A subresultant theory for multivariate polynomials. In Watt, S. M., ed., *Proc. 1991 Int. Symp. on Symbolic and Algebraic Computation*, pp. 79–85. New York, ACM Press.

Gritzmann, P., Klee, V. (1994). On the complexity of some basic problems in computational convexity II: volume and mixed volumes, polytopes: abstract, convex, and computational (Scarborough, ON, 1993), pp. 373–466, NATO. *Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, **440**. Dordrecht, Kluwer.

Gritzmann, P., Sturmfels, B. (1993). Minkowski addition of polytopes: computational complexity and applications to Gröbner bases. *SIAM J. Discrete Math.*, **6**, 246–269.

Grötschel, M., Lovász, L., Schrijver, A. (1993). Geometric algorithms and combinatorial optimization. In *Algorithms and Combinatorics*, 2, 2nd corrected edition. Berlin, Springer-Verlag, xii+362 pp.

Huber, B., Sturmfels, B. (1997). Bernshtein's theorem in affine space. *Discrete Comput. Geom.*, **17**, 137–141.

Kaltofen, E. (1992). *Polynomial Factorization 1987–1991* In Simon, I., ed., *Proc. Latin '92*, SLNCS **583**, pp. 294–313.

Kaltofen, E. (1995). Effective Noether irreducibility forms and applications. *J. Comput. Syst. Sci.*, **50**, 274–295.

Kapranov, M. M., Sturmfels, B., Zelevinsky, A. V. (1992). Chow polytopes and general resultants. *Duke Math. J.*, **67**, 189–218.

Kempf, G., Knudsen, F., Mumford, D., Saint-Donat, B. (1973). In *Toroidal Embeddings I*, LNM **339**. Berlin, Springer-Verlag.
Khovanskii, A. G. (1977). Newton polyhedra and toroidal varieties. *Funct. Anal. Appl.*, **11**, 289–296.
Khovanskii, A. G. (1991). *Fewnomials*. Providence, RI, AMS Press.
Landau, S., Miller, G. L. (1985). Solvability by radicals is in polynomial time. *J. Comput. System Sci.*, **30**, 179–208.
Lang, S. (1983). *Fundamentals of Diophantine Geometry*. New York, Springer.
Lenstra, H. W., (1997a) On the factorization of Lacunary polynomials, Number Theory in Progress, Proceedings of the 70th Birthday Meeting of A. Schinzel, W. de Gruyter, to appear.
Lenstra, H. W., (1997b) Finding Small Degree Factors of Lacunary Polynomials, Number Theory in Progress, Proceedings of the 70th Birthday Meeting of A. Schinzel, W. de Gruyter, to appear.
Malajovich-Muñoz, G., Zubelli, J. P. (1998). Tangent Graeffe Iteration. *Informes de Matematica Serie B-119, IMPA, Rio de Janeiro, Brasil, June.*
Mayr, E., Meyer, A. (1982). The complexity of the word problem for commutative semigroups and polynomial ideals. *Adv. Math.*, **46**, 305–329.
Mourrain, B., Pan, V. Y. (1998). Asymptotic acceleration of solving multivariate polynomial systems of equations. In *Proceedings of ACM STOC.*
Pedersen, P., Sturmfels, B. (1993). Product formulas for sparse resultants and chow forms. *Math. Z.*, **214**, 377–396.
Preparata, F. P., Shamos, M. I. (1985). *Computational Geometry: An Introduction.* Texts and Monographs in Computer Science. New York, Berlin, Springer-Verlag.
Renegar, J., (1989) On the worst case arithmetic complexity of approximating zeros of systems of polynomials. *SIAM J. Comput.*, **18**, 350–370
Rojas, J. M. (1994). A convex geometric approach to counting the roots of a polynomial system. *Theor. Comput. Sci.*, **133**, 105–140.
Rojas, J. M. (1997a). Toric laminations, sparse generalized characteristic polynomials, and a refinement of Hilbert's tenth problem. *Foundations of Computational Mathematics, selected papers of a conference, held at IMPA in Rio de Janeiro, January 1997.* Berlin, Springer-Verlag.
Rojas, J. M. (1997b). Affine elimination theory, extended abstract. In *Proc. Conf. in Honor of the 60th birthday of David A. Buchsbaum, Northeastern University, October, 1997.*
Rojas, J. M. (1998a). Intrinsic near quadratic complexity bounds for real multivariate root counting. In *Proc. Sixth Annual European Symposium on Algorithms*, LNCS **1461**. Berlin, Springer-Verlag.
Rojas, J. M. (1998b). The geometry of elimination I: degree formulae and the vanishing of resultants. Preprint, City University of Hong Kong.
Rojas, J. M. (1998c). The geometry of elimination II: affine elimination theory and better nullstellensatze. Preprint, City University of Hong Kong.
Rojas, J. M. (1999a). Toric intersection theory for affine root counting. *J. Pure Appl. Alg.*, **136**.
Rojas, J. M., On the complexity of diophantine geometry in low dimensions, *Proc. of the 31st Annual Symp. on Theory of Computing (STOC ,99), May 1–4, 1999, Atlanta, Georgia.* ACM Press, New York.
Rojas, J. M., Wang, X. (1996). Counting affine roots of polynomial systems via pointed Newton polytopes. *J. Complexity*, **12**, 116–133.
Schneider, R. (1994). *Convex Bodies: The Brunn-Minkowski Theory*, volume 44, Encyclopedia of Mathematics and its Applications. Cambridge, Cambridge University Press.
Shub, M. (1993). Some remarks on Bézout's theorem and complexity theory, from topology to computation. In *Proceedings of Smalefest*, pp. 443–455. Berlin, Springer-Verlag.
Sturmfels, B. (1991). Sparse Elimination Theory. In Eisenbud, D. and Robbiano, L., eds, *Proc. Computat. Algebraic Geom. and Commut. Algebra 1991, Cortona, Italy*, 1993, pp. 377–396. Cambridge, Cambridge University Press.
Sturmfels, B. (1994). On the Newton polytope of the resultant. *J. Alg. Combin.*, **3**, 207–236.
Sturmfels, B. (1998). Introduction to resultants, applications of computational algebraic geometry (San Diego, CA, 1997), 25–39. In *Proc. Symp. Applied Mathematics, 53.* Providence, RI, American Mathematical Society.
van der Waerden, B. L. (1950). *Modern Algebra,* 2nd edn. F. Ungar, New York, 1950.
Wampler, C. W. (1992). Bezout number calculations for multi-homogeneous polynomial systems. *Appl. Math. Comput.*, **51**, 143–157.
Weyman, J., Zelevinsky, A. (1994). Determinantal formulas for multigraded resultants. *J. Algebraic Geom.*, **3**, 569–597.
Ye, Y. (1994). Combining binary search and Newton's method to compute real roots for a class of real functions. *J. Complexity*, **10**, 271–280.