SPARSE TRACE TESTS

TAYLOR BRYSIEWICZ AND MICHAEL BURR

Abstract. We establish how the coefficients of a sparse polynomial system influence the sum (or the trace) of its zeros. As an application, we develop numerical tests for verifying whether a set of solutions to a sparse system is complete. These algorithms extend the classical trace test in numerical algebraic geometry. Our results rely on both the analysis of the structure of sparse resultants as well as an extension of Esterov's results on monodromy groups of sparse systems.

1. Introduction

The coordinate-wise sum of a finite set of points $S \subseteq \mathbb{C}^n$ is called its trace. When S is a subset of a general linear section of an irreducible variety, the behavior of its trace as the section moves determines whether S comprises the whole section. The following lemma makes this precise.

Lemma 1 ([15, Theorem 3.6])**.** Fix an irreducible variety $X \subseteq \mathbb{C}^n$ and a generic pencil of affine linear spaces L_t of complementary dimension. The trace of $X \cap L_t$ moves affine linearly in t. Conversely, the trace of any nonempty proper subset of $X \cap L_t$ moves nonlinearly in t.

The (classical) trace test [10, 12, 15] is a fundamental algorithm in numerical algebraic geometry [16, 17] which is used to verify the affine linear behavior of the trace in Lemma 1. Our main result is a sparse analogue to Lemma 1. More precisely, we identify a subset of the coefficients of a sparse polynomial system such that the trace is an affine linear function of this collection of coefficients. Conversely, we show that under simple conditions on the support, the traces of nonempty incomplete solution sets are nonlinear functions of this collection of coefficients. We use our results to produce what we call sparse trace tests.

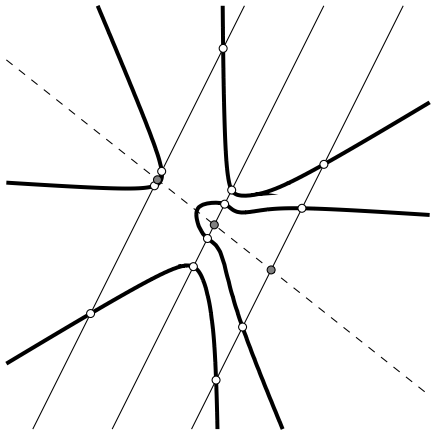

Figure 1. The centroids (gray) of the intersection points (white) of the quartic curve with three parallel lines. These averages move in a line (dashed). We use averages instead of sums in this figure to keep the image small.

Burr was partially supported by grants from the National Science Foundation (CCF-1527193 and DMS-1913119).

The proof of the forward direction of Lemma 1 rests upon the following elementary result dating back to Newton. The trace of the zeros of $f = c_0 + c_1x + c_2x^2 + \cdots + c_{d-1}x^{d-1} + c_dx^d \in \mathbb{C}[x]$ equals $-c_{d-1}/c_d$, whenever $c_d \neq 0$ and zeros are counted with multiplicity. In particular, the trace of the zeros of a polynomial in one variable is an affine linear function of the coefficient c_{d-1} and does not depend on the coefficients of lower-degree monomials, see Figure 2.

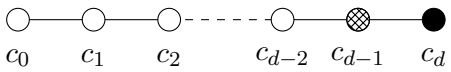

FIGURE 2. The support of a univariate polynomial f. Identifying each point in the figure with the corresponding coefficient of f, the figure displays those which do not affect the trace (white), those which affect the trace affine linearly (cross-hatched), and those which influence the trace nonlinearly (filled).

We generalize Newton's result. Given a tuple $\mathcal{A} = (A_1, \ldots, A_n)$ of monomial supports $A_i \subseteq \mathbb{Z}^n$, we consider sparse polynomial systems $\mathcal{F} = (f_1, \ldots, f_n) \in (\mathbb{C}[x_1, \ldots, x_n])^n$ supported on $\mathcal{A}$. Any such polynomial system is identified with its coefficients in $\mathbb{C}^{\mathcal{A}} := \mathbb{C}^{|\mathcal{A}|}$. The polyhedral geometry of $\mathcal{A}$ controls many aspects of the solutions to $\mathcal{F} = 0$ in the algebraic torus $(\mathbb{C}^\times)^n$ (see [3]), and we study the trace of this solution set. As in the univariate case, the trace of these solutions is a rational function of the coefficients of $\mathcal{F}$. Using a simple discrete geometric construction, we identify a large collection of monomials of $\mathcal{A}$ for which the trace is an affine linear function of the corresponding collection of coefficients. We illustrate an example in Figure 3 where the similarity to Figure 2 is immediate.

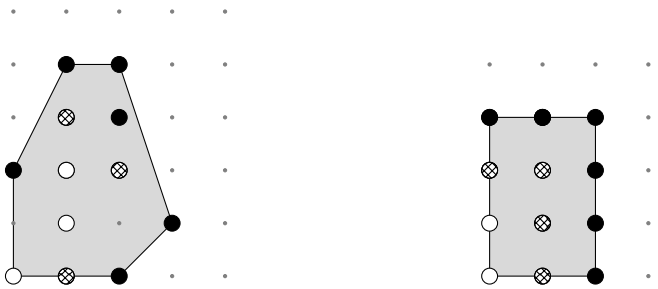

FIGURE 3. The support of $\mathcal{F} = (f_1, f_2)$. Identifying each point in the figure with the corresponding coefficient in $\mathcal{F}$, our results determine a collection of those which do not affect the trace (white), a collection of those which may affect the trace affine linearly (cross-hatched), and a collection of those which may influence the trace nonlinearly (filled).

The backward direction of Lemma 1 is more sophisticated. The proof in [12] relies on the fact that the monodromy group of the branched cover $\pi : \{(x, t) \in \mathbb{C}^n_x \times \mathbb{C}_t \mid x \in X \cap L_t\} \to \mathbb{C}_t$ is the full symmetric group. In the sparse setting, the relevant branched cover is the map

$$\pi_{\mathcal{A}} : \{(x, \mathcal{F}) \in (\mathbb{C}^\times)^n \times \mathbb{C}^{\mathcal{A}} \mid \mathcal{F}(x) = 0\} \to \mathbb{C}^{\mathcal{A}},$$

restricted to those sparse systems where some subset of the coefficients are generic and fixed, and the rest vary. Given a set $\mathcal{A}$ of supports, we provide a simple condition on the varying coefficients which guarantees that the monodromy group of this restriction is the full symmetric group.

The classical and sparse trace tests are instances of what we call completeness tests, which may be used to decide if a nonempty subset of a zero-dimensional variety is proper. When implemented

using numerical computations, such tests return the correct answer almost surely, and we call them numerical completeness tests. For the classical trace test, the zero-dimensional polynomial system is the pair consisting of the equations for an irreducible variety and a generic linear system of complementary dimension. We expect completeness tests to be developed beyond the sparse case; for example, to generalizations of witness sets in numerical algebraic geometry [10, 18].

Our results, although motivated by applications in numerical algebraic geometry, have consequences in symbolic algebra. For example, our results show how to reduce a polynomial system while preserving its trace. This reduction leads to smaller systems which improves the efficiency of both numerical and symbolic computations.

1.1. **Overview.** Since our main results can be stated, understood, and applied without proof, we present them in Section 2 along with the relevant ideas and notation. There, we present two algorithms, which we call sparse trace tests. The detailed background and proofs reside in Sections 3 and 4. In particular, we use sparse resultants in Section 3 to determine coefficients of a sparse system which do not appear in a formula for the trace. In Section 4 we investigate the monodromy groups of restrictions of sparse polynomial systems to establish when traces of incomplete solution sets move nonlinearly. Section 4 also addresses exceptional cases for which the sparse trace tests cannot be applied. In Section 5, we present a gallery of examples showcasing applications of our theory and algorithms.

2. Notation, background, and main results

We introduce the required background and notation pertaining to sparse polynomial systems and state our main results. We present two sparse trace tests (Algorithms 1 and 2), and outline a proof of correctness for these algorithms, relying upon results from Sections 3 and 4.

2.1. **Sparse polynomial systems.** Let e_j denote the j-th standard coordinate vector of the integer lattice $\mathbb{Z}^n$ and $\mathbf{0}$ the origin of $\mathbb{Z}^n$. For a finite subset $A \subseteq \mathbb{Z}^n$, we consider (Laurent) polynomials of the form

$$f(x) = f(x_1, \ldots, x_n) = \sum_{\alpha \in A} c_\alpha x_1^{\alpha_1} \cdots x_n^{\alpha_n} = \sum_{\alpha \in A} c_\alpha x^\alpha \in \mathbb{C}[x^{\pm 1}] := \mathbb{C}[x_1, \ldots, x_n, x_1^{-1}, \ldots, x_n^{-1}],$$

where $c_\alpha \in \mathbb{C}$. We say that f is *supported on* A and its *support* is $\operatorname{supp}(f) = \{\alpha \in A \mid c_\alpha \neq 0\}$. Identifying f with its coefficients $\{c_\alpha\}_{\alpha \in A}$, we write $f \in \mathbb{C}^A := \mathbb{C}^{|A|}$. We define $L[A]$ to be the lattice generated by all differences of points in A. Given any sublattice L of $\mathbb{Z}^n$, its rank is $\operatorname{rk}(L)$, and we define $\operatorname{rk}(A) := \operatorname{rk}(L[A])$. When L is full rank, we write $[\mathbb{Z}^n : L]$ for the index of L in $\mathbb{Z}^n$.

For a fixed collection $\mathcal{A} = (A_1, \ldots, A_N)$, a tuple $\mathcal{F} = (f_1, \ldots, f_N)$ of Laurent polynomials in $\mathbb{C}[x^{\pm 1}]$ is called a *sparse polynomial system supported on* $\mathcal{A}$ whenever each f_i is supported on A_i. In this case, we write $f_i = \sum_{\alpha \in A_i} c_{i,\alpha} x^\alpha$. As with a single polynomial, we identify $\mathcal{F}$ with its coefficients in $\mathbb{C}^{\mathcal{A}} := \mathbb{C}^{A_1} \times \cdots \times \mathbb{C}^{A_N}$. We apply set operations to collections of supports coordinate-wise. For example, given $\mathcal{B} = (B_1, \ldots, B_N)$ and $\mathcal{C} = (C_1, \ldots, C_N)$, we write $\mathcal{B} \cup \mathcal{C}$ for $(B_1 \cup C_1, \ldots, B_N \cup C_N)$. When $\mathcal{A} = \mathcal{B} \sqcup \mathcal{C}$, we may decompose $\mathcal{F} = \mathcal{F}_\mathcal{B} + \mathcal{F}_\mathcal{C} \in \mathbb{C}^\mathcal{B} \times \mathbb{C}^\mathcal{C} = \mathbb{C}^\mathcal{A}$ in order to refer to the coefficients of $\mathcal{F}$ indexed by $\mathcal{B}$ and $\mathcal{C}$ individually. We define $L[\mathcal{A}]$ to be the lattice generated by $\bigcup_{i=1}^n L[A_i]$. When each A_i has full rank, we say that $\mathcal{A}$ is *abundant*.

The set of zeros in $\mathbb{C}^n$ of a polynomial system is denoted by $\mathcal{V}(\mathcal{F})$. In the sparse setting, we consider zeros of $\mathcal{F} \in \mathbb{C}^\mathcal{A}$ in the *algebraic torus* $(\mathbb{C}^\times)^n := \{x \in \mathbb{C}^n \mid x_i \neq 0 \text{ for } i = 1, \ldots, n\}$, and so we decorate the notation $\mathcal{V}$ as $\mathcal{V}^\times(\mathcal{F}) := \{x \in (\mathbb{C}^\times)^n \mid f_i(x) = 0 \text{ for } i = 1, \ldots, n\}$.

To refer to the entire family of sparse polynomial systems supported on $\mathcal{A}$, we replace the coefficients $c_{i,\alpha}$ with parameters $u_{i,\alpha}$ and write $\mathcal{F}(u) \in \mathbb{C}[u_{i,\alpha}][x^{\pm 1}]$. This system is called the *universal polynomial system over* $\mathcal{A}$. The following *incidence variety* encodes the structure of the

family of polynomial systems $\mathbb{C}^{\mathcal{A}}$ and their zeros:

$$X_{\mathcal{A}} := \{(x, \mathcal{F}) \mid x \in \mathcal{V}^{\times}(\mathcal{F}), \mathcal{F} \in \mathbb{C}^{\mathcal{A}}\} \subseteq (\mathbb{C}^{\times})^n \times \mathbb{C}^{\mathcal{A}}.$$

In other words, $X_{\mathcal{A}}$ is the zero set of $\mathcal{F}(u)$ in $(\mathbb{C}^{\times})^n \times \mathbb{C}^{\mathcal{A}}$. Any particular solution set $\mathcal{V}^{\times}(\mathcal{F})$ is naturally identified with the fibre over $\mathcal{F}$ with respect to the projection $\pi_{\mathcal{A}}$ on the second factor. The fibre over a point $x \in (\mathbb{C}^{\times})^n$ with respect to the projection $\pi_{\mathcal{V}}$ on the first factor is identified with all sparse systems having x as a solution.

We say that $\mathcal{A}$ is a *square* set of supports when $N = n$. Similarly, systems supported on square sets of supports are called *square systems*. When $\mathcal{A}$ is square, its mixed volume $\mathrm{MV}(\mathcal{A})$ is defined to be the mixed volume of the convex hulls $\{\mathrm{conv}(A_i)\}_{i=1}^n$. Unless otherwise stated, $\mathcal{A}$ refers to a *square set of supports with* $\mathrm{MV}(\mathcal{A}) > 0$. We repeat these assumptions in all results for completeness.

The discrete geometry of $\mathcal{A}$ controls much of the complex geometry of $X_{\mathcal{A}}$. In particular, for square systems, the following theorem relates the mixed volume of $\mathcal{A}$ to the cardinality of a generic fibre of $\pi_{\mathcal{A}}$.

Theorem 2 (BKK [1, 11])**.** Let $\mathcal{A}$ be a square set of supports with $\mathrm{MV}(\mathcal{A}) > 0$. The cardinality of $\mathcal{V}^{\times}(\mathcal{F})$ is $\mathrm{MV}(\mathcal{A})$ for all $\mathcal{F}$ outside of a Zariski closed subset of $\mathbb{C}^{\mathcal{A}}$.

Geometrically, Theorem 2 states that when $\mathrm{MV}(\mathcal{A}) > 0$, the map $\pi_{\mathcal{A}}$ is a *branched cover* of degree $\mathrm{MV}(\mathcal{A})$. The Zariski closed subset of $\mathbb{C}^{\mathcal{A}}$ for which the fibres do not have cardinality $\mathrm{MV}(\mathcal{A})$ is called the *branch locus* of $\pi_{\mathcal{A}}$. We call a system $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$ *Bernstein-generic* whenever it is not in this branch locus. Bernstein-generic systems have exactly $\mathrm{MV}(\mathcal{A})$-many isolated zeros, each of multiplicity one, in the algebraic torus. More generally, systems which are not Bernstein-generic have at most $\mathrm{MV}(\mathcal{A})$-many isolated zeros in the algebraic torus, counted with multiplicity.

Restricted to the complement of the branch locus, the map $\pi_{\mathcal{A}}$ is topologically a $\mathrm{MV}(\mathcal{A})$-to-one covering space. Consequently, a loop $\gamma : [0, 1] \to \mathbb{C}^{\mathcal{A}}$ based at $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$ which avoids the branch locus may be uniquely lifted to $\mathrm{MV}(\mathcal{A})$-many paths in $X_{\mathcal{A}}$. These paths connect points in $\pi_{\mathcal{A}}^{-1}(\mathcal{F})$, inducing a permutation on the fibre. The collection of all such permutations form a subgroup $G(\pi_{\mathcal{A}})$ of the symmetric group on $\pi_{\mathcal{A}}^{-1}(\mathcal{F})$. The group $G(\pi_{\mathcal{A}})$ encodes global information about the symmetries of the branched cover and is called the *monodromy group* of $\pi_{\mathcal{A}}$. It does not depend on the base point $\mathcal{F}$ and is well-defined up to relabeling the points of $\pi_{\mathcal{A}}^{-1}(\mathcal{F})$.

2.2. **Traces.** Given a subset $S \subseteq \mathbb{C}^n$, its coordinate-wise sum $\Sigma(S) := (\Sigma_1(S), \ldots, \Sigma_n(S))$ is called its *trace*. In the numerical algebraic geometry literature (see, for example, [15]), the coordinate-wise *average* $\mu(S) := (\mu_1(S), \ldots, \mu_n(S)) = \frac{\Sigma(S)}{|S|}$ of S has been traditionally called its trace. To avoid ambiguity, we call $\mu(S)$ the *centroid* of S.

Without loss of generality, *we focus on the first coordinate* $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}))$. Analogous statements about the other coordinates of $\Sigma(\mathcal{V}^{\times}(\mathcal{F}))$ may be gleaned from our results by the interested reader.

The authors of [5] show that $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}(u)))$ is a rational function of the coefficients of $\mathcal{F}(u)$, expressed in terms of sparse resultants (see Section 3 for additional details). We assume that this fraction has been reduced so that the numerator and denominator do not share any nontrivial factors. This formula involves some, but not all, of the coefficients of $\mathcal{F}(u)$. We distinguish monomials in $\mathcal{A}$ depending on how they appear in the formula for $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}(u)))$.

Definition 3. Given a square set of supports $\mathcal{A}$ with $\mathrm{MV}(\mathcal{A}) > 0$, we define the following:

- The *unnecessary support* $\mathcal{U} \subseteq \mathcal{A}$ consists of those monomials whose coefficients do not appear in the formula for $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}(u)))$. The complement $\mathcal{N} = \mathcal{A} \backslash \mathcal{U}$ of $\mathcal{U}$ is the *necessary support*.
- The *affine linear support* $\mathcal{AL} \subseteq \mathcal{N} \subseteq \mathcal{A}$ consists of those monomials in the necessary support whose coefficients only appear in the numerator, and do so to degree one.
- The *nonlinear support* $\mathcal{NL} \subseteq \mathcal{N} \subseteq \mathcal{A}$, consists of all other necessary monomials, that is, $\mathcal{NL} = \mathcal{N} \backslash \mathcal{AL}$.

These sets partition $\mathcal{A}$ as $\mathcal{A} = \mathcal{U} \sqcup \mathcal{N} = \mathcal{U} \sqcup (\mathcal{A}\mathcal{L} \sqcup \mathcal{N}\mathcal{L})$. Given two sparse systems $\mathcal{F}, \mathcal{G} \in \mathbb{C}^{\mathcal{A}}$ and a subset $\mathcal{C} \subseteq \mathcal{A}$, we write $\mathcal{F} \approx_{\mathcal{C}} \mathcal{G}$ when $\mathcal{F}_{\mathcal{C}} = \mathcal{G}_{\mathcal{C}}$, that is, the coefficients of $\mathcal{F}$ and $\mathcal{G}$ agree over the support $\mathcal{C}$.

Definition 4. A subset $\mathcal{B} \subseteq \mathcal{A}$ is called *trace-affine-linear (TAL)* if the trace $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}))$ is an affine linear function of the coefficients indexed by $\mathcal{B}$.

Remark 5. Any singleton consisting of an element of $\mathcal{U} \cup \mathcal{A}\mathcal{L}$ is TAL. Moreover, by definition, any TAL subset must be contained in $\mathcal{U} \cup \mathcal{A}\mathcal{L}$. However, the converse is not true: if two coefficients c_α, c_β appear to degree one in the same monomial in the numerator of the formula for $\Sigma_1(\mathcal{V}^{\times}(\mathcal{F}(u)))$, then the set $\{\alpha, \beta\}$ is not TAL despite being contained in $\mathcal{U} \cup \mathcal{A}\mathcal{L}$.

The following lemma provides the analogue of the forward direction of Lemma 1. It is a direct consequence of Definitions 3 and 4.

Lemma 6. Suppose $\mathcal{A}$ is a square set of supports with $\mathrm{MV}(\mathcal{A}) > 0$ and $\mathcal{B} \subseteq \mathcal{A}$ is a TAL subset of $\mathcal{A}$. Fix Bernstein-generic $\mathcal{F}, \mathcal{G} \in \mathbb{C}^{\mathcal{A}}$ such that $\mathcal{F} \approx_{\mathcal{A}\setminus\mathcal{B}} \mathcal{G}$ (respectively, $\mathcal{F} \approx_{\mathcal{N}} \mathcal{G}$). Then the trace Σ_1 of $\mathcal{V}^{\times}(t\mathcal{F} + (1-t)\mathcal{G})$ is an affine linear (respectively, constant) function of t, restricted to the set of t's where $t\mathcal{F} + (1-t)\mathcal{G}$ is Bernstein-generic.

Example 7. Let

$$\mathcal{F} = \begin{Bmatrix} f_1 : & 3x^2y^4 + 2x^2y^3 - xy^4 + x^3y + 5x^2y^2 + xy^3 + 2x^2 + 4y^2 + 9x \\ f_2 : & 4x^2y^3 + x^2y^2 + 8xy^3 - 4x^2y - 2xy^2 + 3y^3 + 4x^2 + 5xy + 4x - y - 9 \end{Bmatrix}, \tag{1}$$

and

$$\mathcal{G} = \begin{Bmatrix} g_1 : & 3x^2y^4 + 2x^2y^3 - 3xy^4 + x^3y + 3x^2y^2 - 2xy^3 - 3xy^2 + 2x^2 + xy + y^2 + 4x - 5 \\ g_2 : & 4x^2y^3 + x^2y^2 + 3xy^3 - 4x^2y - 4xy^2 + 5y^3 + 4x^2 + 3xy - 4y^2 - x - 2y - 5 \end{Bmatrix}. \tag{2}$$

These systems are supported on the monomials depicted in Figure 3.

The polynomials f_1 and g_1 agree on the coefficients of $\{x^2y^4, x^2y^3, x^3y, x^2\}$, as do f_2 and g_2 with respect to the coefficients of $\{x^2y^3, x^2y^2, x^2y, x^2\}$. We show in Example 9 that the complement of this collection of coefficients is TAL with respect to x. In Table 1, we calculate Σ_1 and Σ_2 for $t\mathcal{F} + (1-t)\mathcal{G}$ for several values of t. Through direct inspection, we see that Σ_1 appears to be an affine linear function of t, whereas Σ_2 is not. Therefore, by Lemma 6, we conclude that the complement of this set of coefficients is not TAL with respect to y.

t	0	1	2	3	4	5
Σ_1	3.922	-0.578	-5.078	-9.578	-14.078	-18.578
Σ_2	-0.200	-0.523	-8.135	5.772	1.974	1.236

TABLE 1. The traces for $\mathcal{V}^{\times}(t\mathcal{F} + (1-t)\mathcal{G})$ where $\mathcal{F}$ and $\mathcal{G}$ are Systems (1) and (2) in Example 7.

2.3. **Sparse trace tests.** We first recall the classical trace test [12] since it serves as a model for the sparse trace test. The input to the algorithm is an irreducible variety X, a general family of parallel linear spaces L_t of complementary dimension, and a subset $S \subseteq X \cap L_0$. The algorithm uses *homotopy continuation* to construct an analytic continuation $S_t \subseteq X \cap L_t$ of $S = S_0$ as t varies. This process tracks the points of S to points of $X \cap L_t$ for generic $t \in \mathbb{C}$. The values of $\Sigma(S_0)$, $\Sigma(S_{1/2})$, and $\Sigma(S_1)$ are compared, and, after appealing to Lemma 1, these three traces are collinear if and only if $S = X \cap L_0$.

A *completeness test* is an algorithm whose input is a zero-dimensional polynomial system $\mathcal{F}$ and a nonempty subset S of its zeros. It returns the output `pass` if S is *complete*, that is, if $S = \mathcal{V}(\mathcal{F})$. It returns `fail` if $S \neq \mathcal{V}(\mathcal{F})$. The main application of the classical trace test is to decide whether

Algorithm 1: Sparse trace test

Input: • A collection of supports $\mathcal{A}$ in $\mathbb{Z}^n$ with $\mathrm{MV}(\mathcal{A}) > 0$ and $L[\mathcal{A}] = \mathbb{Z}^n$
• A Bernstein-generic polynomial system $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$
• $\emptyset \neq S \subseteq \mathcal{V}^{\times}(\mathcal{F})$
• A collection $\mathcal{B} \subseteq \mathcal{A}$ which is TAL and abundant.

Output: If $S = \mathcal{V}^{\times}(\mathcal{F})$, then `pass`, else `fail`.

Construct $\mathcal{G} \in \mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{A}\setminus\mathcal{B}}$ by choosing generic $\mathcal{G}_{\mathcal{B}} \in \mathbb{C}^{\mathcal{B}}$.
Use homotopy continuation to follow the analytic continuation $S_t \subseteq \mathcal{V}^{\times}(t\mathcal{F} + (1-t)\mathcal{G})$ of $S = S_1$ to $t = 1/2$ and $t = 0$.
Compute $\Sigma_1(S_0)$, $\Sigma_1(S_{1/2})$, and $\Sigma_1(S_1)$.
if $(0, \Sigma_1(S_0))$, $(1/2, \Sigma_1(S_{1/2}))$ *and* $(1, \Sigma_1(S_1))$ *are collinear* **then**
    return `pass`
else
    return `fail`

a nonempty subset S of a general linear section $X \cap L$ of an irreducible variety is complete. In this sense, the classical trace test is a completeness test for irreducible varieties.

Following the approach of the classical trace test, we introduce a *sparse trace test* in Algorithm 1. Using the stronger conditions of Lemma 6, we describe a second sparse trace test, Algorithm 2, which we call the *constant sparse trace test*. We refer to these algorithms as the *sparse trace tests*.

By Lemma 6, the sparse trace tests return `pass` when $S = \mathcal{V}^{\times}(\mathcal{F})$, even when $\mathcal{B}$ is not abundant or $L[\mathcal{A}] \neq \mathbb{Z}^n$. Thus, they return `fail` only when $S \neq \mathcal{V}^{\times}(\mathcal{F})$. However, from a practical point of view, even this one-sided usage requires *a priori* knowledge of some subset $\mathcal{B} \subseteq \mathcal{A}$ which is either TAL or contained in the unnecessary support. In Section 2.4, we give simple discrete geometric descriptions of sets satisfying these relationships. The conditions that $\mathcal{B}$ is abundant and $L[\mathcal{A}] = \mathbb{Z}^n$ make the algorithm two-sided and are easy to check. In Section 2.5 we justify the inclusion of these additional conditions on $\mathcal{B}$.

2.4. **Approximating the support.** The use of either sparse trace test requires valid $\mathcal{A}$ and $\mathcal{B}$ as input. It is straightforward to check the conditions that $\mathcal{B}$ is abundant and $L[\mathcal{A}] = \mathbb{Z}^n$, but finding a candidate for $\mathcal{B}$ is, *a priori*, not obvious. We provide easy-to-compute subsets of $\mathcal{A}$ which are TAL or contained in $\mathcal{U}$, which may be used for $\mathcal{B}$.

Algorithm 2: Constant sparse trace test

Input: • A collection of supports $\mathcal{A}$ in $\mathbb{Z}^n$ with $\mathrm{MV}(\mathcal{A}) > 0$ and $L[\mathcal{A}] = \mathbb{Z}^n$
• A Bernstein-generic polynomial system $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$
• $\emptyset \neq S \subseteq \mathcal{V}^{\times}(\mathcal{F})$
• A collection $\mathcal{B} \subseteq \mathcal{U}$ which is abundant.

Output: If $S = \mathcal{V}^{\times}(\mathcal{F})$, then `pass`, else `fail`.

Construct $\mathcal{G} \in \mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{A}\setminus\mathcal{B}}$ by choosing generic $\mathcal{G}_{\mathcal{B}} \in \mathbb{C}^{\mathcal{B}}$.
Construct $|S|$-many solutions $S' \subseteq \mathcal{V}^{\times}(\mathcal{G})$.
if $\Sigma_1(S) = \Sigma_1(S')$ **then**
    return `pass`
else
    return `fail`

Let $A \subseteq \mathbb{Z}^n$ be finite. The set of *k-offset* points of A in the *first coordinate* is defined to be

$$\text{offset}(A,k) = \{\alpha \in A \mid \alpha + (k+\varepsilon)e_1 \notin \text{conv}(A) \text{ for all } \varepsilon > 0\}.$$

These sets are increasing, that is, if $k < k'$, then $\text{offset}(A,k) \subseteq \text{offset}(A,k')$. We note that $\text{offset}(A,0)$ consists of the points of A which maximize a linear functional $x \mapsto \langle x, \omega\rangle$ for some $\omega \in \mathbb{R}^n$ with $\omega_1 > 0$. For a collection $\mathcal{A}$, we define $\text{offset}(\mathcal{A},k)$ coordinate-wise.

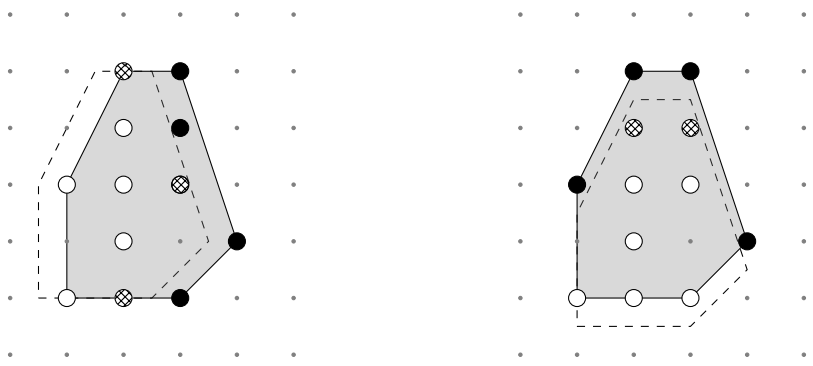

(a) (b)

FIGURE 4. The sets $\text{offset}(A,.5)$ and $\text{offset}(A,1)$ in the (a) first and (b) second coordinates along with the convex hull of A. The points of $\text{offset}(A,.5)$ are filled, the points of $\text{offset}(A,1)\setminus\text{offset}(A,0.5)$ are cross-hatched, and the remaining points are white. The shifted polytope illustrates the behavior for multiple points at once.

The following specialization of Theorem 22 provides discrete geometric constructions of subsets of $\mathcal{A}$ which are either TAL or contained in $\mathcal{U}$.

Lemma 8. Let $\mathcal{A} = (A_1, \ldots, A_n)$ be a collection of supports in $\mathbb{Z}^n$ such that $\text{MV}(\mathcal{A}) > 0$. Then,

- $\mathcal{A}\setminus\text{offset}(\mathcal{A},0.5)$ is TAL and
- $\mathcal{A}\setminus\text{offset}(\mathcal{A},1) \subseteq \mathcal{U}$.

Example 9. We continue the discussion of Systems (1) and (2) from Example 7. Since the coefficients of $\mathcal{F}$ and $\mathcal{G}$ which agree contain $\text{offset}(\mathcal{A},0.5)$ with respect to the first coordinate, Lemmas 6 and 8 imply that the trace $\Sigma_1(\mathcal{V}^\times(t\mathcal{F}+(1-t)\mathcal{G}))$ is an affine linear function in the remaining coefficients. On the other hand, since the coefficients of f_1 and g_1 corresponding to the monomials $\{y^2, xy^4, x^2y^4\}$ disagree, Lemma 8 does not apply with respect to the second coordinate. Table 1 shows the linearity of the trace in the first coordinate and nonlinearity of the trace in the second coordinate.

2.5. **Incomplete solution sets.** We explain how the conditions $\mathcal{B}$ is abundant and $L[\mathcal{A}] = \mathbb{Z}^n$, on $\mathcal{A}$ and $\mathcal{B}$, guarantee that the output `pass` from the sparse trace tests implies the equality $S = \mathcal{V}^\times(\mathcal{F})$. Our argument relies on results about the monodromy group $G(\pi_{\mathcal{A}})$ detailed in Section 4, and is structured as follows: Suppose that $\mathcal{F}$ is Bernstein-generic, and $\mathcal{G}$ in $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{A}\setminus\mathcal{B}}$ is generic. When the monodromy group of $\pi_{\mathcal{A}}$ restricted to the preimage of the line containing $\mathcal{F}$ and $\mathcal{G}$ is the full symmetric group, there is a loop which induces the transposition swapping $s \in S$ and $s' \in \mathcal{V}^\times(\mathcal{F})\setminus S$. When the points in $\mathcal{V}^\times(\mathcal{F})$ have distinct first coordinates, $\Sigma_1(S) \neq \Sigma_1(S\cup\{s'\}\setminus\{s\})$. In these cases, since the trace of the analytic continuation of $S \neq \mathcal{V}^\times(\mathcal{F})$ is continuous along the corresponding monodromy loop, the traces along this loop cannot lie on a line.

Theorem 10. Let $(\mathcal{A},\mathcal{F},S,\mathcal{B})$ be the input to Algorithm 1 or 2. Then the algorithm returns `pass` if and only if $S = \mathcal{V}^\times(\mathcal{F})$.

Proof. If $\mathcal{B}$ is TAL (respectively, $\mathcal{B} \subseteq \mathcal{U}$), then Lemma 6 implies that Algorithm 1 (respectively, Algorithm 2) returns `pass` when $S = \mathcal{V}^\times(\mathcal{F})$ is given as input. It remains to be shown that the algorithms return `fail`, almost surely, when $S \neq \mathcal{V}^\times(\mathcal{F})$.

The conditions on the input to Algorithms 1 or 2 imply that the monodromy group $\pi_{\mathcal{A}}$ restricted to $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{A}\setminus\mathcal{B}}$ is the full symmetric group by Theorem 41. Zariski showed that the further restriction to $t\mathcal{F}+(1-t)\mathcal{G}$ preserves the monodromy group [22]. Theorem 42 implies that the first coordinates of points in $\mathcal{V}^\times(\mathcal{F})$ are distinct. By the argument above the statement of the theorem, the algorithms return `fail`. □

Remark 11. The condition that $\mathcal{B}$ is abundant simplifies the proofs in Section 4. Certainly, less restrictive conditions also suffice, as illustrated by the the classical trace test. Finding precise conditions is left for future research.

Section 4.1 details why the condition $L[\mathcal{A}] = \mathbb{Z}^n$ is necessary. The following example illustrates one type of issue that arises when $L[\mathcal{A}] \neq \mathbb{Z}^n$.

Example 12. Suppose that $\mathcal{F} = (f_1, f_2)$ is a Bernstein-generic system supported on the support $\mathcal{A} = (A, A)$ where $A = \{(0,0), (2,0), (4,0), (3,1), (0,2), (2,2)\}$, as depicted in Figure 5. We note that the index of $L[\mathcal{A}]$ in $\mathbb{Z}^2$ is 2. For any solution $(s_1, s_2) \in \mathcal{V}^\times(\mathcal{F})$, the point $(-s_1, -s_2)$ is also a solution. Hence, if $S \subseteq \mathcal{V}^\times(\mathcal{F})$ consists of pairs of the form (s_1, s_2) and $(-s_1, -s_2)$, then the trace of S is zero. In particular, the trace is constant under a perturbation of any of the coefficients and the sparse trace tests cannot recognize that S might not be complete.

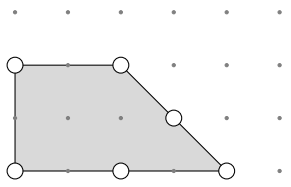

A

FIGURE 5. Support A such that the sparse trace tests cannot be applied on a polynomial system supported on (A, A).

3. Sparse resultants

The efficient application of the sparse trace tests requires finding valid supports $\mathcal{B}$ for input. Lemma 8 provides simple candidates. It is a specialization of Theorem 22, proven here using the theory of sparse resultants.

3.1. **Supports and sparse resultants.** Let $\mathcal{A} = (A_1, \ldots, A_N)$ be a collection of finite subsets $A_i \subseteq \mathbb{Z}^n$. Note that n and N are not required to be equal. For $I \subseteq [N] := \{1, 2, \ldots, N\}$, we write $\mathcal{A}_I$ for the subset $\{A_i\}_{i\in I}$. The *defect* of a subset $\mathcal{A}_I \subseteq \mathcal{A}$ is $\operatorname{def}(\mathcal{A}_I) := \operatorname{rk}(\mathcal{A}_I) - |I|$. When $n = N$, the following theorem of Minkowski characterizes when the mixed volume of $\mathcal{A}$ is nonzero.

Lemma 13. The mixed volume of $\mathcal{A} = (A_1, \ldots, A_n)$ is nonzero if and only if $\operatorname{def}(\mathcal{A}_I) \geq 0$ for all $I \subseteq [n]$.

We use the following generalization of Minkowski's result in Lemma 39 of Section 4.2.

Lemma 14. For supports $\mathcal{A} = (A_1, \ldots, A_k)$, $A_i \subseteq \mathbb{Z}^n$, and generic $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$, the dimension of $\mathcal{V}^\times(\mathcal{F})$ is $n - k$ when $\operatorname{def}(\mathcal{A}_I) \geq 0$ for all $I \subseteq [k]$. Otherwise, $\mathcal{V}^\times(\mathcal{F}) = \emptyset$.

A collection $\mathcal{A}$ is *essential* if its defect is -1 and every proper subset has nonnegative defect. Let $Q = \{\mathbf{0}, e_1\} \subseteq \mathbb{Z}^n$. When $n = N$ and $\mathrm{MV}(\mathcal{A}) > 0$, there is a unique subset $\mathcal{A}' \subseteq \mathcal{A}$ such that $\{Q\} \cup \mathcal{A}'$ is essential [5, 21].

For a square polynomial system $\mathcal{F} = \{f_1, \ldots, f_N\}$ supported on $\mathcal{A}$, we define $Z(\mathcal{A})$ to be the set of polynomial systems in $\mathbb{C}^{\mathcal{A}}$ which have a solution in $(\mathbb{C}^\times)^n$. When $\mathcal{A}$ has a unique essential subset, the Zariski closure of $Z(\mathcal{A})$ is a hypersurface in $\mathbb{C}^{\mathcal{A}}$ defined by an irreducible polynomial in $\mathbb{Z}[u_{i,\alpha}]$, see [21]. In particular, this polynomial is unique, up to a nonzero constant. Following the terminology in [4], we call this polynomial the *$\mathcal{A}$-eliminant* or *sparse eliminant* and denote it by $\mathrm{Elim}_{\mathcal{A}}$. If $\mathcal{A}$ does not have a unique essential subset, then we define $\mathrm{Elim}_{\mathcal{A}} = 1$.

Historically, the sparse eliminant has been referred to as the *sparse resultant*. We use the following redefinition of the sparse resultant (as in [13]) which produces more uniform statements [4]. When $\mathrm{Elim}_{\mathcal{A}} \neq 1$, the restriction of $\pi_{\mathcal{A}}$ to $\pi_{\mathcal{A}}^{-1}(Z(\mathcal{A}))$ is generically $d_{\mathcal{A}}$-to-one. We define the *$\mathcal{A}$-resultant* or *sparse resultant* to be

$$\mathrm{Res}_{\mathcal{A}} = \mathrm{Elim}_{\mathcal{A}}^{d_{\mathcal{A}}}.$$

For the universal polynomial system $\mathcal{F}(u)$ over $\mathcal{A}$, the sparse resultant $\mathrm{Res}_{\mathcal{A}}(\mathcal{F}(u))$ is a polynomial in $\mathbb{Z}[u_{i,\alpha}]$. Therefore, it can be evaluated at specific coefficients $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$. We write this in any of the following ways:

$$\mathrm{Res}_{\mathcal{A}}(\mathcal{F}) = \mathrm{Res}_{\mathcal{A}}(f_0, \ldots, f_n) = \mathrm{Res}_{\mathcal{A}}(\{c_{i,\alpha}\}).$$

3.2. **Approximation of the necessary support.** Let $\mathcal{A}$ be a square set of supports with $\mathrm{MV}(\mathcal{A}) > 0$. In the hidden variable technique, we view $\mathcal{F}(u)$ in the ring $\mathbb{C}[x_1][x_2, \ldots, x_n]$ by treating x_1 as a coefficient. This turns $\mathcal{F}(u)$ into a system of n equations in $n-1$ variables, supported on the collection of projections $\pi(\mathcal{A}) = (\pi(A_1), \ldots, \pi(A_n))$ under the forgetful map $\pi : \mathbb{Z}^n \to \mathbb{Z}^{n-1}$ which forgets the first coordinate. We let $\mathcal{G}(v)$ be the universal system over the support $\pi(\mathcal{A})$.

Lemma 15. If $\mathcal{A}$ is a square set of supports and $\mathrm{MV}(\mathcal{A}) > 0$, then $\pi(\mathcal{A})$ contains a unique essential subset.

Proof. Since $\mathrm{MV}(\mathcal{A}) > 0$, the defect of any subset of $\mathcal{A}$ is at least zero by Lemma 13. On the other hand, since $|\mathcal{A}| = n$, the defect of $\mathcal{A}$ is at most zero and so $\mathrm{rk}(\mathcal{A}) = n$. Since $\mathrm{rk}(\mathcal{A}) = n$, we must have that $\mathrm{rk}(\pi(\mathcal{A})) = n - 1$ implying that $\pi(\mathcal{A})$ has defect -1. Thus, $\pi(\mathcal{A})$ contains at least one essential subset.

To see that this subset must be unique, suppose that $\pi(\mathcal{A}_I)$ and $\pi(\mathcal{A}_J)$ are two distinct essential subsets of $\pi(\mathcal{A})$. The defect of $\pi(\mathcal{A}_I)$ is one less than the defect $\mathcal{A}_I$ when e_1 is contained in the affine span of $\mathcal{A}_I$ and equal to the defect of $\mathcal{A}_I$ otherwise. Since $\mathrm{MV}(\mathcal{A}) > 0$, we conclude that $\mathrm{def}(\mathcal{A}_I) = \mathrm{def}(\mathcal{A}_J) = 0$ and e_1 is in the affine span of both $\mathcal{A}_I$ and $\mathcal{A}_J$. This implies that

$$\mathrm{rk}(\mathcal{A}_{I\cup J}) \leq \mathrm{rk}(\mathcal{A}_I) + \mathrm{rk}(\mathcal{A}_J) - 1 = |I| + |J| - 1 < |I| + |J|.$$

But then $\mathcal{A}_{I\cup J}$ has negative defect, contradicting the hypothesis that $\mathrm{MV}(\mathcal{A}) > 0$. □

Lemma 15 shows that $\mathrm{Res}_{\pi(\mathcal{A})} \in \mathbb{Z}[v_{i,\beta}]$ is not 1. Writing the polynomials f_i of $\mathcal{F}(u)$ as polynomials supported on $\pi(\mathcal{A})$ gives

$$f_i = \sum_{\beta \in \pi(A_i)} h_{i,\beta}(x_1) x^\beta \in \mathbb{C}[x_1][x_2, \ldots, x_n].$$

Thus, evaluating $\mathrm{Res}_{\pi(\mathcal{A})}$ at the system $\mathcal{F}(u)$ amounts to substituting $h_{i,\beta}(x_1)$ for $v_{i,\beta}$, that is,

$$\mathrm{Res}_{\pi(\mathcal{A})}(\mathcal{G}(v))\big|_{v_{i,\beta}=h_{i,\beta}(x_1)} = \mathrm{Res}_{\pi(\mathcal{A})}(\mathcal{F}(u)) \in \mathbb{Z}[u_{i,\alpha}][x_1]. \tag{3}$$

This polynomial vanishes at all points a_1 such that the system $\mathcal{F}(a_1, x_2, \ldots, x_n)$ has a solution in $(\mathbb{C}^\times)^{n-1}$. The fact that a_1 may be zero is reflected in a power of x_1 appearing as a factor of this polynomial, as exhibited in the following lemma.

Lemma 16 ([4, Proposition 4.7 and Theorem 1.4])**.** Let $\mathcal{A}$ be square with $\mathrm{MV}(\mathcal{A}) > 0$. Then there exists $d \in \mathbb{Z}$ such that, up to a constant,

$$\mathrm{Res}_{\pi(\mathcal{A})}(\mathcal{F}(u)) = x_1^d \,\mathrm{Res}_{Q,A_1,\dots,A_n}(z - x_1, f_1, \dots, f_n)|_{z=x_1}.$$

Moreover, for generic $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$,

$$\mathrm{Res}_{\pi(\mathcal{A})}(\mathcal{F}) = x_1^d \prod_{s \in \mathcal{V}^\times(\mathcal{F})} (x_1 - s_1).$$

Lemma 17 ([14, Theorems 4 and 5])**.** As a polynomial in x_1, the degree of $\mathrm{Res}_{\pi(\mathcal{A})}(\mathcal{F}(u))$ is the mixed volume of the convex hulls of $A_i \cup (\{0\} \times \pi(A_i))$.

The convex hull of $A \cup (\{0\} \times \pi(A))$ is often called the *shadow* of the convex hull of A. We define $\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u))$ to be the polynomial $\mathrm{Res}_{Q,A_1,\dots,A_n}(z - x_1, f_1, \dots, f_n)|_{z=x_1}$ in $\mathbb{Z}[u][x_1]$.

Corollary 18. As a polynomial in x_1, the degree of $\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u))$ is $\mathrm{MV}(\mathcal{A})$.

Lemma 19. Let $\mathcal{A} = (A_1, \dots, A_n)$ be a square collection of supports with $\mathrm{MV}(\mathcal{A}) > 0$. Let $\mathcal{F}(u)$ be the universal polynomial system over $\mathcal{A}$. We write the resultant $\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u))$ as

$$\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u)) = q_{\mathrm{MV}(\mathcal{A})}(u) x_1^{\mathrm{MV}(\mathcal{A})} + q_{\mathrm{MV}(\mathcal{A})-1}(u) x_1^{\mathrm{MV}(\mathcal{A})-1} + \cdots + q_1(u) x_1 + q_0(u).$$

Then

$$\Sigma_1(\mathcal{V}^\times(\mathcal{F}(u))) = -\frac{q_{\mathrm{MV}(\mathcal{A})-1}(u)}{q_{\mathrm{MV}(\mathcal{A})}(u)}.$$

Proof. By Lemma 16 and Corollary 18, we have that $\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}) = \prod_{s \in \mathcal{V}^\times(\mathcal{F})}(x_1 - s_1)$ for generic $\mathcal{F} = \mathcal{F}(c) \in \mathbb{C}^{\mathcal{A}}$, up to a constant. Since $\mathcal{F}$ is generic, $q_{\mathrm{MV}(\mathcal{A})}(c) \neq 0$ and so the sum of the x_1 coordinates of points in $\mathcal{V}^\times(\mathcal{F})$ is

$$\Sigma_1(\mathcal{V}^\times(\mathcal{F})) = -\frac{q_{\mathrm{MV}(\mathcal{A})-1}(c)}{q_{\mathrm{MV}(\mathcal{A})}(c)}.$$

Since this equality holds for all generic $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$, the equality also holds for the universal polynomial system. □

Remark 20. D'Andrea and Jeronimo make the quotient in Lemma 19 explicit. For a square collection of supports $\mathcal{A}$ with positive mixed volume, a specialization of formula [5, Theorem 2.3] gives

$$\Sigma_1(\mathcal{V}^\times(\mathcal{F}(u))) = d_{Q,A_1,\dots,A_n} \frac{\frac{\partial \,\mathrm{Elim}_{Q,A_1,\dots,A_n}}{\partial u_{0,e_1}}(1, f_1, \dots, f_n)}{\mathrm{Elim}_{Q,A_1,\dots,A_n}(1, f_1, \dots, f_n)} = \frac{\frac{\partial \,\mathrm{Res}_{Q,A_1,\dots,A_n}}{\partial u_{0,e_1}}(1, f_1, \dots, f_n)}{\mathrm{Res}_{Q,A_1,\dots,A_n}(1, f_1, \dots, f_n)}$$

where $\mathcal{F}(u) = (f_0, \dots, f_n)$ is the universal polynomial system over $(Q, A_1, \dots, A_n)$.

To prove the main result of this section, we first state an elementary result on the coefficients in the composition of a monomial with a collection of polynomials. Let $[x^k]f(x)$ denote the coefficient of x^k in $f(x)$.

Lemma 21. Define $g_i(x) = \sum_{j=0}^{d_i} c_{i,j} x^j \in \mathbb{C}[c][x]$ and fix $\gamma = (\gamma_1, \dots, \gamma_n) \in \mathbb{Z}^n$. The polynomial

$$g^\gamma(x) := g_1(x)^{\gamma_1} \cdots g_n(x)^{\gamma_n}$$

has degree $d = d_1\gamma_1 + \cdots + d_n\gamma_n$, and if $\prod c_{i_k,j_k} x^{d-\delta}$ is a term of $g^\gamma(x)$, then $\sum d_{i_k} - j_k = \delta$.

Consequently, if $d_i - j > \delta$, then $c_{i,j}$ does not appear in $[x^{d-\delta}]g^\gamma(x)$ and each term $\prod c_{i_k,j_k} x^{d-\delta}$ can have at most one $c_{i,j}$ with the property that $d_i - j > \delta/2$.

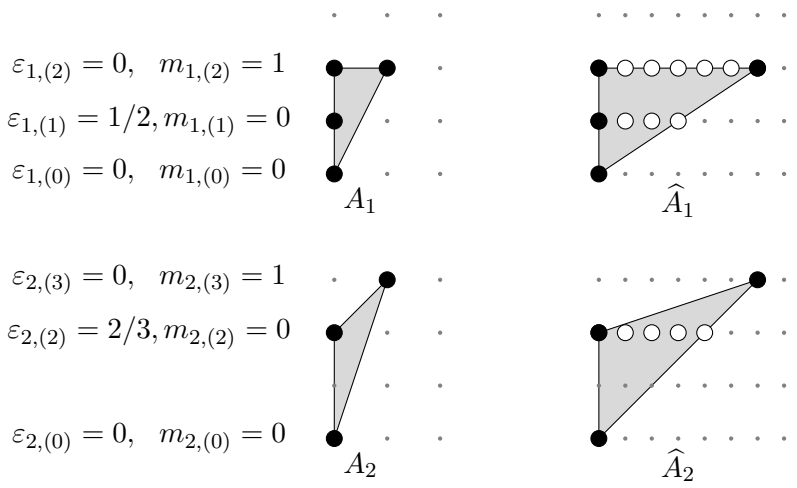

FIGURE 6. Support $\mathcal{A} = (A_1, A_2)$ and $\widehat{\mathcal{A}}$ along with a depiction of $\varepsilon_{i,\beta}$ and $m_{i,\beta}$ for $\mathcal{A}$.

Before stating the main result of this section, we establish some notation. Fix a square collection of supports $\mathcal{A}$ in $\mathbb{Z}^n$. For $\beta \in \pi(A_i)$, we let $(m_{i,\beta}, \beta)$ be the integer point in the Newton polytope of A_i with maximal first coordinate and $\varepsilon_{i,\beta} \in [0, 1)$ be the distance in the direction of e_1 from $(m_{i,\beta}, \beta)$ to the boundary of A_i. Note that there are finitely many $\varepsilon_{i,\beta}$, all of which are rational. We define λ to be the least common multiple of their denominators. Using $\odot$ to denote coordinate-wise multiplication, we define $\widehat{\mathcal{A}} := \mathbb{Z}^n \cap \operatorname{conv}(\lambda e_1 \odot \mathcal{A}) \cap \pi^{-1}(\pi(\mathcal{A}))$, so that $\lambda e_1 \odot \mathcal{A}$ is $\mathcal{A}$ scaled by λ in the direction of e_1. Note that $\pi(\mathcal{A}) = \pi(\widehat{\mathcal{A}})$ and all $\widehat{\varepsilon}_{i,\beta}$ for $\widehat{\mathcal{A}}$ are zero, see Figure 6.

Theorem 22. Let $\mathcal{A} = (A_1, \ldots, A_n)$ be a collection of supports in $\mathbb{Z}^n$ such that $\mathrm{MV}(\mathcal{A}) > 0$. The support of $q_{\mathrm{MV}(\mathcal{A})-\delta}(u)$ is contained in $\operatorname{offset}(\mathcal{A}, \delta)$. Moreover, the function $u \mapsto q_{\mathrm{MV}(\mathcal{A})-\delta}(u)$ is an affine linear function of the coefficients indexed by points in $\mathcal{A} \setminus \operatorname{offset}(\mathcal{A}, \delta/2)$.

Proof. We characterize which monomials in $\{u_{i,\alpha}\}_{\alpha \in A_i}$ may appear in the support of the coefficient $q_{\mathrm{MV}(\mathcal{A})-\delta}(u)$ of $x_1^{\mathrm{MV}(\mathcal{A})-\delta}$ in $\operatorname{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u)) \in \mathbb{Z}[u][x_1]$. Let $R(x_1) = \operatorname{Res}_{\pi(\mathcal{A})}(\mathcal{G}(h(x_1)))$ be the polynomial in $\mathbb{Z}[u][x_1]$ resulting from the substitution $v_{i,\beta} \mapsto h_{i,\beta}(x_1)$ in Equation (3). Let d be the degree of $R(x_1)$ in x_1.

Consider $\widehat{\mathcal{A}}$ as defined prior to the statement of this lemma. We let $\widehat{h_{i,\beta}}(x_1)$, and $\widehat{R}(x_1)$ be the analogues of $h_{i,\beta}(x_1)$ and $R(x_1)$ for $\widehat{\mathcal{A}}$, and $\widehat{d}$ the degree of $\widehat{R}(x_1)$. We note that $\operatorname{Res}_{\pi(\mathcal{A})}(\mathcal{G}(v)) = \operatorname{Res}_{\pi(\widehat{\mathcal{A}})}(\mathcal{G}(v))$. Since $\mathcal{S}(\widehat{\mathcal{A}}) = \lambda e_1 \odot \mathcal{S}(\mathcal{A})$ where $\mathcal{S}$ denotes the corresponding shadow polytope, we have that $\widehat{d} = \lambda d$.

Writing $h_{i,\beta}(x_1) = \sum_{j=0}^{m_{i,\beta}} c_{i,(j,\beta)} x_1^j$, we observe that $c_{i,(j,\beta)}$ appears in $[x_1^{d-\delta}]R(x_1)$ only if $\widehat{c}_{i,(\lambda j,\beta)}$ appears in $[x_1^{\lambda(d-\delta)}]\widehat{R}(x_1)$. Moreover, $[x_1^{\lambda(d-\delta)}]\widehat{R}(x_1)$ is the sum of $[x_1^{\lambda(d-\delta)}]\prod(\widehat{h}_{i,\beta}(x_1))^{\gamma_{i,\beta}}$ over all γ indexing monomials $v^\gamma = \prod v_{i,\beta}^{\gamma_{i,\beta}}$ in $\operatorname{Res}_{\pi(\mathcal{A})}(\mathcal{G}(v))$. Since $\deg_{x_1}(\widehat{h}_{i,\beta}(x_1)) = \lambda(m_{i,\beta} + \varepsilon_{i,\beta})$, Lemma 21 implies that $\widehat{c}_{i,(\lambda j,\beta)}$ appears in $[x_1^{\lambda(d-\delta)}]$ only if

$$\lambda(m_{i,\beta} + \varepsilon_{i,\beta}) - \lambda j \leq \widehat{d} - \lambda(d - \delta) = \lambda\delta.$$

Dividing through by λ gives the necessary condition that

$$m_{i,\beta} + \varepsilon_{i,\beta} - j \leq \delta.$$

We observe that $m_{i,\beta} + \varepsilon_{i,\beta} - j$ is the distance in the first coordinate from the point (j, β) to the boundary of the Newton polytope of A_i. Hence, $[x_1^{d-\delta}]R(x_1)$ may only involve coefficients $u_{i,\alpha}$ where the distance (in the direction of the first coordinate) from α to the boundary of the Newton

polytope of A_i is at most δ, that is, $\operatorname{offset}(\mathcal{A}, \delta)$. Since $R(x_1) = x_1^k \operatorname{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u))$, we have that $[x_1^{d-\delta}]R(x_1) = [x_1^{\mathrm{MV}(\mathcal{A})-\delta}] \operatorname{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F}(u)) = q_{\mathrm{MV}(\mathcal{A})-\delta}(u)$.

Moreover, each term of $[x_1^{d-\delta}]R(x_1)]$ can involve at most one coefficient $u_{i,\alpha}$ having this distance to the boundary greater than $\delta/2$. Hence, the function $u \mapsto [x_1^{d-\delta}]R(x_1)$ is an affine linear function of the coefficients indexed by $\mathcal{A} \setminus \operatorname{offset}(\mathcal{A}, \delta/2)$, completing the proof. □

4. Sparse monodromy

There has been recent progress in the study of monodromy groups of sparse polynomial systems [6, 7, 19]. Esterov showed that there are (essentially) two properties of $\mathcal{A}$ which can cause $G(\pi_{\mathcal{A}})$ to fail to be the full symmetric group. These properties explain the need for restrictions on the input to the sparse trace tests, as we detail in Section 4.1. In Section 4.2, we extend Esterov's result to restrictions of the branched cover $\pi_{\mathcal{A}}$, which completes the proof of correctness for the sparse trace tests.

4.1. Lacunary and triangular supports.

Throughout this section, we assume $\mathcal{A}$ is a square set of supports in $\mathbb{Z}^n$. The following example illustrates why the condition $L[\mathcal{A}] = \mathbb{Z}^n$ is necessary for the sparse trace tests.

Example 23. Let $\mathcal{F}(x_1, x_2)$ be any Bernstein-generic sparse polynomial system supported on $\mathcal{A}$, where $A_1 = \{(0,0), (1,0), (1,2)\}$ and $A_2 = \{(0,0), (1,0), (0,2), (1,2)\}$, see Figure 7. Since every power of x_2 is even, $\mathcal{F}$ may be written in the coordinates $(y_1, y_2) = (x_1, x_2^2)$ for some system $\mathcal{G}(y_1, y_2)$. Thus, the zeros of $\mathcal{F}$ in the torus have the following form:

$$\mathcal{V}^\times(\mathcal{F}) = \{(a_1, b_1), (a_1, -b_1), (a_2, b_2), (a_2, -b_2)\}.$$

The map $(x_1, x_2) \overset{\phi}{\mapsto} (x_1, x_2^2)$ is a two-to-one map from $\mathcal{V}^\times(\mathcal{F})$ to $\mathcal{V}^\times(\mathcal{G})$.

There are two types of proper nontrivial subsets S of $\mathcal{V}^\times(\mathcal{F})$ for which the sparse trace tests erroneously succeed. When S consists of a single fibre of ϕ, the trace $\Sigma_2(S)$ is zero. When S consists of a single point in each fibre of ϕ the trace $\Sigma_1(S)$ is half of $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$. In both cases, if $\Sigma_i(\mathcal{V}^\times(\mathcal{F}))$ moves affine linearly during the sparse trace test, then so does $\Sigma_i(S)$.

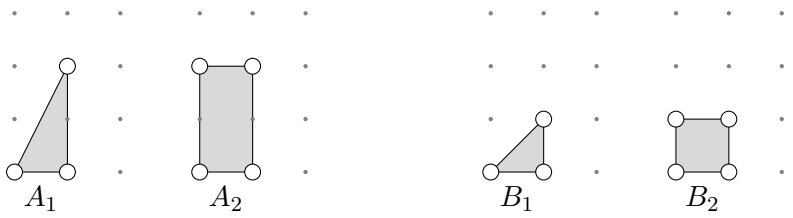

Figure 7. Lacunary support $\mathcal{A} = (A_1, A_2)$ and a lacunary reduction $\mathcal{B} = (B_1, B_2)$ of $\mathcal{A}$.

We now discuss the general case of the structure exhibited in Example 23.

Definition 24. We say a collection of supports $\mathcal{A}$ is *lacunary* if $[\mathbb{Z}^n : L[\mathcal{A}]] > 1$. Similarly, any $\mathcal{F}$ supported on lacunary support is called a *lacunary* system.

For any lacunary support $\mathcal{A}$ in $\mathbb{Z}^n$, there exists a $\mathbb{Z}$-linear map $\Phi : \mathbb{Z}^n \to \mathbb{Z}^n$ and nonlacunary support $\mathcal{B}$ such that $\Phi(\mathcal{B}) = \mathcal{A}$. We call such a $\mathcal{B}$ a *lacunary reduction* of $\mathcal{A}$. Additionally, the sparse polynomial system $\mathcal{G}$ in $\mathbb{C}^{\mathcal{B}}$ with the same coefficient vector as $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$ is called a *lacunary reduction* of $\mathcal{F}$ with respect to Φ. The lacunary reduction has the property that

$$\mathcal{G}(y_1, \ldots, y_n) = \mathcal{G}(\phi(x)) = \mathcal{F}(x)$$

where $\phi : (\mathbb{C}^\times)^n \to (\mathbb{C}^\times)^n$ is the monomial map $\phi(x_i) = x^{\Phi(e_i)} := x_1^{\Phi(e_i)_1} \cdots x_n^{\Phi(e_i)_n}$. The branched cover $\pi_{\mathcal{A}} : X_{\mathcal{A}} \to \mathbb{C}^{\mathcal{A}}$ decomposes as $\pi_{\mathcal{A}} = \pi_{\mathcal{B}} \circ (\phi \times \mathrm{id})$, where $\mathrm{id}(\mathcal{F}) = \mathcal{G} \in \mathbb{C}^{\mathcal{B}}$. The fibre of $\phi \times \mathrm{id}$ over a point $(y, \mathcal{G}) \in X_{\mathcal{B}}$ is

$$(\phi \times \mathrm{id})^{-1}(y, \mathcal{G}) = \{(x, \mathcal{F}) \ \mid \ \phi(x) = y, \ \mathcal{F} = \mathcal{G}\}.$$

This fibre has cardinality $\det(\Phi) = [\mathbb{Z}^n : L[\mathcal{A}]]$ and is identified with the solutions to the binomial system $\phi(x) = y$.

Theorem 25. Suppose $\mathcal{A}$ is lacunary.

- If $e_1 \notin L[\mathcal{A}]$, then $\Sigma_1(\mathcal{V}^\times(\mathcal{F})) = 0$.
- If $e_1 \in L[\mathcal{A}]$, then $\Sigma_1(\mathcal{V}^\times(\mathcal{F})) = [\mathbb{Z}^n : L[\mathcal{A}]]\Sigma_1(\mathcal{V}^\times(\mathcal{G}))$ where $\Phi(\mathcal{B}) = \mathcal{A}$ is a lacunary reduction of $\mathcal{A}$ satisfying $\Phi(e_1) = e_1$ and $\mathcal{G} \in \mathbb{C}^{\mathcal{B}}$ is the corresponding lacunary reduction of $\mathcal{F}$.

Proof. If $e_1 \notin L[\mathcal{A}]$, then there exists a linear map Φ and lacunary reduction $\mathcal{B}$ such that $\Phi(e_1) = ke_1$ for some $k > 1$. Hence, if $(s_1, \ldots, s_n)$ is a solution to some generic $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$, then so is $(\omega_k \cdot s_1, s_2, \ldots, s_n)$ for any k-th root of unity ω_k. Thus, the sum of the first coordinates of the solutions to $\mathcal{F}$ is zero.

If $e_1 \in L[\mathcal{A}]$, then Φ may be chosen so that $\Phi(e_1) = e_1$. Then the fibre over Φ of the point $(y, \mathcal{G})$ consists of $[\mathbb{Z}^n : L[\mathcal{A}]]$ points with identical first coordinate y_1, and the result follows. □

The following corollary highlights how the application of our sparse trace tests on (invalid) lacunary support is one-sided by identifying nonempty proper subsets $S \subsetneq \mathcal{V}^\times(\mathcal{F})$ on which these algorithms return `pass`.

Corollary 26. Suppose $\mathcal{A}$ is lacunary and $\mathcal{F}$ is supported on $\mathcal{A}$.

- If $e_1 \notin L[\mathcal{A}]$, then the trace Σ_1 of a union of fibres over $\phi \times \mathrm{id}$ is zero.
- If $e_1 \in L[\mathcal{A}]$, then the trace Σ_1 of a union of k points in each fibre is $\frac{k}{[\mathbb{Z}^n : L[\mathcal{A}]]}\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$.

We now illustrate a second property of $\mathcal{A}$ which prevents the use of our sparse trace tests. This property never occurs for abundant $\mathcal{A}$.

Example 27. Let $\mathcal{F}(x_1, x_2) = (f_1, f_2)$ be any Bernstein-generic sparse polynomial system supported on $\mathcal{A}$, where $A_1 = \{(0,0), (1,0), (2,0)\}$ and $A_2 = \{(0,0), (1,0), (0,1), (2,1), (1,1), (0,2)\}$, see Figure 8. Since the first coordinate of any point in $\mathcal{V}^\times(\mathcal{F})$ must be a solution to the square system $\mathcal{F}_1(x_1)$, the zeros of $\mathcal{V}^\times(\mathcal{F})$ in the torus have the following form:

$$\mathcal{V}^\times(\mathcal{F}) = \{(a_1, b_1), (a_1, c_1), (a_2, b_2), (a_2, c_2)\}.$$

The map $(x_1, x_2) \overset{\psi}{\mapsto} (x_1)$ is a two-to-one map from $\mathcal{V}^\times(\mathcal{F})$ to $\mathcal{V}^\times(\mathcal{F}_1)$.

If S consists of a single point in each fibre of ψ, then $\Sigma_1(S)$ is half of $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$. As in the lacunary example, the trace $\Sigma_1(S)$ moves affine linearly whenever $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ does.

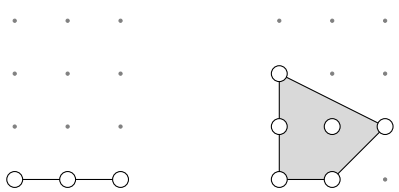

FIGURE 8. Triangular support $\mathcal{A} = (A_1, A_2)$.

Definition 28. We say a collection of supports $\mathcal{A}$ is *triangular* if there exists a proper subset $I \subsetneq [n]$ such that $\mathrm{rk}(\mathcal{A}_I) = |I|$. Similarly, any $\mathcal{F}$ supported on triangular support is called *triangular*.

We note that the condition that $\mathcal{B}$ is abundant in our sparse trace tests implies that $\mathcal{A}$ is not triangular. When $\mathcal{A}$ is triangular, witnessed by $I \subsetneq [n]$, the mixed volume of $\mathcal{A}_I$ is defined to be its mixed volume within its affine span. When $1 < \mathrm{MV}(\mathcal{A}_I) < \mathrm{MV}(\mathcal{A})$, we say that $\mathcal{A}$ is *strictly triangular*.

Suppose that $\mathcal{F}'$ is supported on the triangular support $\mathcal{A}'$, witnessed by the subsystem $\mathcal{A}'_I$. We consider the map $\Phi' : \mathbb{Z}^{|I|} \to \mathbb{Z}^n$ which sends $e_{j_1}, \dots, e_{j_{|I|}}$ to generators of the saturated lattice $\mathrm{span}(L[\mathcal{A}'_I]) \cap \mathbb{Z}^n$, thus identifying $L[\mathcal{A}'_I]$ with a sublattice of $\mathbb{Z}^{|I|}$. By choosing a complement of $L[\mathcal{A}'_I]$ in $\mathbb{Z}^n$, we extend the map Φ' to an invertible change of coordinates $\Phi : \mathbb{Z}^n \to \mathbb{Z}^n$ and define $\phi : (\mathbb{C}^\times)^n \to (\mathbb{C}^\times)^n$ to be the corresponding invertible monomial map.

We define $\mathcal{A} := \Phi^{-1}(\mathcal{A}')$ and have the following isomorphism of incidence varieties:

$$X_{\mathcal{A}'} \xrightarrow{\phi \times \Phi^{-1}} X_{\mathcal{A}}.$$

We note that the system $\mathcal{F}_I = \mathcal{F}'_I(\phi(x))$ is a polynomial system in the variables $x_{j_1}, \dots, x_{j_{|I|}}$.

When $\mathcal{A}$ is triangular, the map $\pi_{\mathcal{A}} : X_{\mathcal{A}} \to \mathbb{C}^{\mathcal{A}}$ decomposes as

$$X_{\mathcal{A}} \to X_{\mathcal{A}_I} \times \mathbb{C}^{\mathcal{A}_{I^c}} \to \mathbb{C}^{\mathcal{A}},$$

where I^c is the complement of I in $[n]$. The first map takes $(x, \mathcal{F}) \mapsto ((x_I, \mathcal{F}_I), \mathcal{F}_{I^c})$, and the second map takes $((x_I, \mathcal{F}_I), \mathcal{F}_{I^c}) \mapsto \mathcal{F}$.

Theorem 29. Suppose $\mathcal{A}$ is triangular, witnessed by $I \subsetneq [n]$. If $e_1 \in L[\mathcal{A}_I]$ then

$$\frac{\mathrm{MV}(\mathcal{A})}{\mathrm{MV}(\mathcal{A}_I)} \Sigma_1(\mathcal{V}^\times(\mathcal{F}_I)) = \Sigma_1(\mathcal{V}^\times(\mathcal{F})),$$

where the mixed volume of $\mathcal{A}_I$ is taken in $L[\mathcal{A}_I]$.

Proof. If $e_1 \in L[\mathcal{A}_I]$ then there exists an invertible $\mathbb{Z}$-linear map $\Phi : \mathbb{Z}^n \to \mathbb{Z}^n$ fixing e_1 such that $L[\Phi^{-1}(\mathcal{A}_I)] \subseteq \langle e_1, \dots, e_{|I|} \rangle$. Since this map fixes e_1, the corresponding monomial map fixes the first coordinates of points in $(\mathbb{C}^\times)^n$. Moreover, since Φ is invertible, it preserves mixed volumes. Hence, without loss of generality, we assume that $L[\mathcal{A}_I] \subseteq \langle e_1, \dots, e_{|I|} \rangle$. Any fibre of $X_{\mathcal{A}} \to X_{\mathcal{A}_I} \times \mathbb{C}^{\mathcal{A}_{I^c}}$ consists of $\frac{\mathrm{MV}(\mathcal{A})}{\mathrm{MV}(\mathcal{A}_I)}$-many points $\{(s_i, \mathcal{F})\}$ where each s_i has the same first coordinate. Therefore, the trace $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is $\frac{\mathrm{MV}(\mathcal{A})}{\mathrm{MV}(\mathcal{A}_I)}$ times the trace $\Sigma_1(\mathcal{V}^\times(\mathcal{F}_I))$. □

Corollary 30. Suppose $\mathcal{A}$ is triangular with $e_1 \in L[\mathcal{A}]$ and $\mathcal{F}$ is supported on $\mathcal{A}$. Then the trace of a union of k points in each fibre of $X_{\mathcal{A}} \to X_{\mathcal{A}_I} \times \mathbb{C}^{\mathcal{A}_{I^c}}$ is $\frac{k \cdot \mathrm{MV}(\mathcal{A}_I)}{\mathrm{MV}(\mathcal{A})} \Sigma_1(\mathcal{V}^\times(\mathcal{F}))$. Thus, if S is a union of k points in each fibre, then the trace is an affine linear function of the coefficients of $\mathcal{F}$, even if S is not complete.

Remark 31. Theorem 29 and Corollary 30 may be thought of as partial triangular analogues to Theorem 25 and Corollary 26 for lacunary supports. As depicted in Example 27, triangularity guarantees only one type of proper subset of $\mathcal{V}^\times(\mathcal{F})$ which causes the sparse trace tests to erroneously succeed (compare to Example 23). Understanding whether additional problematic subsets of $\mathcal{V}^\times(\mathcal{F})$ always exist in the triangular setting is left to further research.

Although the sparse trace tests may not be applied to lacunary or triangular supports, there are settings in which these properties are advantageous. In Section 5 we use these properties to more efficiently compute traces. In [2], the authors give a recursive algorithm for solving such polynomial systems. The only polynomial systems which need to be directly solved, in their method, are those which are neither lacunary nor triangular. Therefore, one may pair their work with our sparse trace tests to establish a method for verifying the completeness of solution sets to sparse systems, even when they are lacunary or triangular.

4.2. **The restricted monodromy problem.** The following result by Esterov shows that lacunary and triangular supports, are not only problematic for the sparse trace tests, but also restrict the monodromy group $G(\pi_{\mathcal{A}})$. The goal of this section is to extend that result to restrictions of $\pi_{\mathcal{A}}$.

Proposition 32. [6, Theorem 1.5] For a square set of supports $\mathcal{A}$, the monodromy group $G(\pi_{\mathcal{A}})$ is the full symmetric group if and only if either of the following are true:

- $\mathcal{A}$ is neither lacunary nor strictly triangular.
- $\mathrm{MV}(\mathcal{A}) = 2$ and $[\mathbb{Z}^n : L[\mathcal{A}]] = 2$

Remark 33. The monodromy group associated to the family $u_{1,0} + u_{1,2}x^2$ of sparse polynomial systems is the full symmetric group, despite the support $\{0, 2\}$ being lacunary. This example shows the necessity of the second condition above.

Throughout this section, we assume that $\mathcal{A} = (A_1, \ldots, A_n)$ is a collection of supports in $\mathbb{Z}^n$ with $\mathrm{MV}(\mathcal{A}) > 0$. We fix $\mathcal{B} \subseteq \mathcal{A}$, $\mathcal{C} = \mathcal{A}\backslash\mathcal{B}$, and set $N = \sum_{i=1}^{n} |B_i|$. We take $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$ to be generic and write $\mathcal{F} = \mathcal{F}_{\mathcal{B}} + \mathcal{F}_{\mathcal{C}} \in \mathbb{C}^{\mathcal{B}} \times \mathbb{C}^{\mathcal{C}}$. We are interested in the monodromy action induced by varying the coefficients indexed by the points in $\mathcal{B}$.

We consider restricted polynomial systems $\mathcal{F}_u := \mathcal{F}(u_{\mathcal{B}}, \mathcal{F}_{\mathcal{C}}) = \mathcal{F}_{\mathcal{B}}(u) + \mathcal{F}_{\mathcal{C}} = (f_{1,u}, \ldots, f_{n,u})$, where

$$f_{i,u}(x) = \underbrace{\sum_{\beta \in B_i} u_{i,\beta} x^{\beta}}_{g_{i,u}(x)} + \underbrace{\sum_{\gamma \in C_i} c_{i,\gamma} x^{\gamma}}_{h_i(x)},$$

and the corresponding restricted branched cover

$$\begin{array}{c} X_{\mathcal{B},\mathcal{F}_{\mathcal{C}}} = \{(x, b) \mid \mathcal{F}_b(x) = 0\} \\ \Big\downarrow {\scriptstyle \pi_{\mathcal{B},\mathcal{F}_{\mathcal{C}}}} \\ \mathbb{C}^{\mathcal{B}} \cong \mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}} \subseteq \mathbb{C}^{\mathcal{A}}. \end{array}$$

The linear space $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}}$ in $\mathbb{C}^{\mathcal{A}}$ is not generic, and so, care must be taken in computing the monodromy group. Following the approach of Harris [9], we establish conditions under which the monodromy group $G(\pi_{\mathcal{B},\mathcal{F}_{\mathcal{C}}})$ is the full symmetric group by showing that

(1) $G(\pi_{\mathcal{B},\mathcal{F}_{\mathcal{C}}})$ contains a simple transposition.
(2) $G(\pi_{\mathcal{B},\mathcal{F}_{\mathcal{C}}})$ is 2-transitive.

Note that when $\mathcal{B} = \mathcal{A}$, the branched cover $\pi_{\mathcal{A},\emptyset}$ is the same as $\pi_{\mathcal{A}}$, so Proposition 32 applies.

Theorem 34. If $\mathcal{A}$ is a square set of supports which is neither lacunary nor triangular and $N > 0$, then the monodromy group $G(\pi_{\mathcal{B},\mathcal{F}_{\mathcal{C}}})$ contains a simple transposition.

Proof. By Proposition 32, $\pi_{\mathcal{A}}$ contains a simple transposition whenever $\mathcal{A}$ is neither lacunary nor triangular. In this setting, the transposition is witnessed by any complex line in $\mathbb{C}^{\mathcal{A}}$ which crosses the $\mathcal{A}$-discriminant transversally. A monodromy loop, within this complex line, around its isolated point of intersection with the $\mathcal{A}$-discriminant induces the transposition. Thus, it is enough to show that there exists such a line in $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}}$ which also crosses the discriminant transversally.

There are only two ways that a generic line in $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}}$ fails to cross the discriminant transversally: the discriminant does not involve the coefficients indexed by $\mathcal{B}$ or its intersection with $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}}$ is singular everywhere.

Since $\mathcal{A}$ is neither lacunary nor triangular, Proposition 32 implies $G(\pi_{\mathcal{A}})$ contains a simple transposition and so the $\mathcal{A}$-discriminant is not singular everywhere. Since $\bigcup_{\mathcal{G}_{\mathcal{C}} \in \mathbb{C}^{\mathcal{C}}} \mathbb{C}^{\mathcal{B}} \times \mathcal{G}_{\mathcal{C}} = \mathbb{C}^{\mathcal{A}}$, the intersection of the discriminant with $\mathbb{C}^{\mathcal{B}} \times \mathcal{F}_{\mathcal{C}}$ for generic $\mathcal{F}_{\mathcal{C}}$ is not singular everywhere. Moreover, by [6, Lemma 1.20] and [6, Lemma 3.9], the defining equation of the discriminant has positive degree in c_α for each $\alpha \in \mathcal{A}$. In particular, this equation involves the coefficients indexed by $\mathcal{B}$. □

Finding conditions which imply that $G(\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}})$ is 2-transitive is more involved, and we follow the approach of [9]. We consider the following incidence correspondence of the fibre-square of $\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}}$,

$$\begin{array}{ccc} & Y = \{(x,y,b) \mid \mathcal{F}_b(x) = \mathcal{F}_b(y) = 0\} & \\ {\scriptstyle p}\swarrow & & \searrow{\scriptstyle \pi} \\ S & & \mathbb{C}^{\mathcal{B}} \end{array}$$

equipped with projections to the sets $S = ((\mathbb{C}^\times)^n)^2$ and $\mathbb{C}^{\mathcal{B}}$. To show $\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}}$ is 2-transitive, we rely on the following elementary result.

Proposition 35. The monodromy group $\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}}$ is 2-transitive if and only if Y has two components of top dimension.

The structure of the argument is as follows: When $\mathcal{B}$ is abundant, we stratify the image $p(Y)$ of Y into subvarieties of S. We compute the dimension of each stratum and show that the fibre dimension is constant over each. We show that there are only two preimages of strata which have top dimension and these come from irreducible strata. Since Y is the union of these preimages, we conclude that Y has at most two components of top dimension.

Suppose $\mathcal{B}$ is abundant and the support $\mathcal{A}$ has been shifted so $\mathbf{0} \in B_i$ for each $i \in [n]$. For all $I \subseteq [n]$, we define

$$V_I = \{(x,y) \in S \;\mid\; x^\alpha = y^\alpha \text{ for all } \alpha \in L[\mathcal{B}_I]\}$$
$$U_I = V_I \backslash \bigcup_{J \supsetneq I} V_J \qquad\qquad W_I = U_I \cap p(Y).$$

Since $\mathcal{B}$ is abundant, $\operatorname{rk}(L[\mathcal{B}_I]) = n$ for any $I \neq \emptyset$. Thus, after an invertible monomial change of coordinates, $L[\mathcal{B}_I] = k_1\mathbb{Z} \oplus \cdots \oplus k_n\mathbb{Z}$, where $k_1, \ldots, k_n$ are the *invariant factors* of $L[\mathcal{B}_I]$. With respect to this lattice, the variety V_I decomposes into components $V_I(\omega)$, one for each tuple $\omega = (\omega_1, \ldots, \omega_n)$ where ω_i is a k_i-th root of unity. Each component $V_I(\omega)$ is an n-dimensional irreducible variety defined via the explicit parametrization

$$\varphi_\omega : (\mathbb{C}^\times)^n \to S$$
$$(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_n, \omega_1 x_1, \ldots, \omega_n x_n).$$

Note that if $V_J(\omega) \cap V_I(\omega') \neq \emptyset$, then they are equal. Hence, the irreducible components of U_I are components of V_I. That is, when $I \neq \emptyset$, the components of U_I are of the form $V_I(\omega)$ for some (but not all) tuples of roots of unity ω. Hence, we refer to the components of U_I as $U_I(\omega)$. The only exception is $U_\emptyset = S \backslash \bigcup_{\emptyset \neq I} V_I$

Lemma 36. For any $s \in W_I$, the dimension of $p^{-1}(s)$ is $N - 2n + |I|$.

Proof. The fibre of p over $s = (x,y)$ is the solution set to the $2n$-many *affine* linear equations

$$\begin{aligned} g_{1,u}(x) &= -h_i(x), \quad g_{2,u}(x) = -h_i(x), \quad \ldots, \quad g_{n,u}(x) = -h_n(x) \\ g_{1,u}(y) &= -h_i(y), \quad g_{2,u}(y) = -h_i(y), \quad \ldots, \quad g_{n,u}(y) = -h_n(y) \end{aligned} \tag{4}$$

in the variables $\{u_{i,\beta} \mid \beta \in B_i\}$. Since $s \in W_I \subseteq p(Y)$, the fibre over s is nonempty, and the dimension of the fibre is given by the dimension of the kernel of the linear map $u \mapsto (g_{i,u}(x), g_{i,u}(y))_{i=1}^n$. Since each condition ($g_{i,u}(x) = 0$ or $g_{i,u}(y) = 0$) defining this kernel only involves the variables $\{u_{i,\beta}\}_{\beta \in B_i}$, the only linear dependencies amongst System (4) are between $g_{i,u}(x) = 0$ and $g_{i,u}(y) = 0$ for some i. Since the $u_{i,\beta}$ are indeterminants, this linear dependency occurs exactly when the vectors $\{x^\beta\}_{\beta \in B_i}$ and $\{y^\beta\}_{\beta \in B_i}$ are linearly dependent. Equivalently, since $\mathbf{0} \in B_i$ is a coordinate of each vector, they must be equal. Hence, the codimension of $p^{-1}(s)$ is $2n - |I|$, and its dimension is $N - 2n + |I|$. □

Corollary 37. The variety W_I is $U_I \cap \mathcal{V}(\{h_i(x) - h_i(y)\}_{i \in I})$. In particular, $W_\emptyset = U_\emptyset$ has dimension $2n$ and p is dominant.

Proof. A point $s = (x, y) \in U_I$ fails to be in W_I whenever $p^{-1}(s)$ is empty, or equivalently, System (4) is inconsistent. This occurs if and only if $g_{i,u}(x) = -h_i(x)$ and $g_{i,u}(y) = -h_i(y)$ describe parallel hyperplanes in the variables $\{u_{i,\beta}\}_{\beta \in B_i}$. This happens when $i \in I$ and $h_i(x) \neq h_i(y)$. □

We now introduce several definitions in preparation for establishing the dimensions of the $W_I(\omega)$ when $I \neq \emptyset$. We assume that an invertible monomial change of coordinates has been applied so that $L[\mathcal{B}_I] = k_1\mathbb{Z} \oplus \cdots \oplus k_n\mathbb{Z}$. Since we are only interested in computing the dimension, we may perform this change of coordinates individually for each I. For a tuple of roots of unity ω indexing a component of U_I, we define the lattice

$$\mathcal{L}_\omega = \{\alpha \in \mathbb{Z}^n \mid \omega^\alpha = 1\}$$

and the supports

$$\mathcal{C}^\omega = (C_I^\omega, \ldots, C_n^\omega) = (C_1 \backslash \mathcal{L}_\omega, \ldots, C_n \backslash \mathcal{L}_\omega)$$

We define $\delta_I(\omega)$ to be the number of supports in $\mathcal{C}_I^\omega$ which are nonempty, that is the number of $i \in I$ such that $C_i^\omega \neq \emptyset$.

Example 38. Let $B_1 = \{(0,0), (1,0), (0,2)\}$, $B_2 = \{(0,0), (2,0), (0,2)\}$, $C_1 = \{(0,1), (1,1), (2,0)\}$, $C_2 = \{(1,1)\}$, and $\mathcal{A} = \mathcal{B} \cup \mathcal{C}$ as depicted in Figure 9. We observe that $(0,0)$, $(1,1)$, and $(2,0)$ are all contained in $\mathcal{L}_{(-1,-1)}$ and so $\mathcal{C}^{(-1,-1)} = (\{(0,1)\}, \emptyset)$. Hence, $\delta_2(-1,-1) = 0$. On the other hand, we have $\mathcal{C}^{(1,-1)} = (\{(0,1), (1,1)\}, \{(1,1)\})$ and $\delta_{12}(1,-1) = 2$. The subset $\mathcal{C}_2^{(1,-1)}$ has negative defect.

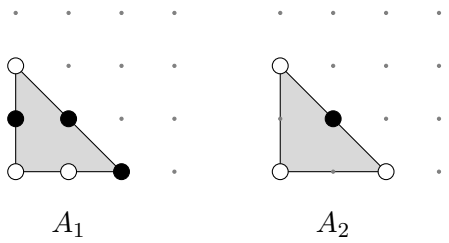

FIGURE 9. Support $\mathcal{A} = (A_1, A_2)$ where the points in $\mathcal{B}$ are white and the points in $\mathcal{C}$ are filled.

Lemma 39. For $\emptyset \neq I \subseteq [n]$ and component $U_I(\omega)$, the variety $W_I(\omega) = U_I(\omega) \cap p(Y)$ is empty if the defect of any subset of $\mathcal{C}_I^\omega$ is negative. Otherwise, $W_I(\omega)$ is the union of irreducible varieties of dimension $n - \delta_I(\omega)$.

Proof. By Corollary 37, $W_I(\omega)$ may be identified with the zero set

$$\mathcal{V}(\{h_i(x_1, \ldots, x_n) - h_i(\omega_1 x_1, \ldots, \omega_n x_n)\}_{i \in I}),$$

its preimage under the parametrization of $U_I(\omega)$. Explicitly, these polynomials are

$$\sum_{\gamma \in C_i} c_{i,\gamma} x^\gamma - \sum_{\gamma \in C_i} c_{i,\gamma} (\omega_1 x_1)^{\gamma_1} \cdots (\omega_n x_n)^{\gamma_n} \qquad \text{for } i \in I. \tag{5}$$

After collecting terms, this is the sparse polynomial system

$$\sum_{\gamma \in \mathcal{C}_I^\omega} c_{i,\gamma} (1 - \omega^\gamma) x^\gamma \qquad \text{for } i \in I. \tag{6}$$

System (5) is supported on C_i, however, any term indexed by $\gamma \in C_i \cap \mathcal{L}_\omega$ does not appear after collecting terms to obtain System (6) because $(1 - \omega^\gamma) = 0$. Hence, System (6) is supported on $\mathcal{C}_I^\omega$.

The key observation is that System (6) is generic with respect to $\mathcal{C}_I^\omega$: the nonzero constant $1-\omega^\gamma$ may be absorbed into the coefficient $c_{i,\gamma}$ by a reparametrization. By Lemma 14, the dimension of W_I is $n-\delta_I(\omega)$ when the defect of every subset of $\mathcal{C}_I^\omega$ is nonnegative and empty otherwise. □

Theorem 40. Suppose $\mathcal{A}$ is a square set of supports with $\mathrm{MV}(\mathcal{A})>0$. If $\mathcal{B}\subseteq\mathcal{A}$ is abundant, then the monodromy group $G(\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}})$ is 2-transitive if and only if $\mathcal{A}$ is not lacunary.

Proof. The backwards direction follows from Proposition 32: If $\mathcal{A}$ is lacunary, then $G(\pi_\mathcal{A})$ is not 2-transitive, so the restriction $G(\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}})$ is not either.

For the forward direction, we compute the dimension of the preimage $p^{-1}(W)$ of a component W of $W_I(\omega)$ for $I\neq\emptyset$. The dimension of W is $n-\delta_I(\omega)$ by Lemma 39. The dimension of the fibre over a point in W is $N-2n+|I|$ by Lemma 36. Hence, the dimension of $p^{-1}(W)$ is $N-n+|I|-\delta_I(\omega)$ which only obtains the value of N when $|I|=n$ and $\delta_I(\omega)=0$.

Suppose there exists $W_I(\omega)$ such that $I=[n]$, $\delta_I(\omega)=0$, and $\omega\neq(1,\dots,1)$. Since $\delta_I(\omega)=0$, each C_i is contained in the lattice $\mathcal{L}_\omega$, a proper sublattice of $\mathbb{Z}^n$ since $\omega\neq(1,\dots,1)$. Note that $\mathcal{L}_\omega$ also contains $\mathcal{B}$ since $L[\mathcal{B}]=k_1\mathbb{Z}\oplus\cdots\oplus k_n\mathbb{Z}$ and the ω_i are k_i-th roots of unity. Thus, $\mathcal{C}$ and $\mathcal{B}$ are both contained in $\mathcal{L}_\omega$ and so $\mathcal{A}=\mathcal{B}\cup\mathcal{C}$ must also be contained in $\mathcal{L}_\omega$. This is a contradiction since $\mathcal{A}$ is not lacunary. Hence, the only $W_{[n]}(\omega)$ containing an irreducible component W for which $\dim(p^{-1}(W))=N$ is the variety $W_{[n]}(1,\dots,1)=\{(x,y)\in S\mid x=y\}$. Its preimage is the diagonal of Y and is irreducible of dimension N.

The only other variety in the stratification of $p(Y)$ is $W_\emptyset$. This is an irreducible variety of dimension $2n$, and the dimension of a fibre of a point in $W_\emptyset$ is $N-2n$. Hence, $p^{-1}(W_\emptyset)$ is irreducible of dimension N. Thus, the only two components of Y of dimension N are $p^{-1}(W_\emptyset)$ and $p^{-1}(W_{[n]}(1,\dots,1))$. □

Putting Theorem 34 and Theorem 40 together, we arrive at our main result.

Theorem 41. Let $\mathcal{A}$ be a nonlacunary square set of supports with $\mathrm{MV}(\mathcal{A})>0$. Suppose $\mathcal{B}\subseteq\mathcal{A}$ is abundant. Then the monodromy group $G(\pi_{\mathcal{B},\mathcal{F}_\mathcal{C}})$ is the full symmetric group.

Finally, using the machinery developed in this section, we provide a simple proof that the first coordinates of the solution set $\mathcal{V}^\times(\mathcal{F})$ are distinct whenever $\mathcal{A}$ is neither lacunary nor triangular.

Lemma 42. If $\mathcal{A}$ is a square set of supports which is neither lacunary nor triangular and $\mathcal{F}\in\mathbb{C}^\mathcal{A}$ is generic, then the first coordinates of $\mathcal{V}^\times(\mathcal{F})$ are distinct.

Proof. We let $\mathcal{B}=\emptyset$ so that Y is the fibre-square of $\pi_\mathcal{A}$. Since $G(\pi_\mathcal{A})$ is 2-transitive by Proposition 32, Y has two components of top dimension N. Independent of the argument in this section requiring $\mathcal{C}$ to be abundant, these components are still $p^{-1}(W_\emptyset)$ and the diagonal $p^{-1}(W_{[n]}(1,\dots,1))$. Any other component of Y cannot surject onto $\mathbb{C}^\mathcal{A}$ via π for dimension reasons. Hence, if $s=(x,y,\mathcal{F})$ is a point in a generic fibre over π with $x_1=y_1$, we must have that $s\in p^{-1}(W_{[n]}(1,\dots,1))$. □

Example 43. Consider the supports $\mathcal{A}$ and $\mathcal{B}$ in Example 38. The inclusion lattice corresponding to the sets $\{V_I\}_{I\subseteq[2]}$ and $\{U_I\}_{I\subseteq[2]}$ are displayed in Figure 10. Hence, the set $S=((\mathbb{C}^\times)^2)^2$ is stratified by $U_{12}(1,1),U_{12}(1,-1),U_2(-1,-1),U_2(-1,1)$, and $U_\emptyset$. By Lemma 39, $W_{12}(1,-1)$ and $W_2(-1,1)$ are both empty since both $\mathcal{B}_{\{2\}}^{(1,-1)}$ and $\mathcal{B}_{\{1\}}^{(-1,1)}$ have negative defects. Hence, the only nonempty components of $p(Y)$ are $W_{12}(1,1)$,$W_2(-1,-1)$, and $W_\emptyset$. Since $\delta_{12}(1,1)=0$ and $\delta_2(-1,-1)=0$, we have that $W_I(\omega)=U_I(\omega)$. The dimension counts for $p^{-1}(W_{12}(1,1)),p^{-1}(W_2(-1,-1))$, and $p^{-1}(W_\emptyset)$ appear in Table 2. We remark that $p^{-1}(W_2(-1,-1))$ is actually a subvariety of $p^{-1}(W_\emptyset)$, and not its own component.

5. Examples and extensions

We collect a gallery of examples that illustrate and extend our theorems.

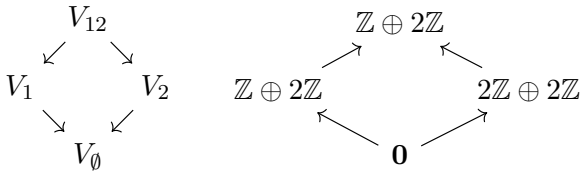

FIGURE 10. The inclusion poset of $\{V_I\}_{I\subseteq\{1,2\}}$ and the inclusion poset of the corresponding lattices $\{L[\mathcal{B}_I]\}_{I\subseteq\{1,2\}}$.

Stratum	$\dim(U_I(\omega))$	$\dim(W_I(\omega))$	Fibre dimension	$\dim(p^{-1}(W_I(\omega))$
$U_{12}(1,1)$	2	2	4	6
$U_{12}(1,-1)$	2	-1	-1	-1
$U_2(-1,1)$	2	-1	-1	-1
$U_2(-1,-1)$	2	2	3	5
$U_\emptyset$	4	4	2	6

TABLE 2. Summary of dimension counts for various strata of S.

5.1. **Computing the trace by reducing to the necessary support.** Fix $\mathcal{A}$ to be a square collection of supports with positive mixed volume, and $\mathcal{F}\in\mathbb{C}^{\mathcal{A}}$ to be Bernstein-generic. Then $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is equal to $\Sigma_1(\mathcal{V}^\times(\mathcal{F}_\mathcal{N}))$ since, by definition, this trace depends only on the coefficients of $\mathcal{F}$ in the necessary support $\mathcal{N}\subseteq\mathcal{A}$. The advantage of reducing to $\mathcal{F}_\mathcal{N}$ is that $\mathrm{MV}(\mathcal{N})$ is often significantly smaller than $\mathrm{MV}(\mathcal{A})$. Numerically, one may perform a direct computation of $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ by solving a system of much smaller degree. Algebraically, resultant computations on $\mathcal{F}_\mathcal{N}$ are likely to be much easier than resultant computations on $\mathcal{F}$.

Example 44. Consider a polynomial system $\mathcal{F}=(f_1,\dots,f_n)$ in n variables, where $\deg(f_i)=d_i$ for $i=1,\dots,n$. Then $\mathcal{F}$ is a Bernstein-generic sparse polynomial system supported on

$$\mathit{\Delta}=(d_1\Delta_n, d_2\Delta_n,\dots,d_n\Delta_n),$$

where $\mathit{\Delta}_n$ consists of the $n+1$ vertices of the standard n-dimensional simplex and $d\Delta_n$ consists of all lattice points in the d-th dilate of Δ_n.

A naive numerical computation of $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ involves computing $\prod_{i=1}^n d=\mathrm{MV}(\Delta)$-many solutions, by Bézout's theorem. However, by Lemma 8, this trace only depends on the terms of each f_i of degree d_i and d_i-1. The mixed volume of $\mathcal{N}$ is equal to $\prod_{i=1}^n d_i-\prod_{i=1}^n(d_i-1)$, considerably smaller than that of $\mathcal{A}$. A root of $\mathcal{V}(\mathcal{F})\subseteq\mathbb{C}^n$ at the origin of multiplicity $\prod_{i=1}^n(d_i-1)$ accounts for the difference.

For example, two generic bivariate polynomials $\{f_1,f_2\}$ of degree 5 have 25 common roots in the torus, whereas the system $\{g_1,g_2\}$ obtained by ignoring all terms of degree 3 or less has only $9=25-16$ common solutions in the torus (see Figure 11). The remaining 16 solutions are supported at the origin of $(\mathbb{C}^\times)^2$.

Example 45. Continuing Example 44, we compare timings of computing the traces of $\{f_1,f_2\}$ and $\{g_1,g_2\}$ by sampling random integers uniformly between -10 and 10 for coefficients. Since the solutions to $\{g_1,g_2\}$ which are not in the torus are at the origin, they do not influence the trace, even though they affect the resultant. We use the implementation in the Resultants package [20] in the Macaulay2 computer algebra system [8]. The experiments were carried out on Clemson's Palmetto server on an Intel Xeon E5-2680 v3 CPU at 2.50GHz with 126 GB of ram and running CentOS linux.

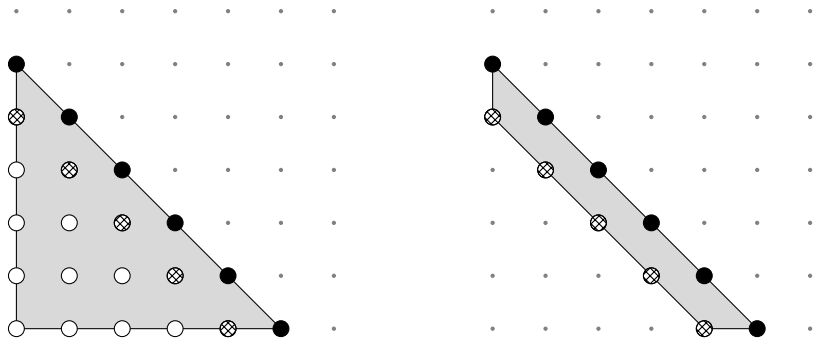

(a) (b)

FIGURE 11. (a) Support $5\Delta_2$. (b) Support $5\Delta_2 \backslash 3\Delta_2$.

Degree	5	10	15	20	25
$\{f_1, f_2\}$	0.0301	0.931	13.7	98.5	494
$\{g_1, g_2\}$	0.0203	0.201	1.33	5.15	17.4

TABLE 3. Timings (in seconds) for computing the trace of two polynomial systems using the hidden variable resultant. One is a generic system $\{f_1, f_2\}$ of bivariate quintics. The other, $\{g_1, g_2\}$, is obtained from $\{f_1, f_2\}$ by ignoring terms of degree less than four.

Rather than restricting the total degrees of each polynomial in a sparse polynomial system (as in Example 44), one can restrict the multidegrees of each polynomial. We illustrate our approach on such systems in the bivariate setting.

Example 46. Consider a Bernstein-generic bivariate system $\{f_1, f_2\}$ supported on

$$\square = (\square_{k_1,\ell_1}, \square_{k_2,\ell_2})$$

where $\square_{k,\ell} = ([0,k] \times [0,\ell]) \cap \mathbb{Z}^2$ (see Figure 12). The mixed volume of $\square$ is $k_1\ell_2 + k_2\ell_1$. The restricted polynomial system $\{g_1, g_2\}$, obtained by setting the coefficients of unnecessary monomials equal to zero, has support $(\square_{1,\ell_1}, \square_{1,\ell_2})$ after translation. This support has mixed volume $\ell_1 + \ell_2$. The translation of the supports corresponds to dividing each g_i by x_1^{k-1} to eliminate a component of $\mathcal{V}(g_1, g_2) \subseteq \mathbb{C}^2$ of multiplicity $(k-1)\ell_2 + (k_2-1)\ell_1$ supported on $x_1 = 0$. Note that, generically, $\Sigma_2(\mathcal{V}^\times(f_1, f_2)) \neq \Sigma_2(\mathcal{V}^\times(g_1, g_2))$.

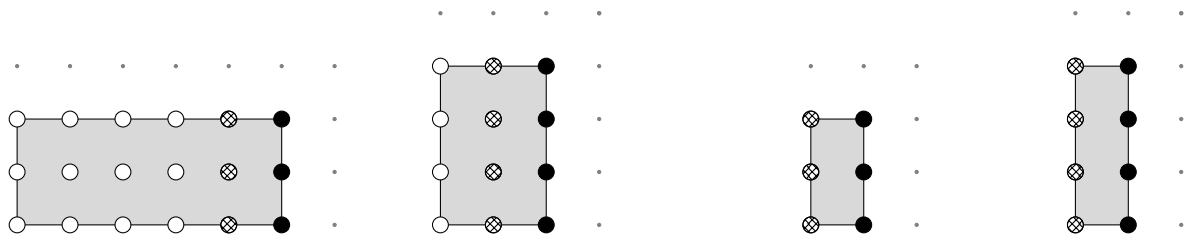

(a) (b) (c) (d)

FIGURE 12. Supports of f_1, f_2, g_1, and g_2 as in Example 46.

Even though the sparse trace tests cannot be directly applied to lacunary or triangular supports, one may use this special structure to compute traces more quickly (see Theorems 25 and 29).

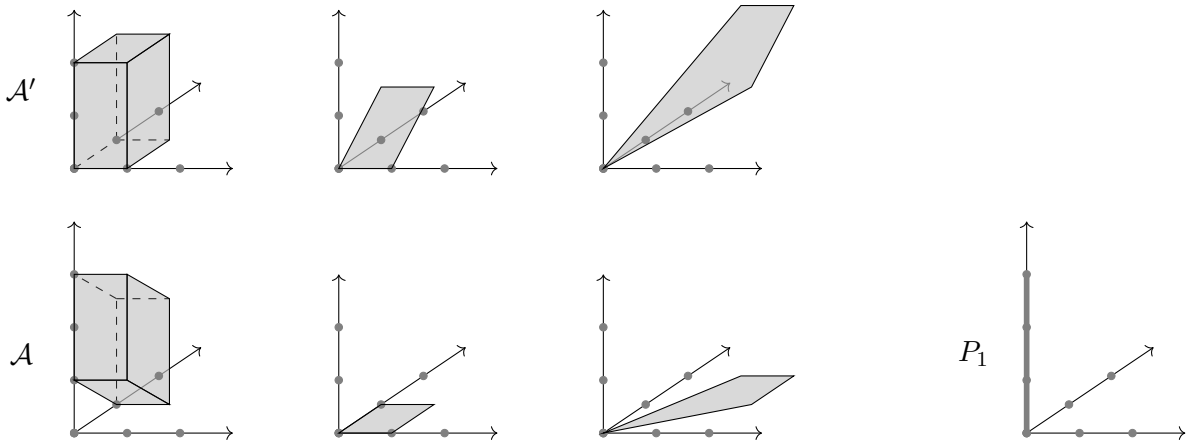

FIGURE 13. The Newton polytopes of $\mathcal{F}'$ in the variables (x_1, x_2, x_3) along with the Newton polytopes of $\mathcal{F}$ in the variables $(x_1, x_2, x_2x_3) = (y_1, y_2, y_3)$ (after multiplying f_1 by y_3 to clear denominators). After solving the subsystem $\{f_2, f_3\}$ in (y_1, y_2), back-substitution into f_1 involves solving a univariate cubic in y_3 supported on P_1.

Example 47. Consider a Bernstein-generic sparse polynomial system $\mathcal{F}' = (f_1', f_2', f_3')$ supported on $\mathcal{A}' = (A_1', A_2', A_3')$ consisting of all of the lattice points contained in the Newton polytopes displayed in Figure 13. The system $\mathcal{F}'$ is triangular, witnessed by the subset $\{2, 3\}$. Let $\mathcal{A}$ be the preimage of $\mathcal{A}'$ under the linear map $\Phi(\alpha_1, \alpha_2, \alpha_3) = (\alpha_1, \alpha_2 + \alpha_3, \alpha_3)$. If $\mathcal{F} = (f_1, f_2, f_3)$ is the corresponding polynomial system supported on $\mathcal{A}$, then $\mathcal{F}(y_1, y_2, y_3) = \mathcal{F}(x_1, x_2x_3, x_3) = \mathcal{F}'(x_1, x_2, x_3)$. The system $\mathcal{F}$ is obviously triangular since $\{f_2, f_3\}$ is a square system in y_1 and y_2.

Since $e_1 \in L[A_{\{2,3\}}]$ and Φ was chosen to fix e_1, the proof of Theorem 29 implies that $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is a multiple of $\Sigma_1\left(\mathcal{V}^\times\left(\mathcal{F}_{\{2,3\}}\right)\right)$. Namely, $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is three times the trace of the subsystem since each solution $(s_1, s_2) \in \mathcal{V}^\times\left(\mathcal{F}_{\{2,3\}}\right)$ extends to three solutions to $\mathcal{V}^\times(\mathcal{F})$ ($f_1(s_1, s_2, y_3)$ is a univariate polynomial of degree three). Geometrically, this is illustrated in Figure 13: as the projection P_1 of A_1 onto $L\left[\mathcal{A}_{\{2,3\}}\right]^\perp$ is a line segment of length three. Therefore, the trace of $\mathcal{V}^\times(\mathcal{F}')$ can be calculated directly from the trace of $\mathcal{V}^\times(f_2, f_3)$ by computing the length of the projection of A_1' onto a complement of $L\left[\mathcal{A}'_{\{2,3\}}\right]$.

5.2. **Applications to nongeneric systems.** Suppose $\mathcal{F} \in \mathbb{C}^{\mathcal{A}}$ is not a Bernstein-generic system, but has $d < \mathrm{MV}(\mathcal{A})$ isolated solutions in the torus. The trace $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is still a ratio of coefficients of $\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F})$. Since each $q_k(\mathcal{F}) = [x_1^k]\,\mathrm{Res}^{(1)}_{Q,\mathcal{A}}(\mathcal{F})$ is an affine linear function of the set of coefficients given in Theorem 22, this trace may be computed just as in the previous section. However, the set of unnecessary coefficients may be smaller in these Bernstein-deficient cases.

Example 48. Suppose that $\mathcal{F} = (f_1, f_2)$ consists of a pair of dense bivariate polynomials of degree 5 which have 24 solutions in $\mathbb{C}^2$, all of which are in the torus. Consequently, $\mathcal{F}$ is not Bernstein-generic: there is one solution at infinity. Using the notation in Theorem 22, the trace $\Sigma_1(\mathcal{V}^\times(\mathcal{F}))$ is $\frac{-q_{23}(\mathcal{F})}{q_{24}(\mathcal{F})}$. By Theorem 22, this is an affine linear function of the coefficients in $\mathcal{A}\setminus\mathrm{offset}(\mathcal{A}, 1)$, expressed in Figure 14 by the white and cross-hatched points in $\mathbb{Z}^2$.

As in Example 44, setting the unnecessary coefficients equal to zero produces a sparse polynomial system $\mathcal{G}$ with $\Sigma_1(\mathcal{V}^\times(\mathcal{F})) = \Sigma_1(\mathcal{V}^\times(\mathcal{G}))$, even though $|\mathcal{V}^\times(\mathcal{G})| = 15 < 24 = |\mathcal{V}^\times(\mathcal{F})|$. Each polynomial in $\mathcal{G}$ is supported on the second set of monomials in Figure 14.

5.3. **Computing other traces.** Rather than computing the trace $\sum_{x\in\mathcal{V}^\times(\mathcal{F})} x_1 = \Sigma_1(\mathcal{V}^\times(\mathcal{F}))$, one may be interested in computing the sum of a different monomial x^γ over the solutions $\mathcal{V}^\times(\mathcal{F})$. Such a function is also called a *trace* [5]. The trace of x^γ may be computed by applying an invertible

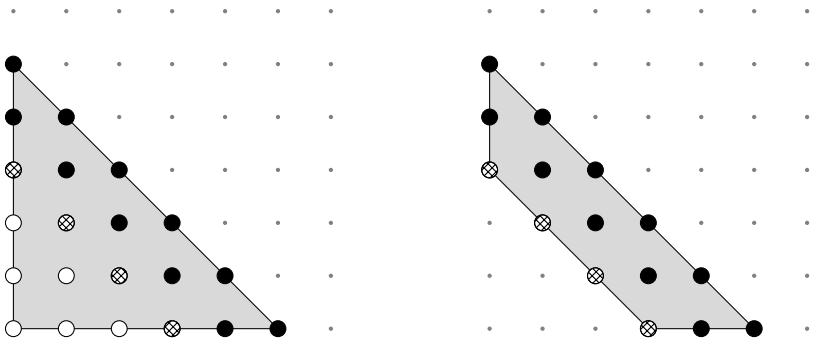

(a) (b)

FIGURE 14. (a) Support of f_1 and f_2 in Example 48. These polynomials are two bivariate quintics with one common solution at infinity and 24 common solutions in $(\mathbb{C}^\times)^2$. (b) Support of the polynomials f_1 and f_2 after truncating the terms of degree less than 3.

monomial change of coordinates $(\mathbb{C}_x^\times)^2 \to (\mathbb{C}_y^\times)^2$ which identifies x^γ with y_1. In the coordinates y, our results apply to the resulting polynomial system $\mathcal{G}(y)$. As the support $\mathcal{A}$ changes under this monomial change of coordinates, so do the sets $\text{offset}(\mathcal{A}, \delta)$. Consequently, we can find subsets of coefficients for which the function $\mathcal{F} \mapsto \sum_{x \in \mathcal{V}^\times(\mathcal{F})} x^\gamma$ is affine linear.

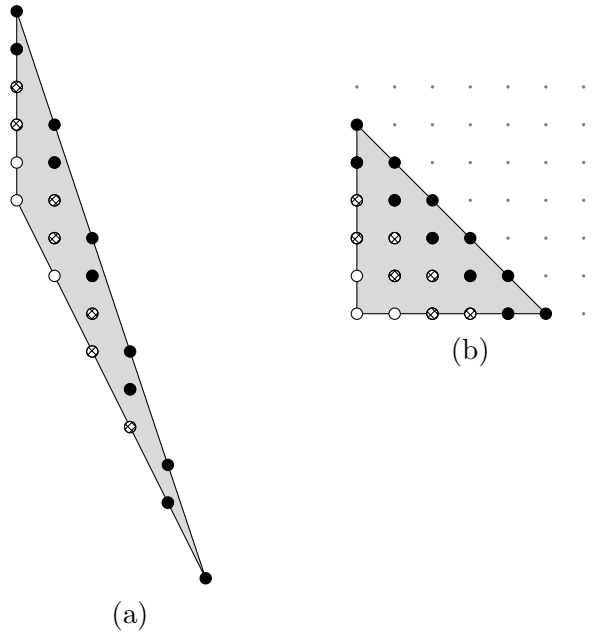

(a) (b)

FIGURE 15. (a) Support for the pair of dense polynomials of degree 5 in Example 49 after a change of coordinates. (b) The support for this pair of dense polynomials of degree 5 shifted back to the original coordinates.

Example 49. Consider the trace of $x_1x_2^2$ over the solutions $\mathcal{V}^\times(\mathcal{F})$ of a Bernstein-generic polynomial system $\mathcal{F} = (f_1, f_2)$ supported on $\mathcal{A} = (5\Delta_2, 5\Delta_2)$. This trace may be computed by applying the monomial substitution $\mathcal{F}(x_1, x_2) \mapsto \mathcal{F}(y_1y_2^{-2}, y_2) = \mathcal{G}(y_1, y_2)$, which induces a map $\mathcal{V}^\times(\mathcal{F}) \to \mathcal{V}^\times(\mathcal{G})$ on varieties which identifies the trace of $x_1x_2^2$ over $\mathcal{V}^\times(\mathcal{F})$ with the trace of y_1

over $\mathcal{V}^\times(\mathcal{G})$. The support of $\mathcal{G}$ is displayed in Figure 15, where our results imply that the white and cross-hatched coefficients influence $\Sigma_1(\mathcal{V}^\times(\mathcal{G}))$ affine linearly. Consequently, the same coefficients influence the trace of $x_1x_2^2$ over $\mathcal{V}^\times(\mathcal{G})$ affine linearly. In particular, the white points in Figure 15 correspond to coefficients which do not influence $\sum_{x\in\mathcal{V}^\times(\mathcal{F})} x_1x_2^2$.

Acknowledgements

We would like to thank Carlos D'Andrea, Gabriella Jeronimo, and Alexander Esterov for helpful discussions. The work began while the authors were in residence at ICERM's semester on Nonlinear Algebra, and continued while the first author was at Texas A&M University and the Max Planck Institute for Mathematics in the Sciences in Leipzig, Germany.

References

[1] D. N. Bernstein. The number of roots of a system of equations. *Functional Analysis and its Applications*, 9:183–185, 1975.

[2] T. Brysiewicz, J. I. Rodriguez, F. Sottile, and T. Yahl. Solving decomposable sparse systems. *Numerical Algorithms*, 88:453–474, 2021.

[3] D. Cox, J. Little, and H. Schenck. *Toric varieties.* Providence, RI: American Mathematical Society (AMS), 2011.

[4] C. D'Andrea and M. Sombra. A Poisson formula for the sparse resultant. *Proceedings of the London Mathematical Society*, 110(4):932–964, 02 2015.

[5] Carlos D'Andrea and Gabriela Jeronimo. Rational formulas for traces in zero-dimensional algebras. *Applicable Algebra in Engineering, Communication and Computing*, 19:495–508, 2008.

[6] A. Esterov. Galois theory for general systems of polynomial equations. *Compositio Mathematica*, 155(2):229–245, 2019.

[7] A. Esterov and L. Lang. Sparse polynomial equations and other enumerative problems whose Galois groups are wreath products. *Selecta Mathematica*, 28(22), 2022.

[8] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at `http://www.math.uiuc.edu/Macaulay2/`.

[9] J. Harris. Galois groups of enumerative problems. *Duke Mathematical Journal*, 46(4):685–724, 1979.

[10] J. D. Hauenstein and J. I. Rodriguez. Multiprojective witness sets and a trace test. *Advances in Geometry*, 20(3):297–318, 2020.

[11] A.G. Kushnirenko. Newton polyhedra and Bézout's theorem. *Akademija Nauk SSSR. Funkcionalc'nyĭ Analiz i ego Priloenija*, 10(3):82–83, 1976.

[12] A. Leykin, J. I. Rodriguez, and F. Sottile. Trace test. *Arnold Mathematical Journal*, 4(1):113–125, 2018.

[13] M. Minimair. Sparse resultant under vanishing coefficients. *Journal of Algebraic Combinatorics*, 18:53–73, 2003.

[14] J. Maurice Rojas. A convex geometric approach to counting the roots of a polynomial system. *Theoretical Computer Science*, 133(1):105–140, 1994.

[15] A. J. Sommese, J. Verschelde, and C. W. Wampler. Symmetric functions applied to decomposing solution sets of polynomial systems. *SIAM Journal on Numerical Analysis*, 40(6):2026–2046, 2002.

[16] A. J. Sommese, J. Verschelde, and Charles W. Wampler. Introduction to numerical algebraic geometry. In *Solving polynomial equations*, volume 14 of *Algorithms Comput. Math.*, pages 301–335. Springer, Berlin, 2005.

[17] A. J. Sommese and C. W. Wampler. Numerical algebraic geometry. In *The mathematics of numerical analysis (Park City, UT, 1995)*, volume 32 of *Lectures in Applied Mathematics*, pages 749–763. Amer. Math. Soc., Providence, RI, 1996.
[18] F. Sottile. General witness sets for numerical algebraic geometry. In *Proceedings of the 45th International Symposium on Symbolic and Algebraic Computation*, ISSAC '20, pages 418–425, New York, NY, USA, 2020. Association for Computing Machinery.
[19] F. Sottile and T. Yahl. Galois groups in enumerative geometry and applications, 2021. arXiv:2108.07905.
[20] Giovanni Staglianò. SparseResultants: computations with sparse resultants. Version 1.1. A *Macaulay2* package available at `https://github.com/Macaulay2/M2/tree/master/M2/Macaulay2/packages`.
[21] B. Sturmfels. On the Newton polytope of the resultant. *Journal of Algebraic Combinatorics*, 3:207–236, 1994.
[22] O. Zariski. A theorem on the poincaré group of an algebraic hypersurface. *Annals of Mathematics*, 38(1):131–141, 1937.

University of Notre Dame, Department of Applied and Computational Mathematics and Statistics, Notre Dame, IN 46556
Email address: `tbrysiew@nd.edu`

Clemson University, School of Mathematical and Statistical Sciences, 220 Parkway Drive, Clemson, SC 29634
Email address: `burr2@clemson.edu`