THEORY OF COMPUTING, Volume 20 (1), 2024, pp. 1–70
www.theoryofcomputing.org

Polynomial Identity Testing via Evaluation of Rational Functions

Ivan Hu Dieter van Melkebeek Andrew Morgan

Received November 2, 2022; Revised April 7, 2024; Published July 18, 2024

Abstract. We introduce a hitting set generator for Polynomial Identity Testing based on evaluations of low-degree univariate rational functions at abscissas associated with the variables. We establish an equivalence up to rescaling with a generator introduced by Shpilka and Volkovich, which has a similar structure but uses multivariate polynomials.

We initiate a systematic analytic study of the power of hitting set generators by characterizing their vanishing ideals, i. e., the sets of polynomials that they fail to hit. We provide two such characterizations for our generator. First, we develop a small collection of polynomials that jointly produce the vanishing ideal. As corollaries, we obtain tight bounds on the minimum degree, sparseness, and partition class size of set-multilinearity in the vanishing ideal. Second, inspired by a connection to alternating algebra, we develop a structured deterministic membership test for the multilinear part of the vanishing ideal. We present a derivation based on alternating algebra along with the required background, as well as one in terms of zero substitutions and partial derivatives, avoiding the need for alternating algebra.

A conference version of this paper appeared in the Proceedings of the 13th Innovations in Theoretical Computer Science Conference [40].

ACM Classification: Theory of computation → Algebraic complexity theory,Theory of computation → Pseudorandomness and derandomization

AMS Classification: 68Q17, 68Q87, 68Q15

Key words and phrases: polynomial identity testing, derandomization, pseudorandomness, lower bounds, vanishing ideal, Gröbner basis

 DOI: 10.4086/toc.2024.v020a001

As evidence of the utility of our analytic approach, we rederive known derandomization results based on the generator by Shpilka and Volkovich and present a new application in derandomization / lower bounds for read-once oblivious algebraic branching programs.

1 Overview

Polynomial identity testing (PIT) is the fundamental problem of deciding whether a given multivariate algebraic circuit formally computes the zero polynomial. PIT has a simple, efficient randomized algorithm that only needs blackbox access to the circuit: Pick a random point and check whether the circuit evaluates to zero on that particular point.

Despite the fundamental nature of PIT and the simplicity of the randomized algorithm, no efficient deterministic algorithm is known—even in the white-box setting, where the algorithm has access to the description of the circuit. The existence of such an algorithm would imply long-sought circuit lower bounds [27, 2, 31]. Conversely, sufficiently strong circuit lower bounds yield blackbox derandomization for all of BPP, the class of decision problems admitting efficient randomized algorithms with bounded error [43, 28]. Although the known results leave gaps between the two directions, they show that PIT constitutes an important stepping stone towards derandomizing BPP, and suggest that derandomizing BPP can be achieved in a blackbox fashion if at all.

Blackbox derandomization of PIT for a class of polynomials C in the variables $x_1, \ldots, x_n$ is equivalent to the efficient construction of a substitution G that replaces each x_i by a low-degree polynomial in a small set of fresh variables such that, for every nonzero polynomial p from C, $p(G)$ remains nonzero [49, Lemma 4.1]. We refer to G as a generator, the fresh variables are its seed, and we say that G hits the class C. If there are l seed variables, and if p and G have degree at most $n^{O(1)}$, then the resulting deterministic PIT algorithm for C makes $n^{O(l)}$ blackbox queries.

Much progress on derandomizing PIT has been obtained by designing such substitutions and analyzing their hitting properties for interesting classes C. Shpilka and Volkovich [48] introduced a generator, now dubbed the Shpilka–Volkovich generator, or "SV generator" for short. It takes as an additional parameter a positive integer l and can be viewed as an algebraic version of l-wise independence in the sense that any selection of l of the original variables can remain independent while the others are forced to zero. The property is realized using Lagrange interpolation with respect to n distinct elements of the underlying field $\mathbb{F}$, one element a_i corresponding to each original variable x_i. We refer to the elements a_i as *abscissas*; they are also parameters of SV.

Definition 1.1 (SV generator)**.** The *Shpilka–Volkovich (SV) Generator* for $\mathbb{F}[x_1, \ldots, x_n]$ is parametrized by the following data:

- A positive integer l.
- For each $i \in [n]$, a distinct abscissa $a_i \in \mathbb{F}$.

The generator SV^l takes as seed l pairs of fresh variables $(y_1, z_1), \ldots, (y_l, z_l)$ and substitutes

$$x_i \leftarrow \sum_{t=1}^{l} z_t \cdot L_i(y_t), \tag{1.1}$$

where the *Lagrange interpolant* L_i is the unique univariate polynomial of degree at most $n-1$ satisfying $L_i(a_i) = 1$ and $L_i(a_j) = 0$ for $j \in [n] \setminus \{i\}$.

SV^1 takes two seed variables, y and z. For any $i \in [n]$, setting $y = a_i$ gets $x_i = z$ while the other variables are set to zero. For larger l, SV^l is the sum of l independent copies of SV^1.

Shpilka and Volkovich proved that SV^1 hits sums of a bounded number of read-once formulas for $l = O(\log n)$ [48], later improved to $l = O(1)$ [41]. The generator for $l = O(\log n)$ has also been shown to hit multilinear depth-4 circuits with bounded top fan-in [32], multilinear bounded-read formulas [7], commutative read-once oblivious algebraic branching programs [15], $\Sigma\mathrm{m} \bigwedge \Sigma\Pi^{O(1)}$ formulas (i.e., sums of terms that are the product of a monomial and a power of a bounded-degree polynomial) [14], circuits with locally-low algebraic rank in the sense of [37], and orbits of simple polynomial classes under invertible linear transformations of the variables [39]. The generator is an ingredient in other hitting set constructions, as well, notably constructions using the technique of low-support rank concentration [4, 3, 26, 25, 46, 10]. It also forms the core of a "succinct" generator that hits a variety of classes, including depth-2 circuits [19].

Vanishing ideal. In this paper, we initiate a systematic study of the power of a generator G through the set of polynomials p such that $p(G)$ vanishes, which we denote by $\mathrm{Van}[G]$. For any fixed generator G, $\mathrm{Van}[G]$ is closed under addition, and for all $q \in \mathbb{F}[x_1, \ldots, x_n]$ and $p \in \mathrm{Van}[G]$, $q \cdot p \in \mathrm{Van}[G]$. By definition, this means that the set $\mathrm{Van}[G]$ has the algebraic structure of an ideal. From now on, we refer to $\mathrm{Van}[G]$ as the *vanishing ideal* of G. Our technical contributions can be understood as precisely characterizing the vanishing ideal of the SV generator.

Characterizations of the vanishing ideal facilitate two objectives:

Derandomization. A generator G hits a class $\mathcal{C}$ of polynomials if and only if $\mathcal{C}$ and $\mathrm{Van}[G]$ have at most the zero polynomial in common. For a class $\mathcal{C}$ defined by a resource bound, G trivially hits $\mathcal{C}$ if the characterization of the nonzero elements in $\mathrm{Van}[G]$ is incompatible with being computable within the resource bound. In other words, derandomization of PIT for $\mathcal{C}$ reduces to proving lower bounds for $\mathrm{Van}[G]$. By developing explicit structure for polynomials in the ideal, lower bounds become more tractable.

More generally, given a characterization of $\mathrm{Van}[G]$, in order to derandomize PIT for a class $\mathcal{C}$ it suffices to design another generator G' that hits merely the polynomials in $\mathcal{C} \cap \mathrm{Van}[G]$. As G hits the remainder of $\mathcal{C}$, combining G with G' yields a generator for all of $\mathcal{C}$. In this way, one may assume—for free—additional structure about the polynomials in $\mathcal{C}$, namely that the polynomials moreover belong to $\mathrm{Van}[G]$.

Lower bounds. If we happen to know that G hits the class $\mathcal{C}$ of polynomials computable within some resource bound, then any expression for a nonzero polynomial in $\mathrm{Van}[G]$

yields an explicit polynomial that falls outside C. Such a statement is often referred to as hardness of representation, and it can be viewed as a lower bound in the model of computation underlying C (assuming the polynomial can be computed in the model at all). Characterizing $\mathrm{Van}[G]$ makes explicit the polynomials to which the lower bound applies.

We illustrate how to make progress on both objectives through our characterizations of the SV generator's vanishing ideal.

Rational function evaluations. Another contribution of our paper is the development of an alternate view of the SV generator, namely as evaluations of univariate rational functions of low degree. We would like to promote the perspective for its intrinsic appeal and applicability. Among other benefits, it facilitates the study of the vanishing ideal.

The transition goes as follows. Recall in Definition 1.1 that the SV generator takes as additional parameters a positive integer l and an arbitrary choice of distinct abscissas $a_i \in \mathbb{F}$ for each of the original variables x_i, $i \in [n]$. When $l = 1$, SV^1 takes as seed two fresh variables, y and z, and can be described succinctly in terms of the Lagrange interpolants L_i for the set of abscissas. Plugging in an explicit expression for the Lagrange interpolants, we have:

$$x_i \leftarrow z \cdot L_i(y) \doteq z \cdot \prod_{j \in [n] \setminus \{i\}} \frac{y - a_j}{a_i - a_j}. \tag{1.2}$$

By rescaling, the denominators on the right-hand side of (1.2) can be cleared, resulting in the following somewhat simpler substitution:

$$x_i \leftarrow z \cdot \prod_{j \in [n] \setminus \{i\}} (y - a_j). \tag{1.3}$$

The vanishing ideals of (1.3) and SV^1 are the same up to rescaling each variable to match the rescaling from (1.2) to (1.3).

More importantly, we apply the change of variables $z \leftarrow z' / \prod_{j \in [n]} (y - a_j)$. The resulting substitution now uses rational functions of the seed:

$$x_i \leftarrow \frac{z'}{y - a_i}. \tag{1.4}$$

The notion of vanishing ideal naturally extends to rational function substitutions. The change of variables from (1.3) to (1.4) establishes that any polynomial vanishing on (1.3) also vanishes on (1.4). The change of variables is invertible (the inverse is $z' \leftarrow z \cdot \prod_{j \in [n]} (y - a_j)$), so any polynomial vanishing on (1.4) also vanishes on (1.3). We conclude that the vanishing ideal of (1.4) is the same as that of SV^1 up to rescaling the variables.

Note that, for fixed y and z', (1.4) may be interpreted as first forming a univariate rational function $f(\alpha) = \frac{z'}{y - \alpha}$ (depending on y and z' but independent of i) and then substituting $x_i \leftarrow f(a_i)$. As y and z' vary, f ranges over all rational functions in α with numerator degree zero and denominator degree one. We denote (1.4) by RFE_1^0, where RFE is a short-hand for

Rational Function Evaluation, 0 bounds the numerator degree, and 1 bounds the denominator degree.

As a generator, RFE_1^0 naturally generalizes to RFE_l^k for arbitrary $k, l \in \mathbb{N}$.

Definition 1.2 (RFE generator)**.** The *Rational Function Evaluation Generator (RFE)* for $\mathbb{F}[x_1, \ldots, x_n]$ is parametrized by the following data:

- A non-negative integer k, the *numerator degree*.
- A non-negative integer l, the *denominator degree*.
- For each $i \in [n]$, a distinct *abscissa* $a_i \in \mathbb{F}$.

The generator RFE_l^k takes as seed a rational function $f \in \mathbb{F}(\alpha)$ such that f can be written as g/h for some $g, h \in \mathbb{F}[\alpha]$ with $\deg(g) \leq k$, $\deg(h) \leq l$, and $h(a_i) \neq 0$ for all $i \in [n]$. From f, it generates the substitution $x_i \leftarrow f(a_i)$ for each $i \in [n]$.

There are multiple ways to parametrize the seed of RFE_l^k using scalars; the flexibility to choose is a source of convenience. We refer to Section 2 for a discussion on different parametrizations, as well as on how large the underlying field $\mathbb{F}$ must be. As is customary in the context of blackbox derandomization of PIT, we assume that $\mathbb{F}$ is sufficiently large, possibly by taking a field extension.

The connection between RFE_1^0 and SV^1 extends as follows. For higher values of l, SV^l is defined as the sum of l independent instantiations of SV^1. The same transformations as above relate SV^l and the sum of l independent instantiations of RFE_1^0. Partial fraction decomposition expresses a (non-degenerate) univariate rational function with numerator of degree $l-1$ and denominator of degree l as a sum of l rational functions with numerators of degree 0 and denominators of degree 1. As a result, SV^l is equivalent in power to RFE_l^{l-1}, up to variable rescaling. See Section 2 for a formal treatment.

For parameter values $k \neq l-1$, there is no SV generator that corresponds to RFE_l^k, but $\mathrm{SV}^{\max(k+1,l)}$ encompasses RFE_l^k (up to rescaling) and uses at most twice as many seed variables. Thus, the RFE-generator and the SV-generator efficiently hit the same classes of polynomials. However, RFE induces simple linear dependencies on the seed variables—as opposed to the nonlinear dependencies produced by SV—which enables our approach for determining the vanishing ideal. The moral is that, even though polynomial substitutions are sufficient for derandomizing PIT, it nevertheless helps to consider rational substitutions. Their use may simplify analysis and arguably yield more elegant constructions.

As another indication of the power of rational substitutions, an alternate interpretation of the RFE generator is that it substitutes the ratio of two linear functions of the seed variables, where the coefficients of the linear functions are powers of the abscissas. A generator that only substitutes linear functions—as opposed to a ratio of linear functions—of the seed variables must have seed length n in order to hit all linear polynomials. This is because if the seed length were less than n, then there exists a nontrivial linear combination of the n variables that becomes zero after substitution. In contrast, the simplest nontrivial case of RFE, RFE_1^0, hits all linear polynomials and only needs a seed of length 2.

Generating set. Our first result describes a small and explicit generating set for the vanishing ideal of RFE. It consists of instantiations of a single determinant expression.

Theorem 1.3 (generating set). *Let $k, l, n \in \mathbb{N}$ and let a_i for $i \in [n]$ be distinct elements of $\mathbb{F}$. The vanishing ideal of RFE_l^k in $\mathbb{F}[x_1, \ldots, x_n]$ for the given choice of abscissas $(a_i)_{i\in[n]}$ is generated by the following polynomials over all choices of $k + l + 2$ indices $i_1, i_2, \ldots, i_{k+l+2} \in [n]$:*

$$\mathrm{EVC}_l^k[i_1, i_2, \ldots, i_{k+l+2}] \doteq \det \begin{bmatrix} a_{i_j}^k & a_{i_j}^{k-1} & \ldots & 1 & a_{i_j}^l x_{i_j} & a_{i_j}^{l-1} x_{i_j} & \ldots & x_{i_j} \end{bmatrix}_{j=1}^{k+l+2}. \tag{1.5}$$

Moreover, for any fixed set $C \subseteq [n]$ of size $k + 1$, the polynomials $\mathrm{EVC}_l^k[C \sqcup L]$ form a generating set of minimum size when L ranges over all $(l + 1)$-subsets of $[n]$ that are disjoint from C, where

$$\mathrm{EVC}_l^k[S] \doteq \mathrm{EVC}_l^k[i_1, i_2, \ldots, i_{|S|}]$$

for $S = \{i_1, \ldots, i_{|S|}\} \subseteq [n]$ with $i_1 < i_2 < \cdots < i_{|S|}$.

The name "EVC" is a shorthand for "Elementary Vandermonde Circulation". Later in this overview and in Section 9 we discuss a representation of polynomials using alternating algebra, with connections to notions from network flow. In this representation, polynomials in the vanishing ideal coincide with circulations, and instantiations of EVC are the elementary circulations.

We refer to the set C in Theorem 1.3 as a *core*. The core C plays a similar role as in a combinatorial sunflower except that, unlike the petals of a sunflower, the sets L do not need to be disjoint outside the core.

Example 1.4. Consider the special case where $k = 0$ and $l = 1$. The generator for $\mathrm{Van}[\mathrm{RFE}_1^0]$ when $i_1 = 1$, $i_2 = 2$, and $i_3 = 3$ is given by

$$\mathrm{EVC}_1^0[1, 2, 3] \doteq \begin{vmatrix} 1 & a_1x_1 & x_1 \\ 1 & a_2x_2 & x_2 \\ 1 & a_3x_3 & x_3 \end{vmatrix} = (a_1 - a_2)x_1x_2 + (a_2 - a_3)x_2x_3 + (a_3 - a_1)x_3x_1.$$

For any fixed $i^* \in [n]$, the polynomials $\mathrm{EVC}_1^0[S]$ form a generating set of minimum size when S ranges over all subsets of $[n]$ of size 3 that contain $C = \{i^*\}$. As an aside, they also constitute minimal polynomials not computable by read-once formulas, which is consistent with the fact that SV^1 hits all read-once formulas (see Theorem 5.7).

In general, the generators EVC_l^k are nonzero, multilinear, homogeneous polynomials of degree $l + 1$, and they have nonzero coefficients for all multilinear monomials of degree $l + 1$. Each generating set of minimum size in Theorem 1.3 yields a Gröbner basis with respect to every monomial order that prioritizes the variables outside C. A Gröbner basis is a special generating set that allows solving ideal-membership queries more efficiently, among other problems in computational algebra [12, 1]. Computing Gröbner bases for general ideals is exponential-space complete [36, 38]. Theorem 1.3 represents a rare instance of a natural and interesting ideal for which we know a small and explicit Gröbner basis. See the end of Section 3 for more background on Gröbner bases.

To gain some intuition about dependencies between the generators EVC_l^k, note that permuting the order of the variables used in the construction of EVC_l^k yields the same polynomial or minus that polynomial, depending on the sign of the permutation. This follows from the determinant structure of EVC_l^k and is the reason why we need to fix the order of the variables in order to obtain a generating set of minimum size. More profoundly, the following relationship holds for every choice of $k+l+3$ indices $i_1, i_2, \ldots, i_{k+l+3} \in [n]$ and every univariate polynomial q of degree at most k:

$$\det\begin{bmatrix} q(a_{i_j}) & a_{i_j}^k & a_{i_j}^{k-1} & \ldots & 1 & a_{i_j}^l x_{i_j} & a_{i_j}^{l-1} x_{i_j} & \ldots & x_{i_j} \end{bmatrix}_{j=1}^{k+l+3} = 0. \tag{1.6}$$

The determinant in (1.6) vanishes because the first column of the matrix is a linear combination of the next $k+1$. A minor expansion across the first column expresses the determinant of the matrix as a linear combination of minors, and each minor is an instantiation of EVC_l^k. Since (1.6) vanishes, the minor expansion yields a linear dependency for every nonzero polynomial q of degree at most k. In fact, when $\{i_1, \ldots, i_{k+l+3}\}$ varies over subsets of $[n]$ containing a fixed core of size $k+1$, the equations (1.6) generate *all* linear dependencies among instances of EVC_l^k.

As corollaries to Theorem 1.3 we obtain the following tight bounds on $\mathrm{Van}[\mathrm{RFE}_l^k]$. The bounds hold for every way to choose the parameters in Definition 1.2, including the abscissas.

Corollary 1.5. *The minimum* degree *of a nonzero polynomial in* $\mathrm{Van}[\mathrm{RFE}_l^k]$ *equals* $l+1$.

Corollary 1.5 proves a conjecture by Fournier and Korwar [20] (additional partial results reported in [35]) that there exists a polynomial of degree $l+1$ in $n = 2l+1$ variables that SV^l fails to hit. The conjecture follows because the generators for $\mathrm{Van}[\mathrm{SV}^l]$ have degree $l+1$ and use $2l+1$ variables. See also Corollary 3.9 in Section 3.

As none of the generators contain a monomial of support l or less, the same holds for every nonzero polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$. This extends the known property that SV^l hits every polynomial that contains a monomial of support l or less [48]. See Proposition 5.1 and Theorem 1.8 for a strengthening in the case of multilinear polynomials.

Corollary 1.6. *The minimum* sparseness, *i. e., number of monomials, of a nonzero polynomial in* $\mathrm{Van}[\mathrm{RFE}_l^k]$ *equals* $\binom{k+l+2}{l+1}$.

The generators EVC_l^k realize the bound in Corollary 1.6 as they exactly contain all multilinear monomials of degree $l+1$ that can be formed out of their $k+l+2$ variables. The claim that no nonzero polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$ contains fewer than $\binom{k+l+2}{l+1}$ monomials requires an additional combinatorial argument (see Lemma 6.1). It is a (tight) quantitative strengthening of the known property that SV^l hits every polynomial with fewer than 2^l monomials [7, 26, 14, 19]. Note that for $k = l-1$ we have that $\binom{k+l+2}{l+1} = \binom{2l+1}{l+1} = \Theta(2^{2l}/\sqrt{l})$. One consequence is that for SV^l to hit all polynomials with m monomials, a seed length of $l = \Omega(\log m)$ is required. In particular, hitting sparse polynomials requires $l = \Omega(\log n)$.

Another consequence deals with set-multilinearity, a common restriction in works on derandomizing PIT and algebraic circuit lower bounds. A polynomial p of degree $l+1$ in a set of variables $\{x_1, \ldots, x_n\}$ is said to be set-multilinear if $[n]$ can be partitioned as $[n] = X_1 \sqcup X_2 \sqcup \cdots \sqcup X_{l+1}$ such that every monomial in p is a product $x_{i_1} \cdot x_{i_2} \cdot \cdots \cdot x_{i_{l+1}}$, where $i_j \in X_j$. Note that set-multilinearity implies multilinearity but not the other way around. As the generators EVC_l^k are not set-multilinear, it is not immediately clear from Theorem 1.3 whether $\mathrm{Van}[\mathrm{RFE}_l^k]$ contains nontrivial set-multilinear polynomials of any degree. However, a variation on the construction of the generators EVC_l^k yields explicit set-multilinear homogeneous polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$ of degree $l+1$ where each X_j has size $k+2$ (see Definition 7.1). We denote them by ESMVC_l^k, where ESMVC stands for "Elementary Set-Multilinear Vandermonde Circulation". ESMVC_l^k contains all monomials of the form $x_{i_1} \cdot x_{i_2} \cdot \cdots \cdot x_{i_{l+1}}$ with $i_j \in X_j$. For any variable partition $(X_1, X_2, \ldots, X_{l+1})$ with $|X_1| = \cdots = |X_{l+1}| = k+2$, ESMVC_l^k is the only set-multilinear polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$ with that variable partition, up to a scalar multiple, and exhibits the following extremal property. See also Theorem 7.4.

Corollary 1.7. *The minimum* partition class size *of a nonzero set-multilinear polynomial of degree* $l+1$ *in* $\mathrm{Van}[\mathrm{RFE}_l^k]$ *equals* $k+2$.

Membership test. Our second characterization of the vanishing ideal of RFE can be viewed as a structured membership test. Given a polynomial p, there is a generic way to test whether p belongs to the vanishing ideal of a generator G, namely by symbolically substituting G into p and verifying that the result simplifies to zero. When G is a polynomial substitution, the well-known transformation of a generator into a deterministic blackbox PIT algorithm yields another test: Verify $p(G) = 0$ for a sufficiently large set of substitutions into the seed variables. By clearing denominators, the same goes for rational substitutions like RFE_l^k.

While the generic test works, one cannot extract G-specific insight into whether or why G hits any particular polynomial. In contrast, our membership test uses the specific structure of G and provides useful insight. Building on the generating set of Theorem 1.3, we state our structured test for membership of multilinear polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$ in terms of partial derivatives and zero substitutions. Several prior papers demonstrated the utility of those operations in the context of derandomizing PIT using the SV generator, especially for syntactically multilinear models [48, 32, 7].

Theorem 1.8 (membership test for multilinear polynomials)**.** *Let $k, l, n \in \mathbb{N}$ and let a_i for $i \in [n]$ be distinct elements of $\mathbb{F}$. A multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ belongs to* $\mathrm{Van}[\mathrm{RFE}_l^k]$ *if and only if both of the following conditions hold:*

1. *There are no monomials of degree l or less, nor of degree $n-k$ or more, in p.*

2. *For all disjoint subsets $K, L \subseteq [n]$ with $|K| = k$ and $|L| = l$, $\partial_L p|_{K \leftarrow 0}$ is zero upon the following substitution for each $i \in [n] \setminus (K \cup L)$, where z denotes a fresh variable:*

$$x_i \leftarrow z \cdot \frac{\prod_{j \in K}(a_i - a_j)}{\prod_{j \in L}(a_i - a_j)}. \tag{1.7}$$

A few technical comments regarding the statement are in order. The first part of condition 1 in Theorem 1.8 generalizes the known property that SV^l hits every multilinear polynomial that contains a monomial of degree l or less [48]. As for the second part, see Proposition 5.1 for more discussion. The two parts together imply that all multilinear polynomials on $n \leq k+l+1$ variables are hit by RFE_l^k.

In condition 2, $\partial_L p|_{K\leftarrow 0}$ denotes the polynomial obtained by taking the partial derivative of p with respect to every variable in L and setting all the variables in K to zero. Because of the multilinearity, the order of the operations does not matter, and the resulting polynomial depends only on variables in $[n] \setminus (K \cup L)$. The substitution (1.7) can be viewed as $x_i \leftarrow f(a_i)$, where

$$f(\alpha) = z \cdot f_{K,L}(\alpha) \doteq z \cdot \frac{\prod_{j\in K}(\alpha - a_j)}{\prod_{j\in L}(\alpha - a_j)}$$

is a valid seed of RFE_l^k for polynomials in the variables x_i, $i \in [n] \setminus (K \cup L)$. Upon substitution, $\partial_L p|_{K\leftarrow 0}$ becomes a univariate polynomial q of degree at most $n-k-l$ in the fresh variable z. In the case where p is homogeneous, q has at most one term, and q is nonzero if and only q is nonzero at $z=1$. In general, for any fixed set Z of $n-k-l+1$ elements of $\mathbb{F}$, q is nonzero if and only if q is nonzero at some $z \in Z$.

Theorem 1.8 can be understood as stating that a multilinear polynomial p is hit by RFE_l^k if and only if p has a monomial supported on few or all-but-few variables, or else there is a set of k zero substitutions, K, and a set of l partial derivatives, L, whose application to p leaves a polynomial that is nonzero after substituting $x_i \leftarrow z \cdot f_{K,L}(a_i)$. By judiciously choosing variables for the zero substitutions and partial derivatives, prior papers managed to simplify polynomials p of certain types and reduce PIT for p to PIT for simpler instances, resulting in efficient recursive algorithms. In Section 5, we develop a general framework for such algorithms and prove correctness directly from Theorem 1.8. Moreover, because Theorem 1.8 is a precise characterization, any argument that SV or RFE hits a class of multilinear polynomials can be converted into one within our framework, i. e., into an argument based on zero substitutions and partial derivatives. Thus, Theorem 1.8 shows that these tools harness the complete power of SV and RFE for multilinear polynomials.

Applications. We illustrate the utility of our characterizations of the vanishing ideal of RFE in the two directions mentioned before.

Derandomization. To start, we demonstrate how Theorem 1.8 yields an alternate proof of the result from [41] that SV^1—equivalently, RFE_1^0—hits every nonzero read-once formula F. Whereas the original proof hinges on a clever ad-hoc argument, our proof (described in Section 5) is entirely systematic and amounts to a couple straightforward observations in order to apply Theorem 1.8.

As a proof of concept of the additional power of our characterization for derandomization, we make progress in a well-studied model for algebraic computation, namely read-once oblivious algebraic branching programs (ROABPs). An ROABP consists of a layered digraph, the width of which constitutes an important complexity parameter. We refer to Section 8.1 for more background.

Theorem 1.9 (ROABP hitting property)**.** *For any integer $l \geq 1$, SV^l hits the class of polynomials computed by read-once oblivious algebraic branching programs of width less than $(l/3)+1$ that contain a monomial of degree at most $l+1$.*

To the best of our knowledge, Theorem 1.9 is incomparable to the known results for ROABPs [45, 29, 30, 17, 15, 3, 6, 26, 25, 24, 46, 10]. Without the restriction that the polynomial has a monomial of degree at most $l+1$, Theorem 1.9 would imply a fully blackbox polynomial-time identity test for the class of constant-width ROABPs. No such test has been proven to exist at this time; prior work requires either quasipolynomial time or else opening the blackbox, such as by knowing the order in which the variables are read.

With the restriction, hitting the class in Theorem 1.9 with SV^l represents fairly specialized progress. This is because SV^{l+1} is well-known to hit every polynomial containing a monomial of support $l+1$ or less, and thus hits the class in Theorem 1.9, irrespective of the restriction on ROABP width. That said, the method of proof of Theorem 1.9 diverges significantly from prior uses of the SV generator and therefore may be of independent interest. We elaborate on the method more when we discuss the techniques of this paper, but for now, we point out that most prior uses of the SV generator rely on combinatorial arguments, i. e., arguments that depend only on which monomials are present in the polynomials to hit. Theorem 1.9 necessarily goes beyond this because there is a polynomial in $\mathrm{Van}[\mathrm{SV}^l]$ of degree $l+1$ that has the same monomials as a polynomial computed by an ROABP of width 2, which by Theorem 1.9 is not in $\mathrm{Van}[\mathrm{SV}^l]$ for $l \geq 4$. Namely, any instance of ESMVC_l^{l-1} contains exactly all the monomials of the form $x_{i_1} \cdot x_{i_2} \cdot \cdots \cdot x_{i_{l+1}}$ with $(i_1, \ldots, i_{l+1}) \in X_1 \times \cdots \times X_{l+1}$ for some disjoint sets X_j; the same goes for $\prod_j \sum_{i_j \in X_j} x_{i_j}$, which is computed by an ROABP of width 2.

Lower bounds. Our result for ROABPs also illustrates this direction. Our derandomization result for the class in Theorem 1.9 is equivalent to the following lower bound.

Theorem 1.10 (ROABP lower bound)**.** *For any integer $l \geq 1$, any nonzero polynomial in* $\mathrm{Van}[\mathrm{SV}^l]$ *that contains a monomial of degree at most $l+1$, requires ROABP width at least* $(l/3)+1$.

Such a lower bound is interesting because there are appealing polynomials meeting the conditions, in particular the generators EVC_l^{l-1} as well as ESMVC_l^{l-1}. Other hardness of representation results follow in a similar manner from prior hitting properties of SV in the literature. The following lower bounds apply to computing both EVC_l^{l-1} and ESMVC_l^{l-1}:

- Any syntactically multilinear formula must have at least $\Omega(\log(l)/\log\log(l))$ reads of some variable [7, Theorem 6.3].
- Any sum of read-once formulas must have at least $\Omega(l)$ summands [41, Corollary 5.2].
- There exists an order of the variables such that any ROABP with that order must have width at least $2^{\Omega(l)}$ [15, Corollary 4.3].

- Any $\Sigma\mathrm{m} \bigwedge \Sigma\Pi^{O(1)}$ formula must have top fan-in at least $2^{\Omega(l)}$ [14]; see also [19, Lemma 5.12].
- Lower bounds over characteristic zero for circuits with locally-low algebraic rank [37, Lemma 5.2].

Techniques. Many of our results ultimately require showing that, under suitable conditions, RFE hits a polynomial p. A recurring analysis fulfills this role in the proofs of Theorems 1.3, 1.8, and 1.9. We take intuition from the analytic setting (e. g., $\mathbb{F} = \mathbb{R}$) and study the behavior of $p(\mathrm{RFE})$ as a function of the seed's zeroes and poles. When they are near the abscissas of chosen variables of p, the behavior is dominated by the contributions of the monomials of p for which the variables with abscissas near zeros have minimal degree and the variables with abscissas near poles have maximal degree. This allows us to analyze a first approximation to $p(\mathrm{RFE})$ by "zooming in" on the contributions of the monomials in which the chosen variables have extremal degrees. If the first approximation is nonzero, then we can conclude that RFE hits p. We capture the technique in our Zoom Lemma (Lemma 4.3). Formal Laurent series can express the analytic intuition purely algebraically. We provide a proof from first principles that does not require any background in Laurent series and works over all fields.

Theorem 1.3 states the equality $I = \mathrm{Van}[\mathrm{RFE}_l^k]$ of two ideals, where I denotes the ideal generated by all instantiations of EVC_l^k, and $\mathrm{Van}[\mathrm{RFE}_l^k]$ the vanishing ideal of RFE_l^k.

- The inclusion $\subseteq$ follows from linearizing the defining equations of RFE_l^k (Lemma 3.1). The technique mirrors the use of resultants to compute implicit equations for rational plane curves. This is where the switch from SV to RFE helps.
- To establish the inclusion $\supseteq$ we first show that the equivalence class of any polynomial p modulo I contains a representative r whose monomials exhibit the combinatorial structure of a core (Lemma 3.7). If $p \notin I$, r is nonzero. The core structure of r then allows us to apply the zooming-in mechanism such that the resulting first approximation to r is nonzero, in which case the Zoom Lemma tells us that RFE_l^k hits r (Lemma 3.8). By the inclusion $\subseteq$, we conclude that RFE_l^k hits p.

The proof of Theorem 1.8 also relies on the Zoom Lemma. Membership to the ideal is equivalent to the vanishing of all coefficients of the expansion of $p(\mathrm{RFE})$. The application of the Zoom Lemma can be viewed as determining a small number of coefficients sufficient to guarantee that their vanishing implies all coefficients vanish. The restriction to multilinear polynomials p allows us to express the zoomed-in contributions of p as the result of applying partial derivatives and zero-substitutions.

Theorem 1.9 makes use of the characterization of the minimum width of a read-once oblivious algebraic branching program computing a polynomial p as the maximum rank of the monomial coefficient matrices of p for various variable partitions [42]. The result is effectively about polynomials p that are homogeneous of degree $l + 1$, in which case the monomial coefficient matrices have a block-diagonal structure with $l + 2$ blocks. An application of the Zoom Lemma in the contrapositive yields linear equations between elements of consecutive blocks under the assumption that SV^l fails to hit p. When some block is zero, the equations yield a Cauchy system on the rows or columns of its neighboring blocks. Based on the fact that Cauchy systems have full rank and exploiting the specific structure, we deduce several constraints on the row-space/column-space of the neighboring blocks. A careful analysis and case analysis based on the number of zero blocks yields a rank lower bound of at least $(l/3)+1$ for a well-chosen partition of the variables.

We point out that, in the preceding application, the Zoom Lemma is instantiated several times in parallel to form a large system of equations on the coefficients of p, and the whole system is necessary for the analysis. This stands in contrast to most prior work using SV, which can be cast as using knowledge of how p is computed to guide a search for a *single* fruitful instantiation of the Zoom Lemma.

Alternating algebra representation. The inspiration for several of our results stems from expressing the polynomials EVC_l^k using concepts from alternating algebra (also known as exterior algebra or Grassmann algebra). In fact, the membership test for the ideal generated by the instantiations of EVC_l^k in Theorem 1.8 is based on the relationship $\partial^2 = 0$ from alternating algebra. Our original statement and proof of the theorem made use of that framework, but we managed to eliminate the alternating algebra afterwards. Still, as we find the perspective insightful and potentially helpful for future developments, we describe the connection briefly here and in more detail in Section 9. We explain the intuition for the simple case where the degree of the polynomial p equals $l+1$. In that setting, belonging to the ideal generated by the polynomials EVC_l^k is equivalent to being in their linear span.

The alternating algebra $\Lambda^*(U)$ of a vector space U over a field $\mathbb{F}$ consists of the closure of U under an additional binary operation, referred to as "wedge" and denoted $\wedge$, which is bilinear, associative, and satisfies

$$u \wedge u = 0 \tag{1.8}$$

for every $u \in U$. This determines a well-defined algebra. When the characteristic of $\mathbb{F}$ is not 2, (1.8) can equivalently be understood as anti-commutativity:

$$u_1 \wedge u_2 = -(u_2 \wedge u_1) \tag{1.9}$$

for every $u_1, u_2 \in U$. For any characteristic and $u_1, u_2, \ldots, u_t \in U$,

$$u_1 \wedge u_2 \wedge \cdots \wedge u_t \tag{1.10}$$

is nonzero iff the u_i's are linearly independent, and any permutation of the order of the vectors in (1.10) yields the same element of $\Lambda^*(U)$ up to a sign. The sign equals the sign of the permutation, whence the name "alternating algebra." If U has a basis $V = \{v_1, \ldots, v_n\}$ of size n, then a basis for $\Lambda^*(U)$ can be formed by all 2^n expressions of the form (1.10), where the u_i's range over all subsets of V and are taken in some fixed order. Considering the elements of V as vertices, the basis elements of $\Lambda^*(U)$ can be thought of as the oriented simplices of all dimensions that can be built from V.

Anti-commutativity arises naturally in the context of network flow, where V denotes the vertices of the underlying graph, and a wedge $v_1 \wedge v_2$ of level $t = 2$ represents one unit of flow from v_1 to v_2. Equation (1.9) reflects the fact that one unit of flow from v_1 to v_2 cancels with one unit of flow from v_2 to v_1. The adjacent levels $t = 1$ and $t = 3$ also have natural interpretations in the flow setting: v_1 (the element of $\Lambda^*(U)$ of the form (1.10) with $t = 1$) represents one unit of surplus flow at v_1 (the vertex of the graph), and $v_1 \wedge v_2 \wedge v_3$ abstracts a circulation of one unit along the directed cycle $v_1 \to v_2 \to v_3 \to v_1$.

The different levels are related by so-called boundary maps, which are linear transformations that map a simplex to a linear combination of its subsimplices of one dimension less. The maps are parametrized by a linear weight function $w : U \to \mathbb{F}$, and defined on the vertices by

$$\partial_w : v_1 \wedge v_2 \wedge \cdots \wedge v_t \mapsto \sum_{i=1}^{t} (-1)^{i+1} w(v_i)\, v_1 \wedge \cdots \wedge v_{i-1} \wedge v_{i+1} \wedge \cdots \wedge v_t, \tag{1.11}$$

an expression resembling the minor expansion of a determinant along a column $[w(v_i)]_{i=1}^t$. In the flow setting, using $w \equiv 1$, applying ∂_1 to $v_1 \wedge v_2$ yields $v_2 - v_1$, the superposition of demand at v_1 and surplus at v_2 corresponding to one unit of flow from v_1 to v_2. Likewise, ∂_1 sends the abstract elementary circulation $v_1 \wedge v_2 \wedge v_3$ to the superposition of the three edge flows that make up the directed 3-cycle $v_1 \to v_2 \to v_3 \to v_1$. A linear combination p of terms (1.10) with $t = 2$ represents a valid circulation iff it satisfies conservation of flow at every vertex, which can be expressed as $\partial_1(p) = 0$, i. e., p is in the kernel of ∂_1. An equivalent criterion is for p to be the superposition of circulations around directed 3-cycles, which can be expressed as p being in the image of ∂_1. The relationship $\mathrm{im}(\partial_w) = \ker(\partial_w)$ between the image and the kernel of a boundary map holds for any nonzero w, and generalizes to composed boundary maps: For any linearly independent $w_1, \ldots, w_{k+1}$, it holds that

$$\mathrm{im}\left(\partial_{w_{k+1}} \circ \partial_{w_k} \circ \cdots \circ \partial_{w_1}\right) = \bigcap_{r=1}^{k+1} \ker\left(\partial_{w_r}\right). \tag{1.12}$$

When $w_1, \ldots, w_{k+1}$ are linearly dependent, $\partial_{w_{k+1}} \circ \cdots \circ \partial_{w_1}$ is the zero map.

In the context of RFE, the set V consists of a distinct vertex v_i for each variable x_i, and simplices correspond to multilinear monomials. The anti-commutativity of $\wedge$ coincides with the fact that swapping two arguments to EVC_l^k means swapping two rows in (1.5), which changes the sign of the determinant. Using boundary maps, $\mathrm{EVC}_l^k[i_1, i_2, \ldots, i_{k+l+2}]$ can be viewed as $\partial_\omega(v_{i_1} \wedge v_{i_2} \wedge \cdots \wedge v_{i_{k+l+2}})$, where $\partial_\omega \doteq \partial_{w_{k+1}} \circ \partial_{w_k} \circ \cdots \circ \partial_{w_1}$ and $w_r(v_i) \doteq (a_i)^{r-1}$. By (1.12), this means that EVC_l^k is in the kernel of ∂_{w_r} for each $r \in [k+1]$, or equivalently, in the kernel of ∂_w for each $w : U \to \mathbb{F}$ of the form $w(v_i) = q(a_i)$ where q is a polynomial of degree at most k. In fact, (1.12) implies that the linear span of the generators EVC_l^k consists exactly of the polynomials of degree $l+1$ in this kernel. The linear span coincides with the polynomials of degree $l+1$ in the ideal generated by the polynomials EVC_l^k. For multilinear polynomials, being in the kernel can be expressed in terms of zero substitutions and partial derivatives as in Theorem 1.8. This yields an alternate route for deriving our membership test for multilinear polynomials of degree

$d = l + 1$ in the ideal generated by the instantiations of EVC_l^k, which by Theorem 1.3 agrees with $\mathrm{Van}[\mathrm{RFE}_l^k]$. In the basic case where $k = 0$ and $l = 1$, only the weight function $w \equiv 1$ needs to be considered and the kernel requirement coincides with flow conservation. We refer to Section 9 for the general multilinear case of arbitrary degree.

Related recent work and further research. We propose to systematically investigate the power of generators by characterizing their vanishing ideals. As we demonstrated for SV and RFE, such characterizations can exhibit both strengths and weaknesses of the generator.

Specific other generators of interest include Klivans–Spielman [34] and generators based on the matrix rank condenser by Gabizon and Raz [21, 33, 16]. A related direction is figuring out how vanishing ideals are affected when manipulating generators. Examples include the RFE generator with pseudorandom abscissas, or work that relates the vanishing ideal of a combination of generators to the vanishing ideals of the constituent generators. In particular, a combination of SV with Klivans–Spielman appears in the literature [32, 15, 19], where the latter is used to effectively hit sparse polynomials, which our results show that SV does not.

The generator SV^l is the canonical example of an l-wise independent generator in the algebraic setting. Understanding the power of l-wise independent generators more broadly, e. g., as formalized in [18, 39], could lead to useful insights for derandomizing PIT. This work demonstrates explicit polynomials like EVC_l^{l-1} and ESMVC_l^{l-1} that are not automatically hit by l-wise independence as they are not hit by SV^l. Is there a deeper underlying reason related to l-wise independence?

A generator hits all polynomials from a resource-bounded class iff no nonzero polynomial in the vanishing ideal can be computed within those resources. Chatterjee and Tengse [11] recently showed the following generic limitation: The vanishing ideal of any generator computable by algebraic circuits of polynomial size in the number of variables contains a nonzero polynomial computable in VPSPACE. From this perspective, our results exhibit a weakness of SV and RFE in that their vanishing ideals contain nonzero polynomials from the presumably much smaller class VBP. In fact, EVC_l^k is a polynomial depending on only $k + l + 2$ variables and is computable by a branching program of size polynomial in the number of variables. Thus, in order to hit all branching programs of size s, SV and RFE require a seed length $k + l + 2 = s^{\Omega(1)}$.

A related question is whether the generators we have identified have minimal (or approximately minimal) complexity in the vanishing ideal. Andrews and Forbes [8] recently established such a result for a generator that substitutes an $n \times m$ matrix of variables with the product of $n \times l$ and $l \times m$ matrices of variables for small l. The vanishing ideal of their generator is straightforwardly generated by $(l + 1) \times (l + 1)$ minors. For this vanishing ideal the authors manage to show that every nonzero element is at least as hard as computing $\Theta(l^{1/3}) \times \Theta(l^{1/3})$ determinants (under simple reductions and in the sense of border complexity).

Lastly, we list some avenues for improving specific aspects of our results. Theorem 1.8 represents an elementary deterministic membership test in the vanishing ideal of RFE_l^k for *multilinear* polynomials. Can the elementary test can be extended to all polynomials? From the alternating algebra perspective, the test relies an the convenient one-to-one correspondence between multilinear polynomials and elements of the alternating algebra. For general polynomials, this correspondence is no longer one-to-one, and the resulting membership test is nondeterministic.

Another target is eliminating degree restrictions for our characterizations of specialized classes of polynomials, in particular in Theorem 1.9 for ROABPs. Removing the degree restriction for ROABPs would result in a full blackbox derandomization of constant-width ROABPs. An alternative possibility is that, through better analysis of the vanishing ideal, it turns out that RFE has limitations in derandomizing constant-width ROABPs.

Organization. We start in Section 2 with formal aspects of the RFE generator that have been omitted from the informal discussion thus far. We construct the generating set for the vanishing ideal (Theorem 1.3) in Section 3, followed by the Zoom Lemma in Section 4. The ideal membership test (Theorem 1.8) is developed in Section 5. We present the results on sparseness in Section 6, and the ones on set-multilinearity in Section 7. Background on ROABPs and our result on derandomizing PIT for ROABPs (Theorem 1.9) are covered in Section 8. We end our paper in Section 9 with a further discussion of the alternating algebra representation and an alternate derivation of the membership test for multilinear polynomials in the ideal generated by the instances of EVC_l^k.

2 RFE Generator

We defined the RFE generator in the overview but omitted some of the formal details. In this section, we discuss different parametrizations of RFE as well as how to obtain deterministic blackbox PIT algorithms from a generator and how large the underlying field $\mathbb{F}$ must be. We also state and establish the precise relationship between RFE_l^{l-1} and SV^l.

In Definition 1.2, we defined RFE as a set of substitutions formed by varying the seed f over certain rational functions with coefficients in $\mathbb{F}$. Meanwhile, our analyses proceed by parametrizing f by scalars, abstracting the scalar parameters as fresh formal variables, and calculating in the field of rational functions in those variables. The approaches are equivalent over large enough fields, and the flexibility to choose is a source of convenience. Here are some natural parametrizations of f:

Coefficients. Select scalars $g_0, \ldots, g_k, h_0, \ldots, h_l \in \mathbb{F}$ and set

$$f(\alpha) = \frac{g_k \alpha^k + g_{k-1}\alpha^{k-1} + \cdots + g_1\alpha + g_0}{h_l \alpha^l + h_{l-1}\alpha^{l-1} + \cdots + h_1\alpha + h_0},$$

ignoring choices of $h_0, \ldots, h_l$ for which the denominator vanishes at some abscissa.

Evaluations. Fix two collections, $B = \{b_1, \ldots, b_{k+1}\}$ and $C = \{c_1, \ldots, c_{l+1}\}$, each of distinct scalars from $\mathbb{F}$. Then select scalars $g_1, \ldots, g_{k+1}$ and $h_1, \ldots, h_{l+1}$ and set

$$f(\alpha) = \frac{g(\alpha)}{h(\alpha)}$$

where g is the unique degree-k polynomial with $g(b_1) = g_1$, $g(b_2) = g_2, \ldots, g(b_{k+1}) = g_{k+1}$, and h is defined similarly with respect to C. Choices of $h_1, \ldots, h_{l+1}$ that lead h to vanish at some abscissa are ignored.

Note that an explicit formula for g and h in terms of the parameters can be obtained using the Lagrange interpolants with respect to B and C.

Roots. Select scalars $z, s_1, \ldots, s_{k'}, t_1, \ldots, t_{l'} \in \mathbb{F}$ for some $k' \leq k$ and $l' \leq l$ and set

$$f(\alpha) = z \cdot \frac{(\alpha - s_1) \cdot \cdots \cdot (\alpha - s_{k'})}{(\alpha - t_1) \cdot \cdots \cdot (\alpha - t_{l'})},$$

where $\{t_1, \ldots, t_{l'}\}$ is disjoint from the set of abscissas.

In fact, it is no loss of power to restrict to $k' = k$ and $l' = l$.

Hybrids are of course possible, too. For example, Proposition 2.2 below uses the evaluations parametrization for the numerator and roots parametrization for the denominator.

The following lemma justifies that, for any polynomial p, as long as $\mathbb{F}$ is large enough, $p(\mathrm{RFE})$ vanishes with respect to a particular parametrization of RFE if and only if it vanishes with respect to RFE as defined in Definition 1.2. The lemma is an immediate consequence of the well-known analogous result for polynomials, sometimes referred to as the Polynomial Identity Lemma [44, 13, 50, 47, 9].

Lemma 2.1. *Let $\mathbb{F}$ be field, and $f = g/h \in \mathbb{F}(\tau_1, \ldots, \tau_l)$ be a rational function in l variables with $\deg(g) \leq d$ and $\deg(h) \leq d$. Let $S \subseteq \mathbb{F}$ be finite. Then the probability that f vanishes or is undefined when each τ_i is substituted by a uniformly random element of S is at most $2d/|S|$.*

Proof. The rational function f vanishes or is undefined if and only if the polynomial $p \doteq= gh$ vanishes, which happens with probability at most $\deg(p)/|S|$ according to the Polynomial Identity Lemma. ◻

In particular, if $\mathbb{F}$ is infinite, then, for all polynomials p, all the above parametrizations and Definition 1.2 are equivalent for the purposes of hitting p; when p is fixed, the equivalence holds provided $|\mathbb{F}| \geq \mathrm{poly}(n, \deg(p))$. Quantitative bounds on the number of substitutions to perform when testing whether RFE hits p in the blackbox algorithm likewise follow from Lemma 2.1. As is customary in the context of blackbox derandomization of PIT, if $\mathbb{F}$ is not large enough, then one works instead over a sufficiently large extension of $\mathbb{F}$.

We now formally state and argue the close relationship between RFE_l^{l-1} and SV^l that we sketched in Section 1.

Proposition 2.2. *Let l and n be positive integers. There is an invertible diagonal transformation $A : \mathbb{F}^n \to \mathbb{F}^n$ such that, for any polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$, $p(\mathrm{SV}^l) = 0$ if and only if $(p \circ A)(\mathrm{RFE}_l^{l-1}) = 0$.*

In particular, the vanishing ideals of RFE_l^{l-1} and of SV^l are the same up to the rescaling of Proposition 2.2.

Proof of Proposition 2.2. Let $\widehat{\mathbb{F}}$ be the field of rational functions in indeterminates $\upsilon_1, \ldots, \upsilon_l, \zeta_1, \ldots, \zeta_l$ over $\mathbb{F}$. A polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ has $p(\mathrm{SV}^l) = 0$ if and only if p vanishes at the point

$$\left(\sum_{t=1}^{l} \zeta_t \prod_{j \in [n] \setminus \{i\}} \frac{\upsilon_t - a_j}{a_i - a_j} \ : \ i \in [n] \right) \in \widehat{\mathbb{F}}^n. \tag{2.1}$$

Set $A : \mathbb{F}^n \to \mathbb{F}^n$ to be the diagonal linear transformation that divides the coordinate for x_i by $\prod_{j \in [n] \setminus \{i\}} (a_i - a_j)$. A is invertible. Applying A^{-1} to (2.1) yields the point

$$\left(\sum_{t=1}^{l} \zeta_t \prod_{j \in [n] \setminus \{i\}} (\upsilon_t - a_j) \ : \ i \in [n] \right) = \left(\sum_{t=1}^{l} \left(\zeta_t \prod_{j \in [n]} (\upsilon_t - a_j) \right) \frac{1}{\upsilon_t - a_i} \ : \ i \in [n] \right). \tag{2.2}$$

p vanishes at (2.1) if and only if $p \circ A$ vanishes at (2.2). Now let $\widehat{\mathbb{F}}'$ be the field of rational functions in indeterminates $\tau_1, \ldots, \tau_l, \sigma_1, \ldots, \sigma_l$ over $\mathbb{F}$. After the change of variables

$$\zeta_t \leftarrow \frac{1}{\prod_{j \in [n]} (\tau_t - a_j)} \cdot \frac{-\sigma_t}{\prod_{s \neq t} (\tau_t - \tau_s)} \quad \text{and} \quad \upsilon_t \leftarrow \tau_t$$

(2.2) becomes

$$\left(\sum_{t=1}^{l} \frac{\sigma_t}{\left(\prod_{s \neq t} \tau_t - \tau_s\right)} \frac{1}{a_i - \tau_t} \ : \ i \in [n] \right) = \left(\frac{\sum_{t=1}^{l} \sigma_t \prod_{s \neq t} \frac{a_i - \tau_s}{\tau_t - \tau_s}}{\prod_{t=1}^{l} a_i - \tau_t} \ : \ i \in [n] \right) \in \widehat{\mathbb{F}}'^n. \tag{2.3}$$

Since the change of variables is invertible, $p \circ A$ vanishes at (2.2) if and only if it vanishes at (2.3).

Now, viewing $\sigma_1, \ldots, \sigma_l, \tau_1, \ldots, \tau_l$ as seed variables, observe that the right-hand side of (2.3) is $\mathrm{RFE}_l^{l-1}(g/h)$ where g is parametrized by evaluations ($g(\tau_t) = \sigma_t$) and h is parametrized by roots ($\tau_1, \ldots, \tau_l$). It follows that $p \circ A$ vanishes at (2.3) if and only if $(p \circ A)(\mathrm{RFE}_l^{l-1}) = 0$. □

3 Generating Set

In this section, we establish Theorem 1.3, our characterization of the vanishing ideal of RFE in terms of an explicit generating set. For every $k, l \in \mathbb{N}$, we develop a template, EVC_l^k, for constructing polynomials that belong to the vanishing ideal of RFE_l^k such that all instantiations collectively generate the vanishing ideal.

The template can be derived in the following fashion. Fix any seed f of RFE_l^k, and write it as $f = g/h$ where $g(\alpha) = \sum_{d=0}^{k} g_d \alpha^d$ and $h(\alpha) = \sum_{d=0}^{l} h_d \alpha^d$ are respectively polynomials of degree k and l. For any $i \in [n]$, the polynomial $g(a_i)/h(a_i) - x_i \in \mathbb{F}[x_1, \ldots, x_n]$ vanishes by definition at $\mathrm{RFE}_l^k(f)$. While this polynomial varies with f, it does so uniformly. Specifically, after rescaling to $g(a_i) - h(a_i)x_i$, the polynomial depends only *linearly* on the coefficients of g and h. We exploit this uniformity to construct a polynomial that vanishes at $\mathrm{RFE}_l^k(f)$ but that now is *independent* of f. Since f is arbitrary, the constructed polynomial belongs to the vanishing ideal of RFE_l^k.

The construction begins by expressing the vanishing of each $g(a_i) - h(a_i)x_i$ at $\mathrm{RFE}_l^k(f)$ as the following system of equations. Abbreviating

$$\begin{aligned}\vec{g} &\doteq \begin{bmatrix} g_k & g_{k-1} & \dots & g_1 & g_0 \end{bmatrix}^\intercal \\ \vec{h} &\doteq \begin{bmatrix} h_l & h_{l-1} & \dots & h_1 & h_0 \end{bmatrix}^\intercal ,\end{aligned}$$

we write

$$\begin{bmatrix} a_i^k & a_i^{k-1} & \dots & 1 & a_i^l x_i & a_i^{l-1} x_i & \dots & x_i \end{bmatrix}_{i \in [n]} \cdot \begin{bmatrix} \vec{g} \\ -\vec{h} \end{bmatrix} = 0. \tag{3.1}$$

Written this way, (3.1) has the form of a homogeneous system of linear equations. There is one equation for each $i \in [n]$ and one unknown for each of the $k + l + 2$ parameters of the seed f. The system's coefficient matrix has no dependence on f, but for any f, substituting $\mathrm{RFE}_l^k(f)$ into $x_1, \dots, x_n$ yields a system that has a nontrivial solution, namely the vector in (3.1).

Consider, then, the determinant of the square subsystem of (3.1) formed by any $k + l + 2$ rows. It is a polynomial in $\mathbb{F}[x_1, \dots, x_n]$. Because the coefficient matrix in (3.1) is independent of f, the determinant is independent of f. Because the subsystem has a nonzero solution after substituting $\mathrm{RFE}_l^k(f)$ for any f, the determinant vanishes after substituting $\mathrm{RFE}_l^k(f)$ for any f. We conclude that the determinant belongs to the vanishing ideal of RFE_l^k.

Recalling that the determinant for the subsystem consisting of rows $i_1, \dots, i_{k+l+2}$ is identically $\mathrm{EVC}_l^k[i_1, i_2, \dots, i_{k+l+2}]$, we have established:

Lemma 3.1. *For every $k, l \in \mathbb{N}$ and $i_1, i_2, \dots, i_{k+l+2} \in [n]$, $\mathrm{EVC}_l^k[i_1, \dots, i_{k+l+2}] \in \mathrm{Van}[\mathrm{RFE}_l^k]$.*

As we explain at the end of this section, the above derivation is where our use of RFE in lieu of SV plays a critical role. Before moving on, we also point a few elementary properties and an provide an explicit expression for the coefficients of EVC_l^k as products of Vandermonde determinants in the abscissas and a sign term. We introduce the following notation for the underlying Vandermonde matrices.

Definition 3.2 (Vandermonde matrix)**.** For $T = \{i_1, \dots, i_t\} \subseteq [n]$ with $i_1 < \dots < i_t$, we abbreviate the Vandermonde matrix built from a_i for $i \in T$ in increasing order as

$$A_T \doteq \begin{bmatrix} a_{i_1}^{t-1} & \cdots & 1 \\ \vdots & & \vdots \\ a_{i_t}^{t-1} & \cdots & 1 \end{bmatrix}. \tag{3.2}$$

The sign term makes use of the following standard combinatorial quantity.

Definition 3.3 (cross inversions)**.** For $A, B \subseteq [n]$, $\mathrm{XInv}(A, B) \doteq |\{(a, b) \in A \times B \mid a > b\}|$ denotes the number of *cross inversions* between A and B.

Proposition 3.4. *For every $k, l \in \mathbb{N}$, EVC_l^k is skew-symmetric in that, for any $i_1, \dots, i_{k+l+2} \in [n]$ and permutation π of $[k + l + 2]$,*

$$\mathrm{EVC}_l^k[i_1, \dots, i_{k+l+2}] = \mathrm{sign}(\pi) \cdot \mathrm{EVC}_l^k[i_{\pi(1)}, \dots, i_{\pi(k+l+2)}].$$

For any $S \subseteq [n]$ with $|S| = k + l + 2$, $\mathrm{EVC}_l^k[S]$ is a nonzero, multilinear, and homogeneous polynomial of total degree $l + 1$, and every multilinear monomial of degree $l + 1$ in $x_{i_1}, \ldots, x_{i_{k+l+2}}$ appears with a nonzero coefficient. More specifically, for $S = \{i_1, \ldots, i_{k+l+2}\} \subseteq [n]$ with $i_1 < i_2 < \ldots, i_{k+l+2}$,

$$\mathrm{EVC}_l^k[S] = \sum_{\substack{K \sqcup L = S \\ |L| = l+1}} \gamma_{K,L} \cdot \prod_{i \in L} x_i, \tag{3.3}$$

where

$$\gamma_{K,L} \doteq (-1)^{\mathrm{XInv}(K,L)} \cdot \det(A_K) \cdot \det(A_L). \tag{3.4}$$

Proof. All assertions follow from elementary properties of determinants, that Vandermonde determinants are nonzero unless they have duplicate rows, and the following computation. The coefficient $\gamma_{K,L}$ can be obtained by plugging in 0 for x_i with $i \in K$, and 1 for x_i with $i \in L$. For K^* consisting of the first $k + 1$ elements of S and L^* of the last $l + 1$, this yields the determinant

$$\begin{vmatrix} a_{i_1}^k & \cdots & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ a_{i_{k+1}}^k & \cdots & 1 & 0 & \cdots & 0 \\ * & \cdots & * & a_{i_{k+2}}^l & \cdots & 1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ * & \cdots & * & a_{i_{k+l+2}}^l & \cdots & 1 \end{vmatrix}, \tag{3.5}$$

which equals the product of the Vandermonde matrices $\det(A_{K^*})$ and $\det(A_{L^*})$, and confirms the expression for γ_{K^*,L^*} as $\mathrm{XInv}(K^*, L^*) = 0$. For general K and L, we obtain a determinant with the same shape as (3.5) after rearranging the rows such that the ones involving a_i for $i \in K$ appear first and in order, and the ones involving a_i for $i \in L$ appear last and in order. By skew-symmetry, the rearrangement induces an additional factor of $\mathrm{sign}(\pi) = (-1)^{\mathrm{Inv}(\pi)}$, where π denotes the underlying permutation of $[k + l + 2]$ and $\mathrm{Inv}(\pi)$ denotes the number of inversions of π, which equals $\mathrm{XInv}(K, L)$. □

Lemma 3.1 shows that the polynomials $\mathrm{EVC}_l^k[i_1, \ldots, i_{k+l+2}]$ belong to the vanishing ideal of RFE_l^k. In fact, various subsets of them *generate* the vanishing ideal. To prove that a certain subset does so, we establish the following two steps:

1. Modulo the ideal I generated by the subset, every polynomial p is equal to a polynomial r with a particular combinatorial structure (Lemma 3.7).

2. Every nonzero polynomial r with that structure is hit by RFE_l^k (Lemma 3.8).

Together, these show that every polynomial in the vanishing ideal of RFE_l^k is equal to the zero polynomial modulo the ideal I. We conclude that the ideals coincide, i. e., the vanishing ideal of RFE_l^k is generated by the subset of instantiations of EVC_l^k that define I.

For the subset of instantiations $\mathrm{EVC}_l^k[C \sqcup L]$ where $C \subseteq [n]$ is any fixed subset of size $k+1$ and $L \subseteq [n]$ ranges over all subsets of size $l+1$ disjoint from C, the combinatorial structure bridging the two steps is that the polynomial is cored.

Definition 3.5 (monomial support and cored polynomial)**.** The *support* of a monomial $m \in \mathbb{F}[x_1, \ldots, x_n]$, denoted $\mathrm{supp}(m)$, is the set of indices $i \in [n]$ such that m depends on x_i. For $c, t \in \mathbb{N}$, a polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ is said to be (c,t)-*cored* if there exists $C \subseteq [n]$, called the *core*, such that $|C| \leq c$ and for every monomial m of p, $|\mathrm{supp}(m) \setminus C| \leq t$. For any subset $C \subseteq [n]$ and monomial $m = \prod_{i \in [n]} x_i^{d_i}$, we call $\prod_{i \in C} x_i^{d_i}$ the C-part of m, and $\prod_{i \in [n] \setminus C} x_i^{d_i}$ the non-C-part of m.

The crux for the first step is the following property, which allows us to gradually get closer to a $(k+1, l)$-cored polynomial.

Proposition 3.6. *Let $k, l, n \in \mathbb{N}$, let C be a $(k+1)$-subset of $[n]$, and let I denote the ideal generated by the polynomials $\mathrm{EVC}_l^k[C \sqcup L]$ where L ranges over all $(l+1)$-subsets of $[n] \setminus C$. Consider a monomial $m \in \mathbb{F}[x_1, \ldots, x_n]$ such that $|\mathrm{supp}(m) \setminus C| > l$. Modulo I, m is equal to a linear combination of monomials whose non-C-parts have lower degree than the non-C-part of m.*

Proof. Let L be a subset of $\mathrm{supp}(m) \setminus C$ of size $l+1$. Let m' be the monomial such that $m = m' \cdot x^L$, where $x^L \doteq \prod_{i \in L} x_i$. By Proposition 3.4, x^L is a monomial of $\mathrm{EVC}_l^k[C \sqcup L]$, and every other monomial of $\mathrm{EVC}_l^k[C \sqcup L]$ has non-C-part of degree at most l. It follows that $m' \cdot \mathrm{EVC}_l^k[C \sqcup L]$ can be written as $c \cdot m + r$, where c in a nonzero element in $\mathbb{F}$ and every monomial in r has non-C-part of lower degree than m does. Since ideals are closed under multiplication by any other polynomial, $m' \cdot \mathrm{EVC}_l^k[C \sqcup L] \in I$. Thus, we have $0 \equiv c \cdot m + r \bmod I$, which can be rewritten as $m \equiv -c^{-1} \cdot r \bmod I$. ◻

Proposition 3.6 leads to the following formalization of the first step of our approach.

Lemma 3.7. *Let $k, l, n \in \mathbb{N}$, let C be a $(k+1)$-subset of $[n]$, and let I be the ideal generated by the polynomials $\mathrm{EVC}_l^k[C \sqcup L]$ where L ranges over all $(l+1)$-subsets of $[n] \setminus C$. Modulo I, every polynomial is equal to a $(k+1, l)$-cored polynomial with core C.*

Proof. For any polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$, Proposition 3.6 allows us to systematically eliminate any monomial m in p that violates the $(k+1, l)$-cored condition, without changing p modulo I. The process may introduce other monomials, but those monomials all have non-C-parts of degree lower than m does. This means that the process cannot continue indefinitely. When it ends, the remaining polynomial is $(k+1, l)$-cored with core C and is equivalent to p modulo I. ◻

The second step of our approach is formalized in Lemma 3.8.

Lemma 3.8. *Let $k, l, n \in \mathbb{N}$ and let $r \in \mathbb{F}[x_1, \ldots, x_n]$ be nonzero and $(k+1, l)$-cored. Then RFE_l^k hits r.*

We prove Lemma 3.8 from the Zoom Lemma in Section 4. Assuming it, we have all ingredients for the proof of Theorem 1.3.

Proof of Theorem 1.3. The combination of Lemma 3.1, Lemma 3.7, and Lemma 3.8 shows that, for every core $C \subseteq [n]$ of size $k+1$, the vanishing ideal $\mathrm{Van}[\mathrm{RFE}_l^k]$ is generated by the polynomials $\mathrm{EVC}_l^k[C \sqcup L]$ where L ranges over all $(l+1)$-subsets of $[n] \setminus C$. The generators are all homogeneous of minimum degree $l+1$, and each generator has a monomial that occurs in none of the other generators (namely the product of the variables in L). Therefore, the generating set has minimum size since it forms a vector space basis of the degree-$(l+1)$ part of $\mathrm{Van}[\mathrm{RFE}_l^k]$. □

As an aside, we justify along the same lines the claim from Section 1 that all linear dependencies among instances of EVC_l^k are generated by the equations (1.6) when $\{i_1, \ldots, i_{k+l+3}\}$ ranges over all subsets of $[n]$ containing a core C of size $k+1$. A similar reduction strategy modulo those equations allows us to rewrite

$$\sum_{\substack{S \subseteq [n] \\ |S| = k+l+2}} c_S \cdot \mathrm{EVC}_l^k[S] = 0$$

such that the range of the subsets S is reduced to a $(k+1, l+1)$-cored subclass with core C. By linear independence, the only equation of that form is the trivial one with all $c_S = 0$.

Gröbner basis. We end this section with a short discussion on Gröbner bases. This part is not essential for understanding the remainder of the paper; the reader may feel free to skip it. Readers who want to know more may refer to [1, 12].

Gröbner bases are useful for solving several computational problems involving ideals, including determining whether a given polynomial p belongs to an ideal I given by a finite set G of generators. The setup presumes a total order $\geq$ on monomials with the following properties:

- For all monomials m, we have $m \geq 1$, where 1 denotes the empty monomial.
- For monomials m_1, m_2, m, we have that if $m_1 \geq m_2$, then $m_1 \cdot m \geq m_2 \cdot m$.

Assuming such a monomial ordering $\geq$, every nonzero polynomial has a unique monomial that is maximal in $\geq$, which we call the leading monomial.

For a given a polynomial p, we can compute a G-reduced form of p by repeatedly applying the following reduction step, starting from $f = p$: Find $g \in G$ such that the leading monomial of g divides some monomial m of f, and then subtract a suitable multiple of g from f so as to produce a new value of f that does not contain m as a monomial. If several such g and m exist, pick any. The process continues until no suitable g and m can be found, which the properties of the ordering $\geq$ guarantee to happen at some point. The final f is called a G-reduced form of p, which may or may not be unique.

A natural algorithm to determine membership of p in I is to compute a G-reduced form f of p and conclude that $p \in I$ if and only if $f = 0$. The algorithm generalizes computing the remainder in univariate polynomial division, with some differences being that there are multiple choices of g, and the monomial m does not need to be the leading monomial of f. A positive conclusion, $p \in I$, is always correct because reduction does not affect membership and 0 trivially belongs to the ideal. However, the algorithm can have false negatives, namely when $p \in I$ has a nonzero G-reduced form, i. e., the reduction process reaches some $f \in I$ that does not have a monomial m divisible by the leading monomial of some $g \in G$.

A *Gröbner basis* G for I is a finite set of generators satisfying the additional constraint that every nonzero element of I has a monomial m that is divisible by the leading monomial of some element of G. In this case, the above algorithm for deciding membership in I is always correct. In fact, this gives another characterization of when a finite generating set G is a Gröbner basis. Yet another characterization is that every polynomial p has a unique G-reduced form f.

In the overview, we claimed that the set G of polynomials $\mathrm{EVC}_l^k[C \sqcup L]$ form a Gröbner basis for $\mathrm{Van}[\mathrm{RFE}_l^k]$, where $C \subseteq [n]$ is a fixed core of size $k+1$ and L ranges over the $(l+1)$-subsets of $[n] \setminus C$. This holds with respect to any monomial ordering such that, for every L, $x^L \doteq \prod_{i \in L} x_i$ is the leading monomial of $\mathrm{EVC}_l^k[C \sqcup L]$. Examples of such orderings include all lexicographic orderings where the variables outside C have higher priority than the variables inside C. Lemma 3.8 implies that every nonzero polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$ has a monomial with more than l variables outside of C, which is to say that the monomial is divisible by x^L for some $L \subseteq [n] \setminus C$ of size $l+1$. As x^L is the leading monomial of $\mathrm{EVC}_l^k[C \sqcup L]$, we conclude that every nonzero polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$ has a monomial that is divisible by the leading term of some element of G, i. e., G is a Gröbner basis.

One can interpret Proposition 3.6 as performing a reduction step of p by G. Lemma 3.7 keeps performing this step until it is no longer possible, yielding a G-reduced form f of p that is $(k+1, l)$-cored with core C. Lemma 3.8 implies that any two $(k+1, l)$-cored representatives modulo I of the same polynomial p coincide, so every polynomial p has a unique G-reduced form. This is another way to see that the set G is a Gröbner basis.

Instantiation for SV. By the connection between SV and RFE, a generating set for $\mathrm{Van}[\mathrm{RFE}_l^{l-1}]$ induces a generating set for $\mathrm{Van}[\mathrm{SV}^l]$. We provide an explicit expression as an instantiation of Theorem 1.3 and Proposition 3.4.

Corollary 3.9. *Let $l, n \in \mathbb{N}$ and let a_i for $i \in [n]$ be distinct elements of $\mathbb{F}$. For any fixed set $C \subseteq [n]$ of size l, the polynomials $\mathrm{EVCSV}^l[C \sqcup L]$ form a generating set of minimum size for $\mathrm{Van}[\mathrm{SV}^l]$ when L ranges over all $(l+1)$-subsets of $[n]$ that are disjoint from C. Here, for any $S = \{i_1, \dots, i_{2l+1}\} \subseteq [n]$ with $i_1 < \cdots < i_{2l+1}$,*

$$\mathrm{EVCSV}^l[S] \doteq \sum_{\substack{T \subseteq S \\ |T| = l+1}} \gamma'_{S \setminus T, T} \cdot \prod_{i \in T} x_i,$$

where

$$\gamma'_{S \setminus T, T} \doteq (-1)^{\mathrm{XInv}(S \setminus T, T)} \cdot \left(\prod_{i \in T} \prod_{j \in [n] \setminus \{i\}} (a_i - a_j) \right) \cdot \det(A_{S \setminus T}) \cdot \det(A_T). \tag{3.6}$$

Proof. By Proposition 2.2, for any polynomial $p \in \mathbb{F}[x_1, \dots, x_n]$, $p(\mathrm{SV}^l) = 0$ iff $(p \circ A)(\mathrm{RFE}_l^{l-1}) = 0$, where $A : \mathbb{F}^n \to \mathbb{F}^n$ is the invertible transformation that divides each variable x_i by $\prod_{j \in [n] \setminus \{i\}} (a_i - a_j)$. Since the vanishing ideal of RFE_l^{l-1} coincides with the ideal generated by the polynomials $\mathrm{EVC}_l^{l-1}[C \sqcup L]$, it follows that the vanishing ideal of SV^l coincides with the ideal generated by the polynomials

$$\mathrm{EVC}_l^{l-1}[C \sqcup L] \circ A^{-1}. \tag{3.7}$$

For $T \subseteq S \doteq C \sqcup L$ with $|T| = l+1$, the coefficient of $\prod_{i \in T} x_i$ in $\mathrm{EVC}_l^{l-1}[S]$ is given by $\gamma_{S \setminus T, T}$. By the construction of A, A^{-1} multiplies x_i by $\prod_{j \in [n] \setminus \{i\}} (a_i - a_j)$. Thus, the coefficient of $\prod_{i \in T} x_i$ in (3.7) equals $\gamma'_{S \setminus T, T}$ given by (3.6) and $\mathrm{EVC}_l^{l-1}[C \sqcup L] \circ A^{-1} = \mathrm{EVCSV}^l[S]$. □

In comparison to the coefficients $\gamma_{S \setminus T, T}$ of the polynomial $\mathrm{EVC}_l^{l-1}[S]$, the coefficients $\gamma'_{S \setminus T, T}$ of $\mathrm{EVCSV}^l[S]$ contain an additional term, namely the middle term on the right-hand side of (3.6). As a consequence, each coefficient of $\mathrm{EVCSV}^l[S]$ depends on *all* abscissas $a_1, \ldots, a_n$, whereas the coefficients of $\mathrm{EVC}_l^k[S]$ only depend on the abscissas with indices in S. This reflects a difference in setup between the two generators: The substitution for a variable x_i is a multivariate function of all abscissas in SV versus a univariate function of the abscissa a_i only in RFE. The difference represents one reason why RFE is more convenient to work with than SV, even though both have essentially the same power.

A more important reason is our derivation of the generating set EVC_l^k in Lemma 3.1. Our approach hinges on the fact that the substitutions for a variable x_i induce linear equations involving the seed variables $g_k, \ldots, g_0, h_l, \ldots, h_0$, with coefficients being expressions in terms of the polynomial variables $x_1, \ldots, x_n$ and abscissas $a_1, \ldots, a_n$. Collecting $k + l + 2$ of such equations yields as many linear constraints as unknowns, which suffices to derive a nontrivial element of the vanishing ideal. The substitutions (1.1) for x_i made by SV^l similarly induce linear equations, though not between the mere seed variables $y_1, z_1, \ldots, y_l, z_l$ but between monomials in the seed variables, namely the constant monomial and the monomials $z_t y_t^d$ for $t \in [l]$ and $d \in \{0, \ldots, n-1\}$. In contrast to the setting of RFE, even if we collect all of those equations, namely n linear equations in $nl + 1$ unknowns, this does not give us enough information to derive a nontrivial element of the vanishing ideal.

4 Zoom Lemma

Throughout the paper we make repeated use of a key technical tool, the Zoom Lemma. The lemma allows us to zoom in on the contributions of the monomials in a polynomial p that have prescribed degrees in a subset of the variables. We introduce the following terminology for prescribing degrees.

Definition 4.1 (degree pattern)**.** Let $J \subseteq [n]$. A *degree pattern* with domain J is a J-indexed tuple $d \in \mathbb{N}^J$ of nonnegative integers. A degree pattern d *matches* a monomial $m \in \mathbb{F}[x_1, \ldots, x_n]$ if, for every $j \in J$, m has degree exactly d_j in x_j. We say that d is *in* p if d matches some monomial in p.

For any fixed J, every polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ can be written uniquely in the form

$$p = \sum_{d \in \mathbb{N}^J} p_d \cdot x^d$$

where $x^d \doteq \prod_{j\in J} x_j^{d_j}$ and p_d depends only on variables not indexed by J. We refer to p_d as the coefficient of x^d in p.

The notation p_d can be viewed as a generalization of the common one for the coefficient of degree d of a univariate polynomial p.

Our technique allows us to zoom in on the contributions of the coefficients p_d of degree patterns d that satisfy the following additional constraint.

Definition 4.2 (extremal degree pattern)**.** Let $K, L \subseteq [n]$. A degree pattern $d^* \in \mathbb{N}^{K\cup L}$ is (K, L)-*extremal* in a polynomial $p \in \mathbb{F}[x_1, \dots, x_n]$ if d^* is the unique degree pattern $d \in \mathbb{N}^{K\cup L}$ in p that satisfies both

(i) $d_j \leq d^*_j$ for all $j \in K$, and

(ii) $d_j \geq d^*_j$ for all $j \in L$.

The notion of extremality in Definition 4.2 is closely related to standard notions of minimality and maximality of tuples of numbers. A J-tuple d^* is minimal in a set D of such tuples if the only tuple $d \in D$ that satisfies $d_j \leq d^*_j$ for all $j \in J$, is d^* itself. A maximal tuple is defined similarly by replacing $\leq$ by $\geq$. Minimality is equivalently $(J, \emptyset)$-extremality, and maximality is equivalently $(\emptyset, J)$-extremality.

When K and L intersect, note that only degree patterns $d \in \mathbb{N}^{K\cup L}$ with $d_j = d^*_j$ for all $j \in K \cap L$ affect whether d^* is (K, L)-extremal.

The above terminology lets us state our key technical lemma succinctly.

Lemma 4.3 (Zoom Lemma)**.** *Let $K, L \subseteq [n]$, let $p \in \mathbb{F}[x_1, \dots, x_n]$, and let $d^* \in \mathbb{N}^{K\cup L}$ be a degree pattern that is (K, L)-extremal in p. If the coefficient p_{d^*} is nonzero upon the substitution*

$$x_i \leftarrow z \cdot \frac{\prod_{j\in K\setminus L}(a_i - a_j)}{\prod_{j\in L\setminus K}(a_i - a_j)} \qquad \forall i \in [n] \setminus (K \cup L) \tag{4.1}$$

where z is a fresh variable, then RFE^k_l hits p for any $k \geq |K|$ and $l \geq |L|$.

Note that the result of substituting (4.1) into p_{d^*} is a univariate polynomial q in z. In the case where p is homogeneous, q has a single monomial, so q is nonzero iff q is nonzero at $z = 1$. In general, it suffices for q to be nonzero at some point $z \in \mathbb{F}$. As for the conclusion, the most interesting settings in Lemma 4.3 are $k = |K|$ and $l = |L|$. This is because the range of RFE^k_l is contained in the range of $\mathrm{RFE}^{k'}_{l'}$ for $k' \geq k$ and $l' \geq l$. Also, whereas many or our instantiations of the Zoom Lemma have K and L disjoint, this is not necessary for the lemma to hold.[1]

Let us first see how the Zoom Lemma allows us to complete the proof of Theorem 1.3. There are several ways to do so; we present a fairly generic way.

[1]In fact, allowing K and L to overlap is useful in Section 5 (see Proposition 5.10) and Section 8 (see Proposition 8.8).

Proof of Lemma 3.8 from the Zoom Lemma. Let $C \subseteq [n]$ denote a core of size at most $k+1$ for r. We construct subsets $K, L \subseteq [n]$ with $|K| \leq k$ and $|L| \leq l$, and a degree pattern d^* with domain $K \cup L$ that is (K, L)-extremal in r such that r_{d^*} is nonzero upon the substitution (4.1). The Zoom Lemma then implies that RFE_l^k hits r.

The construction consists of two steps. First, we pick $i^* \in C$ arbitrarily. (We can assume without loss of generality that C is nonempty because if C is a core, then so is C with an additional element.) We also set $K \doteq C \setminus \{i^*\}$, and let d_+ be a degree pattern with domain K that matches a monomial in r and that is minimal among all such degree patterns. The existence of d_+ follows from the fact that r is nonzero. By construction, $|K| \leq (k+1) - 1 = k$ and r_{d_+} is nonzero.

Second, we pick a degree pattern d_- with domain $[n] \setminus C$ that matches a monomial in r_{d_+} and that is maximal among all such degree patterns. The existence of d_- follows from the fact that r_{d_+} is nonzero. Let L denote the set of indices $j \in [n] \setminus C$ on which d_- is positive. The hypothesis that C is a $(k+1, l)$-core for r implies that $|L| \leq l$. By construction, the restriction of d_- to the domain L is maximal among the degree patterns with domain L in r_{d_+}.

Note that K and L are disjoint, because $K \subseteq C$ and $L \subseteq [n] \setminus C$. We define d^* as the degree pattern with domain $K \sqcup L$ that agrees with d_+ on K and with d_- on L. The minimality and maximality properties of d_+ and d_- imply that d^* is (K, L)-extremal in r. As there is at least one monomial in r that agrees with the degree pattern d^*, the coefficient r_{d^*} is nonzero. Since K includes all of C but i^*, r_{d^*} cannot depend on variables indexed by C other than x_{i^*}. By the maximality of d_- on $[n] \setminus C$ and the fact that L contains all indices in $[n] \setminus C$ on which d_- is positive, r_{d^*} cannot depend on any variable in $[n] \setminus C$. Thus, r_{d^*} is a nonzero polynomial that depends only on x_{i^*}. It follows that substituting (4.1) into r_{d^*} yields a nonzero polynomial in z. □

Before giving a formal proof of the Zoom Lemma, we provide some intuition for the mechanism behind it, and we explain how the choice of the substitution (4.1) and the extremality requirement arise. We consider $k = |K|$ and $l = |L|$, and focus on the setting of homogeneous polynomials p, in which case we can set $z = 1$ without loss of generality.

We start with the special case where (i) $\ell = 0$, or equivalently $L = \emptyset$, and (ii) the degree pattern $d^* \in \mathbb{N}^K$ is zero in every coordinate, so x^{d^*} is the constant monomial 1. We can zoom in on p_{d^*} by setting all variables x_j for $j \in K$ to zero. The generator RFE_0^k allows us to do so by picking a seed f such that $f(a_j) = 0$ for all $j \in K$, namely

$$f(\alpha) \doteq \prod_{j \in K} (\alpha - a_j). \tag{4.2}$$

The evaluation of p at $\mathrm{RFE}(f)$ coincides with the evaluation of p_{d^*} at $\mathrm{RFE}(f)$, which is precisely (4.1) with $z = 1$. If the evaluation is nonzero, then evidently RFE_0^k hits p, as desired.

In order to handle more general degree patterns $d^* \in \mathbb{N}^K$, we introduce a fresh parameter ξ_j for each $j \in K$, and replace a_j in (4.2) by $a_j - \xi_j$, i. e., we consider the seeds

$$\hat{f}(\alpha) \doteq \prod_{j \in K} (\alpha - a_j + \xi_j), \tag{4.3}$$

where the hat indicates a dependency on the fresh parameters. For each i, $\hat{f}(a_i)$ is a multivariate polynomial in ξ_j, $j \in K$, and $\mathrm{RFE}(\hat{f})$ applies the substitution $x_i \leftarrow \hat{f}(a_i)$ for each $i \in [n]$. The critical property is that $\hat{f}(a_i)$ contains the factor ξ_i for $i \in K$ but not for $i \notin K$. More precisely,

$$\hat{f}(a_i) = \begin{cases} \hat{c}_i \cdot \xi_i & i \in K \\ \hat{c}_i & i \notin K \end{cases},$$

where $\hat{c}_i \doteq \prod_{j \in K \setminus \{i\}} (a_i - a_j + \xi_j)$ is a multivariate polynomial in the parameters ξ with nonzero constant term, namely $c_i \doteq \prod_{j \in K \setminus \{i\}} (a_i - a_j)$.

For any monomial m with matching degree pattern $d \in \mathbb{N}^K$, we have

$$m(\mathrm{RFE}(\hat{f})) = m(\hat{c}) \cdot \xi^d = m_d(\hat{c}) \cdot \hat{c}^d \cdot \xi^d.$$

Here we see that, when $m(\mathrm{RFE}(\hat{f}))$ is expanded as a linear combination of monomials in the ξ_j, the combination contains only monomials divisible by ξ^d and the coefficient of ξ^d is nonzero (namely c^d).

In the expansion of $p(\mathrm{RFE}(\hat{f}))$, the coefficient of ξ^{d^*}

(a) has a contribution $m(c) = m_{d^*}(c) \cdot c^{d^*}$ from each monomial m in p that matches d^*, and

(b) may have contributions from other monomials m in p but only from those whose degree pattern on K is smaller than d^*, i. e., only if $\deg_j(m) \le d_j^*$ for all $j \in K$.

By adding the contributions of all monomials m with degree pattern d^* we obtain

$$p_{d^*}(\mathrm{RFE}(\hat{f})) \cdot \mathrm{RFE}(\hat{f})^{d^*} = p_{d^*}(\hat{c}) \cdot \hat{c}^{d^*} \cdot \xi^{d^*}.$$

By properties (a) and (b) above, we conclude that the coefficient of the monomial ξ^{d^*} in $p(\mathrm{RFE}(\hat{f}))$:

(a′) has a contribution of $p_{d^*}(c) \cdot c^{d^*}$ from the monomials matching d^*, and

(b′) cannot have any additional contributions provided that there are no degree patterns on K in p that are smaller than d^*.

For a degree pattern d^* in p, condition (b′) can be formulated as the minimality of d^* among the degree patterns on K in p, which is exactly the requirement that d^* is (K, L)-extremal in p for $L = \emptyset$. Under this condition we conclude that the coefficient of the monomial ξ^{d^*} in $p(\mathrm{RFE}(\hat{f}))$ equals $p_{d^*}(c) \cdot c^{d^*}$. Note that c^{d^*} is nonzero. Since $p_{d^*}(c)$ only depends on the components c_i for $i \in [n] \setminus K$, and those components agree with (4.1) for $z = 1$, the coefficient of the monomial ξ^{d^*} in $p(\mathrm{RFE}(\hat{f}))$ is nonzero if and only if p_{d^*} is nonzero at the point (4.1) with $z = 1$. Thus, for a homogeneous polynomial p, the hypotheses of the lemma imply that $p(\mathrm{RFE}(\hat{f}))$ is a nonzero polynomial in the parameters ξ. It follows that a random setting of the parameters ξ yields a seed f' for RFE_0^k such that $p(\mathrm{RFE}(f'))$ is nonzero. This shows that RFE_0^k hits p.

The symmetric case $k = 0$ can be obtained from the case $l = 0$ by transforming $x_i \mapsto x_i^{-1}$ for each $i \in [n]$. The transformation maps a seed f for RFE_l^0 into a seed $\tilde{f}$ for RFE_0^l, wherein the zeroes of $\tilde{f}$ come from the poles of f. Given a polynomial $p(x_1, \ldots, x_n)$, we similarly transform the variables and clear denominators to obtain the polynomial $\tilde{p}(x_1, \ldots, x_n) \doteq p(x_1^{-1}, \ldots, x_n^{-1}) \cdot x^g$, where g is any degree pattern with domain $[n]$ for which g_i is at least the degree of x_i in p for every $i \in [n]$. We apply the previous case of the Zoom Lemma to $\tilde{p}$ and obtain the new case of the Zoom Lemma for p. Note that a monomial with degree pattern $\tilde{d}$ in $\tilde{p}$ corresponds to a monomial with degree pattern $d = g - \tilde{d}$ in p. It follows that $\widetilde{d^*}$ is minimal in $\tilde{p}$ iff d^* is maximal in p, which is exactly the (K, L)-extremality requirement of the Zoom Lemma in the case where $K = \emptyset$.

The above arguments for the special cases $l = 0$ and $k = 0$ carry through for arbitrary polynomials p under the assumption that p_{d^*} is nonzero upon the substitution (4.1) with $z = 1$, i. e., that the univariate polynomial $q(z)$ obtained by substituting (4.1) in p_{d^*} is nonzero at $z = 1$. The homogeneity of p was only used to conclude that if q is nonzero, then q is nonzero at $z = 1$. To handle polynomials p where q may be nonzero but zero at $z = 1$, we run the above argument with an arbitrary value of $z \in \mathbb{F}$ where q is nonzero. We can do so by including an additional factor of z on the right-hand sides of (4.2) and (4.3), i. e., by considering $f(\alpha) \doteq z \cdot \prod_{j \in K}(\alpha - a_j)$ and $\hat{f}(\alpha) \doteq z \cdot \prod_{j \in K}(\alpha - a_j + \xi_j)$, respectively. Both expressions correspond to valid seeds for RFE_0^k in the roots parametrization.

The case for general k and l follows in a similar fashion, introducing parameters for the zeroes as well as the poles of the seed f, considering the monomial in those parameters with degree pattern determined by d^*, and clearing denominators.

Proof of Zoom Lemma. Let K, L, p, and d^* be as in the lemma statement. Fix z to a value in $\mathbb{F}$ such that p_{d^*} is nonzero upon the substitution (4.1). Such a value exists by the hypothesis of the lemma (for large enough $\mathbb{F}$). Since the range of RFE_l^k is contained in the range of $\mathrm{RFE}_{l'}^{k'}$ for $k' \geq k$ and $l' \geq l$, it suffices to show that RFE_l^k hits p for $k = |K|$ and $l = |L|$. Let ξ_j for each $j \in K$ and η_j for each $j \in L$ be fresh indeterminates. We denote by $\widehat{\mathbb{F}}$ the field of rational functions in those indeterminates with coefficients in $\mathbb{F}$, and by V the subset of elements that, when written in lowest terms, have denominators with nonzero constant terms. Let $\Phi : V \to \mathbb{F}$ map each element of V to the result of substituting $\xi_j \leftarrow 0$ for each $j \in K$ and $\eta_j \leftarrow 0$ for each $j \in L$. The result is always well-defined.

Define $\hat{f} \in \widehat{\mathbb{F}}(\alpha)$ as follows:

$$\hat{f}(\alpha) \doteq z \cdot \frac{\prod_{j \in K}(\alpha - a_j + \xi_j)}{\prod_{j \in L}(\alpha - a_j + \eta_j)}.$$

The substitution $\mathrm{RFE}(\hat{f})$ effects $x_i \leftarrow \hat{f}(a_i) \in \widehat{\mathbb{F}}$ for each $i \in [n]$. We claim that $p(\mathrm{RFE}(\hat{f}))$ is nonzero. This suffices to conclude that RFE_l^k hits p, because substituting ξ_j and η_j by a random scalar from $\mathbb{F}$ transforms $\hat{f}$ into a seed f' such that, with high probability, f' is a valid seed for RFE_l^k and $p(\mathrm{RFE}(f')) \neq 0$. Henceforth we show that $p(\mathrm{RFE}(\hat{f})) \neq 0$.

For each $i \in [n]$, there exists $\hat{c}_i \in V$ with $\Phi(\hat{c}_i) \neq 0$ such that

$$\hat{f}(a_i) = \begin{cases} \hat{c}_i \cdot \frac{\xi_i}{\eta_i} & i \in K \cap L \\ \hat{c}_i \cdot \xi_i & i \in K \setminus L \\ \hat{c}_i \cdot \frac{1}{\eta_i} & i \in L \setminus K \\ \hat{c}_i & i \notin K \cup L \end{cases}, \tag{4.4}$$

namely

$$\hat{c}_i = z \cdot \frac{\prod_{j \in K \setminus \{i\}} (a_i - a_j + \xi_j)}{\prod_{j \in L \setminus \{i\}} (a_i - a_j + \eta_j)}.$$

For $i \notin K \cup L$, $\Phi(\hat{c}_i)$ is moreover the value substituted into x_i by (4.1).

Let D denote the set of all degree patterns $d \in \mathbb{N}^{K \cup L}$ that match a monomial in p. We have that

$$p = \sum_{d \in D} p_d \cdot x^d. \tag{4.5}$$

For $d \in D$, define $\hat{q}_d$ to be the result of substituting $x_i \leftarrow \hat{c}_i$ into p_d for each $i \in [n]$.

Combining (4.4) and (4.5), we obtain

$$p(\mathrm{RFE}(\hat{f})) = \sum_{d \in D} \hat{q}_d \cdot \hat{c}^d \cdot \frac{\xi^{d|_K}}{\eta^{d|_L}}, \tag{4.6}$$

where $d|_K$ and $d|_L$ respectively are the restrictions of d onto the domains K and L respectively. Fix any function $\psi : [k+l] \to K \cup L$ such that ψ establishes a bijection between $\{1, \ldots, k\}$ and K and establishes a bijection between $\{k+1, \ldots, k+l\}$ and L. For $j \in \{1, \ldots, k\}$, let ζ_j be an alias for $\xi_{\psi(j)}$, and for $j \in \{k+1, \ldots, k+l\}$, let ζ_j be an alias for $\eta_{\psi(j)}$. For each $d \in \mathbb{N}^{K \cup L}$, define a corresponding $\delta \in \mathbb{Z}^{k+l}$ given by $\delta_j = d_{\psi(j)}$ for $j \in \{1, \ldots, k\}$ and $\delta_j = -d_{\psi(j)}$ for $j \in \{k+1, \ldots, k+l\}$. Let $\Delta \subseteq \mathbb{Z}^{k+l}$ consist of the δ corresponding to each $d \in D$. Finally, for each $d \in D$ with corresponding $\delta \in \Delta$, define $\hat{c}_\delta \doteq \hat{q}_d \cdot \hat{c}^d$, capturing the first two factors in the d-th term of (4.6). Rewritten in this notation, (4.6) becomes

$$\sum_{\delta \in \Delta} \hat{c}_\delta \cdot \prod_{j=1}^{k+l} \zeta_j^{\delta_j}. \tag{4.7}$$

Our hypothesis that d^* is (K, L)-extremal in p says that the only $d \in D$ such that $d_j \leq d_j^*$ for every $j \in K$ and $d_j \geq d_j^*$ for every $j \in L$, is $d = d^*$. Translated into a condition on the element $\delta^* \in \Delta$ corresponding to d^*, the hypothesis says that δ^* is minimal in Δ. Our other hypothesis states that p_{d^*} does not vanish upon substituting (4.1). As (4.1) equates to substituting $x_i \leftarrow \Phi(\hat{c}_i)$ for $i \notin K \cup L$, this hypothesis equivalently states that $\Phi(\hat{q}_{d^*})$ is nonzero. Since for each $j \in K \cup L$ we have $\Phi(\hat{c}_j) \neq 0$, we conclude that $\Phi(\hat{c}_{\delta^*}) \neq 0$. That $p(\mathrm{RFE}(\hat{f}))$ is nonzero now follows from the next proposition. □

Proposition 4.4. *Let $\widehat{\mathbb{F}} = \mathbb{F}(\zeta_1, \ldots, \zeta_r)$ be the field of rational functions in indeterminates $\zeta_1, \ldots, \zeta_r$, let $V \subseteq \widehat{\mathbb{F}}$ consist of the rational functions whose denominator has nonzero constant term, and let $\Phi : V \to \mathbb{F}$ be the function that maps each rational function in V to its value after substituting $\zeta_j \leftarrow 0$ for all $j \in [r]$. Let*

$$s = \sum_{\delta \in \Delta} \hat{c}_\delta \cdot \prod_{j=1}^{r} \zeta_j^{\delta_j}$$

where $\Delta \subseteq \mathbb{Z}^r$ is some finite set, and we have $\hat{c}_\delta \in V$ for every $\delta \in \Delta$. If there exists $\delta^* \in \Delta$ that is minimal in Δ and for which $\Phi(\hat{c}_{\delta^*}) \neq 0$, then $s \neq 0$.

Proof. By clearing denominators, we may assume without loss of generality that, for every $\delta \in \Delta$ and every $j \in [r]$, $\delta_j \geq 0$, and that, for every $\delta \in \Delta$, $\hat{c}_\delta$ is a *polynomial* in $\zeta_1, \ldots, \zeta_r$. In this case, all quantities in the sum for s are polynomials in $\zeta_1, \ldots, \zeta_r$. The minimality hypothesis on δ^* implies that the coefficient of $\prod_{j=1}^{r} \zeta_j^{\delta_j^*}$ in the monomial expansion of s is precisely the constant coefficient of $\hat{c}_{\delta^*}$, and the hypothesis $\Phi(\hat{c}_{\delta^*}) \neq 0$ asserts that this coefficient is nonzero. □

5 Membership Test

In this section we develop the structured membership test for the vanishing ideal $\mathrm{Van}[\mathrm{RFE}_l^k]$ given in Theorem 1.8. We begin with some basic results regarding membership to $\mathrm{Van}[\mathrm{RFE}_l^k]$ and then develop a criterion for multilinear polynomials.

Basic properties. It is well-known that $\mathrm{Van}[\mathrm{SV}^l]$ does not contain any polynomial with a monomial of support at most l, i. e., a monomial involving at most l variables. We generalize the lower bound on the support to $\mathrm{Van}[\mathrm{RFE}_l^k]$ and also establish an upper bound in the case of multilinear polynomials. Note that for multilinear monomials support conditions translate into degree conditions.

Proposition 5.1. *If a polynomial p contains a monomial of support at most l, then RFE_l^k hits p. If a multilinear polynomial p in the variables $x_1, \ldots, x_n$ contains a monomial of support at least $n-k$, then RFE_l^k hits p.*

The known proofs of the first part for SV^l make use of partial derivatives. We establish the generalization for RFE_l^k using our generating set in Theorem 1.3, whose analysis hinges on the Zoom Lemma. A similar argument works for the second part, but we opt to establish it via a black-box reduction to the first part for multilinear polynomials. The approach illustrates the utility of our generalization of SV^l to RFE_l^k since even for SV^l we need to consider settings of the parameters k and l other than $k = l - 1$.

Proof. For the first part, by Proposition 3.4 none of the polynomials EVC_l^k contain a monomial of support l or less. The same holds for the nonzero polynomials in the ideal generated by these polynomials, which by Theorem 1.3 equals $\mathrm{Van}[\mathrm{RFE}_l^k]$. Thus, every polynomial that has a monomial of support at most l, is hit by RFE_l^k.

For the second part, consider $q(x_1,\dots,x_n) \doteq x_1 \cdot \dots \cdot x_n \cdot p(1/x_1,\dots,1/x_n)$. Note that if p is a multilinear polynomial in the variables $x_1,\dots,x_n$, then so is q. If a multilinear p has a monomial of support at least $n-k$, then q has a monomial of support at most k. By the first part, RFE^l_k hits q. Since the mapping $x_i \leftarrow 1/x_i$ transforms RFE^l_k into RFE^k_l, we conclude that RFE^k_l hits p. □

Another feature of SV that generalizes to RFE is that the generator separates the homogeneous components of a given polynomial p. The feature allows us to reduce the general case of testing membership in $\mathrm{Van}[\mathrm{RFE}^k_l]$ to the homogeneous case, as was already effectively used in the proof of the Zoom Lemma.

Proposition 5.2. *For any polynomial p, p vanishes upon substituting* RFE *if and only if every homogeneous component of p vanishes upon substituting* RFE*.*

Proof. For any seed f for RFE and any scalar z, the rescaled substitution $z \cdot \mathrm{RFE}(f)$ is in the range of RFE, namely as $\mathrm{RFE}(z \cdot f)$. It follows (provided that $\mathbb{F}$ is sufficiently large) that $p(\mathrm{RFE})$ vanishes if and only if $p(\zeta \cdot \mathrm{RFE})$ vanishes, where ζ is a fresh indeterminate. We now consider the expansion of $p(\zeta \cdot \mathrm{RFE})$ as a polynomial in ζ. With $p^{(d)}$ as the degree-d homogeneous component of p, we have

$$p(\zeta \cdot \mathrm{RFE}) = \sum_d p^{(d)}(\zeta \cdot \mathrm{RFE}) = \sum_d \zeta^d \cdot p^{(d)}(\mathrm{RFE}).$$

The coefficient of ζ^d, $p^{(d)}(\mathrm{RFE})$, has no dependence on ζ. We deduce that $p(\zeta \cdot \mathrm{RFE})$ is the zero polynomial if and only if $p^{(d)}(\mathrm{RFE})$ vanishes for every d. □

Criterion for multilinear polynomials. We now develop the full membership test for multilinear polynomials given in Theorem 1.8. Condition 2 in Theorem 1.8 is closely related to the Zoom Lemma. Note that for multilinear polynomials and disjoint K and L, $\partial_L p|_{K\leftarrow 0}$ coincides with the coefficient p_{d^*} where d^* is the degree pattern with domain $K \sqcup L$, 0 in the positions of K, and 1 in the positions of L. Moreover, since p is multilinear, the condition that d^* be (K, L)-extremal in p is automatically satisfied: The only multilinear monomial m with support in $K \sqcup L$ with $\deg_{x_i}(m) \le d^*_i = 0$ for all $i \in K$ and $\deg_{x_i}(m) \ge d^*_i = 1$ for all $i \in L$ is $m = x^{d^*}$. This leads to the following specialization of the Zoom Lemma for multilinear polynomials with disjoint K and L.

Lemma 5.3 (Zoom Lemma for multilinear polynomials)**.** *Let $K, L \subseteq [n]$ be disjoint, and let $p \in \mathbb{F}[x_1,\dots,x_n]$ be a multilinear polynomial. If $\partial_L p|_{K\leftarrow 0}$ is nonzero upon the substitution*

$$x_i \leftarrow z \cdot \frac{\prod_{j\in K}(a_i - a_j)}{\prod_{j\in L}(a_i - a_j)} \qquad \forall i \in [n] \setminus (K \sqcup L), \tag{5.1}$$

where z is a fresh variable, then RFE^k_l hits p for any $k \ge |K|$ and $l \ge |L|$.

Observe that the substitution (5.1) in Lemma 5.3 coincides with (1.7) in Theorem 1.8. Thus, for a multilinear polynomial p, condition 2 in Theorem 1.8 expresses that there is no way to show that RFE_l^k hits p via an application of Lemma 5.3. The necessity of this condition for membership of p in $\mathrm{Van}[\mathrm{RFE}_l^k]$ is clear. The necessity of Condition 1 in Theorem 1.8 is just the special case of Proposition 5.1 for multilinear polynomials.

What remains to argue is that the combination of condition 1 and condition 2 is sufficient for membership of p in $\mathrm{Van}[\mathrm{RFE}_l^k]$. Equivalently, it remains to argue the following property for every multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ that only contain monomials of degrees between $l+1$ and $n-k-1$: If p is hit by RFE_l^k then there is an application of Lemma 5.3 that exhibits this fact. We actually prove the property for every multilinear p of degree d with $l \leq d \leq n-k$. We do so with a two-step strategy similar to one we used for the part of Theorem 1.3 that RFE_l^k hits every polynomial outside of the ideal generated by instantiations of EVC_l^k:

1. Modulo the ideal I generated by a certain subset of the instantiations of EVC_l^k, p is equal to a cored polynomial r with certain parameters. Since $p \notin \mathrm{Van}[\mathrm{RFE}_l^k]$ and $I \subseteq \mathrm{Van}[\mathrm{RFE}_l^k]$, r needs to be nonzero.

2. For every such cored polynomial r that is nonzero, we can apply Lemma 5.3 to prove that RFE_l^k hits r, i. e., condition 2 fails for r.

By linearity and the necessity of condition 2 for multilinear polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$), we conclude that the condition fails for p, as well.

The crux for the first step in the context of Theorem 1.3 is the transformation in Proposition 3.6, which gradually gets closer to a cored polynomial with the desired parameters. In general, the transformation in Proposition 3.6 does not maintain multilinearity. We show how to tweak the transformation and preserve multilinearity at the expense of an increase in the size of the core.

Proposition 5.4. *Let $k, l, n, d \in \mathbb{N}$, let C be a $(k+d-l)$-subset of $[n]$, and let I denote the ideal generated by the polynomials $\mathrm{EVC}_l^k[K \sqcup L]$ where K ranges over all $(k+1)$-subsets of C and L ranges over all $(l+1)$-subsets of $[n] \setminus K$. Consider a multilinear monomial $m \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most d such that $|\mathrm{supp}(m) \setminus C| > l$. Modulo I, m is equal to a linear combination of multilinear monomials of the same degree as m but whose non-C-parts have lower degree than the non-C-part of m.*

Proof. Consider the subset $L \subseteq \mathrm{supp}(m) \setminus C$ of size $l+1$ in the proof of Proposition 3.6, and $x^L \doteq \prod_{i \in L} x_i$. Since m is multilinear, so is $m' \doteq m/x^L$, and $|\mathrm{supp}(m')| \leq d - |L| = d - l - 1$. Provided $|C| \geq (k+1) + (d-l-1) = k+d-l$, there exists a subset $K \subseteq C$ of size $k+1$ that is disjoint from $\mathrm{supp}(m')$. We substitute $\mathrm{EVC}_l^k[C \sqcup L]$ by $\mathrm{EVC}_l^k[K \sqcup L]$ in the proof. $\mathrm{EVC}_l^k[K \sqcup L]$ is homogeneous and multilinear. By construction $K \sqcup L$ is disjoint from $\mathrm{supp}(m')$, so $\mathrm{EVC}_l^k[K \sqcup L]$ does not depend on any variables that m' depends on. It follows that $m' \cdot \mathrm{EVC}_l^k[K \sqcup L]$ is multilinear and homogeneous of the same degree as m, and so is r in the proof. ◻

Applying Proposition 5.4 repeatedly in a similar way as Proposition 3.6 in the proof of Lemma 3.7 yields the following formalization of the first step in the setting of Theorem 1.8.

Lemma 5.5. *Let $k, l, n, d \in \mathbb{N}$, let C be a $(k+d-l)$-subset of $[n]$, and let I denote the ideal generated by the polynomials $\mathrm{EVC}_l^k[K \sqcup L]$ where K ranges over all $(k+1)$-subsets of C and L ranges over all $(l+1)$-subsets of $[n] \setminus K$. Modulo I, every multilinear polynomial p of degree at most d in $\mathbb{F}[x_1, \ldots, x_n]$ is equal to a $(k+d-l, l)$-cored multilinear polynomial with core C that is either zero or else has the same degree as p.*

The following refinement of Lemma 3.8 from the context of Section 3 represents the corresponding second step in the context of Theorem 1.8. This is where the degree constraint comes into play.

Lemma 5.6. *Let $k, l, n, d \in \mathbb{N}$ with $l \leq d \leq n-k$. Let r be a nonzero multilinear polynomial of degree d in $\mathbb{F}[x_1, \ldots, x_n]$ that is $(d+k-l, l)$-cored. There are disjoint sets $K, L \subseteq [n]$ with $|K| = k$ and $|L| = l$ so that $\partial_L r|_{K \leftarrow 0}$ is nonzero upon the substitution (5.1).*

Proof. Let C denote the core of size at most $d+k-l$, and let m be a monomial of r of degree d that maximizes $|\mathrm{supp}(m) \setminus C|$. Let K be a subset of $[n] \setminus \mathrm{supp}(m)$ of size k that contains all of $C \setminus \mathrm{supp}(m)$. Such a set K exists because $|C \setminus \mathrm{supp}(m)| = |C| - |\mathrm{supp}(m) \cap C| \leq |C| - (d-l) \leq k$, and $|[n] \setminus \mathrm{supp}(m)| = n - d \geq k$. Let L be a subset of $\mathrm{supp}(m)$ of size l that contains all of $\mathrm{supp}(m) \setminus C$. Such a set L exists because $|\mathrm{supp}(m) \setminus C| \leq l$ and $|\mathrm{supp}(m)| = d \geq l$. Note that K and L are disjoint.

The monomial m has a nonzero contribution to $\partial_L r|_{K \leftarrow 0}$. In general, a monomial m' has a nonzero contribution to $\partial_L r|_{K \leftarrow 0}$ if and only if $\mathrm{supp}(m')$ is disjoint from K and contains L. The disjointness requirement implies that $\mathrm{supp}(m') \cap C \subseteq C \setminus K = \mathrm{supp}(m) \cap C$, where the equality follows from the choice of K. The inclusion requirement implies that $\mathrm{supp}(m) \setminus C = L \setminus C \subseteq \mathrm{supp}(m') \setminus C$, where the equality follows from the choice of L. In combination with the maximality of $|\mathrm{supp}(m) \setminus C|$ among the monomials of r of degree d, this means that either m' does not have degree d or else $\mathrm{supp}(m') \setminus C = \mathrm{supp}(m) \setminus C$. It follows that the only monomials m' of r of degree d that contribute to $\partial_L r|_{K \leftarrow 0}$ satisfy $\mathrm{supp}(m') \subseteq \mathrm{supp}(m)$. As r only contains multilinear monomials, r has exactly one monomial of degree d that has a nonzero contribution to $\partial_L r|_{K \leftarrow 0}$, namely the monomial m. We conclude that the polynomial $q(z)$ that results from substituting (5.1) into $\partial_L r|_{K \leftarrow 0}$ has a nonzero term of degree $d - l$. ◻

We now have all ingredients to establish Theorem 1.8.

Proof of Theorem 1.8. The necessity of condition 1 and condition 2 for the membership of a multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ in $\mathrm{Van}[\mathrm{RFE}_l^k]$ immediately follows from Proposition 5.1 and Lemma 5.3, respectively. For sufficiency, we need to show that if p only contains monomials of degrees between $l+1$ and $n-k-1$ and is hit by RFE_l^k, then there exist disjoint sets $K, L \subseteq [n]$ with $|K| = k$ and $|L| = l$ so that $\partial_L p|_{K \leftarrow 0}$ is nonzero upon the substitution (5.1).

By Lemma 5.5, we can write p as $p = q + r$, where $q \in \mathrm{Van}[\mathrm{RFE}_l^k]$ and r is a multilinear $(k+d-l, l)$-cored polynomial that is either zero or else has the same degree as p. Since $p \notin \mathrm{Van}[\mathrm{RFE}_l^k]$, the case of zero r is ruled out. Thus r is a multilinear $(k+d-l, l)$-cored polynomial of degree d, where $l+1 \leq d \leq n-k-1$. Lemma 5.6 then yields disjoint sets $K, L \subseteq [n]$ with $|K| = k$ and $|L| = l$ so that $\partial_L r|_{K \leftarrow 0}$ is nonzero upon the substitution (5.1). As both p and r are multilinear, so is $q = p - r$. The contrapositive of Lemma 5.3 implies that $\partial_L q|_{K \leftarrow 0}$ is zero upon the substitution (5.1). It follows that $\partial_L p|_{K \leftarrow 0} = \partial_L q|_{K \leftarrow 0} + \partial_L r|_{K \leftarrow 0}$ is nonzero upon the substitution (5.1). ◻

We conclude this section by detailing the connection between Theorem 1.8 and some prior applications of the SV generator.

Application to read-once formulas. We start with the theorem that SV^1 hits read-once formulas. The original proof in [41] goes by induction on the depth of F. The critical part is the inductive step for the case where the top gate is an addition, say $F = F_1 + F_2$. The argument in [41] involves a clever analysis that uses the variable-disjointness of F_1 and F_2 to show that $F_1(\mathrm{SV}^1)$ and $F_2(\mathrm{SV}^1)$ cannot cancel each other out. We present an alternate proof that has a similar inductive outline but follows a more structured, principled approach based on Theorem 1.8 for the critical part.

Theorem 5.7 ([41])**.** SV^1 *hits read-once formulas.*

Alternate proof. We show by induction on the depth the formula F that if F is nonconstant, then so is $F(\mathrm{SV}^1)$. This suffices because it implies that nonconstant formulas are hit by SV^1, and nonzero constant formulas are hit as the range of SV^1 is nonempty.

The inductive step consists of two cases, depending on whether the top gate is a multiplication gate or an addition gate. The case of a multiplication gate follows from the general property that the product of a nonconstant polynomial with any nonzero polynomial is nonconstant. It remains to consider the case of an addition gate.

For a nonconstant formula F, $F(\mathrm{SV}^1)$ is nonconstant iff SV^1 hits the variable part of F (which is a nonzero polynomial). By Theorem 1.8 with $k = 0$ and $l = 1$, the latter is the case iff at least one of the following two conditions hold:

1. F has a homogeneous component of degree 1 or at least n.

2. For some $L = \{i\} \subseteq [n]$, the derivative $\partial_{x_i} F$ is nonzero upon the substitution (1.7).

Consider a read-once formula F with an addition gate on top: $F = F_1 + F_2$. The variable-disjointness of F_1 and F_2 implies that if condition 1 holds for at least one of F_1 or F_2, then it holds for F. The same is true for condition 2. The inductive step in the case of an addition gate at the top follows. ◻

The case of an addition gate in the above proof has a clean geometric interpretation along the lines of the alternating algebra representation that we discussed in Section 1 for polynomials that are multilinear (which polynomials computed by read-once formulas are). Recall that we can think of the variables as vertices, and multilinear monomials as simplices made from those vertices.[2] A multilinear polynomial is a weighted collection of such simplices with weights from $\mathbb{F}$. In this view, Theorem 1.8 translates to the following characterization: a weighted collection of simplices corresponds to a polynomial in the vanishing ideal of RFE^0_1 iff there

[2]In this setting the orientation of the simplices does not matter.

are no simplices of zero, one, or all vertices (condition 1), and the remaining weights satisfy a certain system of linear equations (condition 2). Crucially, for each equation in the system, there is a vertex such that the equation only involves weights of the simplices *that contain that vertex*, namely the vertex corresponding to the variable x_i where $L = \{i\}$. Meanwhile, the sum of two variable-disjoint polynomials corresponds to taking the vertex-disjoint union of two weighted collections of simplices. It follows directly that if either of the two polynomials violates a requirement besides the "no simplex of zero vertices" requirement, then their sum violates the same requirement. The "no simplex of zero vertices" requirement holds automatically when considering the variable parts, and maps to the special handling of the constant term in the formal proof.

Zero-substitutions and partial derivatives. As mentioned in the overview, several prior papers demonstrated the utility of partial derivatives and zero substitutions in the context of derandomizing PIT using the SV generator, especially for syntactically multilinear models. By judiciously choosing variables for those operations, these papers managed to simplify p and reduce PIT for p to PIT for simpler instances, resulting in an efficient recursive algorithm. Such recursive arguments can be wrapped into a general framework, similar to the one presented in [39] for generic l-independent generators. Whereas the power of the framework in the generic setting remains open, thanks to Theorem 1.8, we can prove that our framework captures the full power of the specific l-independent generator SV^l. More generally, we exhibit a natural reformulation within the framework of any argument that RFE hits a certain class of multilinear polynomials, such as those computable with some bounded complexity in some syntactic model.

For the argument, we assume that we can break up the class in the following way.

Definition 5.8 (grading hypothesis)**.** A class $C = \bigcup_{k,l\in\mathbb{N}} C_{k,l}$ of polynomials satisfies the grading hypothesis if for every $k, l \in \mathbb{N}$ and $p \in C_{k,l}$, at least one of the following holds:

- $k = l = 0$ and p is nonzero.
- $k > 0$ and there is a zero substitution such that the result is in $C_{k-1,l}$.
- $l > 0$ and there is a partial derivative such that the result is in $C_{k,l-1}$.

Under the additional mild assumption of closure under variable rescaling, we obtain a parameter-efficient framework through direct applications of Theorem 1.8.

Proposition 5.9. *Let $C = \bigcup_{k,l\in\mathbb{N}} C_{k,l}$ be a class of polynomials that satisfies the grading hypothesis and such that each $C_{k,l}$ is closed under variable rescaling. If RFE_0^0 hits $C_{0,0}$ then RFE_l^k hits $C_{k,l}$ for every $k, l \in \mathbb{N}$.*

Proof. The proof is by induction on k and l. The base case is $k = l = 0$, where the claim is immediate. When $k > 0$ or $l > 0$, our hypotheses are such that $p \in C_{k,l}$ either simplifies under a zero substitution or a partial derivative. In either case, we show how a violation of the conditions in Theorem 1.8 for a simpler polynomial $p' \in \mathbb{F}[x'_1, \ldots, x'_n]$ translates into a corresponding violation of the conditions for $p \in \mathbb{F}[x_1, \ldots, x_n]$, where each variable x'_i is a rescaling of x_i.

More specifically, by condition 1 of Theorem 1.8, we may assume that p only has homogeneous components with degrees in the range $l+1, \ldots, n-k-1$. We argue in both cases that p' similarly satisfies condition 1 of Theorem 1.8. By the induction hypothesis and closure under variable rescaling, it follows that $\partial'_{L'} p'|_{K' \leftarrow 0}$ (where the prime in ∂' indicates that the partial derivatives are with respect to the primed variables x'_i) is nonzero for some K' and L' under a particular substitution. Out of K' and L' we then construct K and L such that $\partial_L p|_{K \leftarrow 0}$ is nonzero upon the substitution in condition 2 of Theorem 1.8, where variable rescaling between x'_i and x_i enables us to match the substitutions for $\partial'_{L'} p'|_{K' \leftarrow 0}$ and $\partial_L p|_{K \leftarrow 0}$. We provide the remaining details for each case separately.

- If p simplifies under a zero substitution $x_{j^*} \leftarrow 0$, then write p as $p = q x_{j^*} + r$ where q and r are polynomials that do not depend on x_{j^*}, and set $p'(\ldots, x'_i, \ldots) = r(\ldots, x_i, \ldots)$ with $x_i = x'_i \cdot (a_i - a_{j^*})$. By closure under rescaling, $p' \in \mathcal{C}_{k-1,l}$, so by induction p' is hit by RFE_l^{k-1}. We apply Theorem 1.8 to p' with respect to the set of variables $\{x'_1, \ldots, x'_{j^*-1}, x'_{j^*+1}, \ldots, x'_n\}$ and k replaced by $k-1$. As p only has homogeneous components with degrees in the range $l+1, \ldots, n-k-1$, so does p', and condition 1 of Theorem 1.8 holds for p'. This means that condition 2 does not hold for p'. Thus, there must be disjoint $K', L' \subseteq [n] \setminus \{j^*\}$ with $|K'| = k-1$ and $|L'| = l$ so that $\partial'_{L'} p'|_{K' \leftarrow 0}$ is nonzero upon the substitution
$$x'_i \leftarrow z \cdot \frac{\prod_{j \in K'}(a_i - a_j)}{\prod_{j \in L'}(a_i - a_j)}. \tag{5.2}$$
Setting $K = K' \cup \{j^*\}$ and $L = L'$, we have
$$\partial_L p|_{K \leftarrow 0} = \partial_{L'} r|_{K' \leftarrow 0} = \partial'_{L'} p'|_{K' \leftarrow 0} \Big/ \prod_{i \in L'} (a_i - a_{j^*})$$
and the substitution (5.2) induces the substitution (1.7).

- If p simplifies under a partial derivative $\partial_{x_{j^*}}$, then write p as $p = q x_{j^*} + r$ where q and r are polynomials that do not depend on x_{j^*}, and set $p'(\ldots, x'_i, \ldots) \doteq q(\ldots, x_i, \ldots)$ with $x_i = x'_i/(a_i - a_{j^*})$. By closure under rescaling, $p' \in \mathcal{C}_{k,l-1}$, so by induction p' is hit by RFE_{l-1}^k. We apply Theorem 1.8 to p' with respect to the set of variables $\{x'_1, \ldots, x'_{j^*-1}, x'_{j^*+1}, \ldots, x'_n\}$ and l replaced by $l-1$. As p' has homogeneous components of degrees one less than p does, condition 1 of Theorem 1.8 holds for p', so condition 2 must fail. Thus, there are disjoint $K', L' \subseteq [n] \setminus \{j^*\}$ with $|K'| = k$ and $|L'| = l-1$ so that $\partial'_{L'} p'|_{K' \leftarrow 0}$ is nonzero upon the substitution (5.2). Setting $K = K'$ and $L = L' \cup \{j^*\}$, we have
$$\partial_L p|_{K \leftarrow 0} = \partial_{L'} q|_{K' \leftarrow 0} = \partial'_{L'} p'|_{K' \leftarrow 0} \cdot \prod_{i \in L'} (a_i - a_{j^*})$$
and the substitution (5.2) induces the substitution (1.7).

In both cases we conclude that $\partial_L p|_{K \leftarrow 0}$ is nonzero upon the substitution (1.7), which is the sought violation of condition 2 of Theorem 1.8. ◻

We remark that the mild requirement of closure under variable rescaling in Proposition 5.9 can be dropped completely at the cost of reduced efficiency in parameters.[3]

Proposition 5.10. *Let $C = \bigcup_{k,l\in\mathbb{N}} C_{k,l}$ be a class of polynomials that satisfies the grading hypothesis. If* RFE_0^0 *hits* $C_{0,0}$ *then* RFE_{k+l}^{k+l} *hits* $C_{k,l}$ *for every* $k, l \in \mathbb{N}$.

Proof sketch. The strategy is the same as in the proof of Proposition 5.9, but in the inductive step the index i^* is added to both K' and L' instead of just one of the two sets. This obviates the need for rescaling to ensure that the substitutions match. Note that the resulting sets K and L are no longer disjoint, but the general Zoom Lemma accommodates overlapping sets K and L. ◻

Theorem 1.8 tells us that derivatives and zero substitutions suffice to witness when a multilinear polynomial p is hit by SV or RFE. One can ask, if we know more information about p, can we infer *which* derivatives and zero substitutions form a witness? In some cases we know. For example, if p has a low-support monomial $x_1 \cdots x_l$, then it suffices to take derivatives with respect to each of $x_1, \ldots, x_l$. On the other hand, consider that whenever two polynomials p and q are hit by SV, then so is their product pq. Given explicit witnesses for p and q, we do not know how to obtain an explicit witness for the product pq.

6 Sparseness

By Proposition 3.4, the generators EVC_l^k contain exactly $\binom{k+l+2}{l+1}$ monomials. The following result shows that no nonzero polynomial in the vanishing ideal of RFE_l^k has fewer monomials. Corollary 1.6 follows.

Lemma 6.1. *Suppose $p \in \mathbb{F}[x_1, \ldots, x_n]$ is nonzero and has only s monomials with nonzero coefficients. Then, for any k, l such that $\binom{k+l+2}{l+1} > s$,* RFE_l^k *hits p.*

The tactic here is to show that, if p has too few monomials appearing in it, then there is a way to instantiate the Zoom Lemma wherein p_{d^*} is a single monomial and therefore is nonzero upon the substitution (4.1).

Proof. For $i \in [n]$, we define two operations, $\downarrow_i$ and $\uparrow_i$, on nonempty sets of monomials. Applying $\downarrow_i$ to such a set M yields the subset of M consisting of the monomials in which x_i appears with its least degree among all the monomials in M. We define $\uparrow_i$ similarly, except we select the monomials in which x_i appears with its highest degree. We make the following claim:

Claim 6.2. *For any nonempty set of monomials with fewer than $\binom{k+l+2}{l+1}$ monomials, there is a sequence of $\downarrow$ and $\uparrow$ operations, with at most k $\downarrow$ operations and at most l $\uparrow$ operations, such that the resulting set of monomials has exactly one element.*

[3] This is a setting where we exploit the possibility of the sets K and L in the Zoom Lemma to overlap.

The claim implies the lemma as follows. Let M be the set of monomials with nonzero coefficient in p. Apply the claim to M to get a sequence of $\downarrow$ and $\uparrow$ operations resulting in a single monomial m_0. Let K denote the indices used for the $\downarrow$ operations and L the indices used for the $\uparrow$ operations. Let d^* be the degree pattern with domain $K \cup L$ that matches m_0. By how the operators are defined, every monomial m in M satisfies either

- $\deg_{x_i}(m) > d^*_i$ for some $i \in K$ (m was removed by $\downarrow_i$),
- $\deg_{x_i}(m) < d^*_i$ for some $i \in L$ (m was removed by $\uparrow_i$), or
- $\deg_{x_i}(m) = d^*_i$ for every $i \in K \cup L$, in which case $m = m_0$.

Accordingly, d^* is (K, L)-extremal in p and the Zoom Lemma applies. As p_{d^*} is a single monomial, it is nonzero upon the substitution (4.1). As $|K| \leq k$ and $|L| \leq l$, we conclude that p is hit by RFE^k_l.

It remains to prove Claim 6.2. We do this by induction on $|M|$. In the base case, $|M| = 1$, in which case the empty sequence suffices. Otherwise, $|M| > 1$, in which case there is a variable x_i that appears with at least two distinct degrees among monomials in M. The sets $\downarrow_i(M)$ and $\uparrow_i(M)$ are nonempty and disjoint. Since M has size less than $\binom{k+l+2}{l+1} = \binom{k+l+1}{l+1} + \binom{k+l+1}{l}$, either $\downarrow_i(M)$ has size less than $\binom{k+l+1}{l+1}$, or $\uparrow_i(M)$ has size less than $\binom{k+l+1}{l}$. Whichever is the case, the claim follows by applying the inductive hypothesis to it. □

7 Set-Multilinearity

Although the generators EVC^k_l provided by Theorem 1.3 are not set-multilinear, the vanishing ideal of RFE^k_l does contain set-multilinear polynomials. In this section, we construct some of degree $l + 1$ with partition classes of size $k + 2$. In fact, we argue that all set-multilinear polynomials in $\mathrm{Van}[\mathrm{RFE}^k_l]$ of degree $l + 1$ are in the linear span of the ones we construct.

Our construction is a modification of the one for EVC^k_l.

Definition 7.1. Let $k, l, n \in \mathbb{N}$, and let $S_1, \ldots, S_{l+1} \subseteq [n]$ be $l + 1$ disjoint subsets of $k + 2$ indices each. The polynomial ESMVC^k_l is an $(l + 1) \times (l + 1)$ determinant where each entry is itself a $(k + 2) \times (k + 2)$ determinant. We index the rows in the outer determinant by $i = 1, \ldots, l + 1$, and the columns by $d = l, \ldots, 0$. In each (i, d)-th inner matrix, there is one row per $j \in S_i$; it is

$$\begin{bmatrix} a^k_j & a^{k-1}_j & \cdots & a^1_j & a^0_j & a^d_j x_j \end{bmatrix}.$$

The name "ESMVC" is a shorthand for "Elementary Set-Multilinear Vandermonde Circulation". Similar to EVC, the precise instantiation of ESMVC requires one to pick an order for the sets $S_1, \ldots, S_{l+1}$ and an order within each set.

Example 7.2. When $k = 1$ and $l = 2$, ESMVC uses three sets of three variables each. To help convey the structure of the determinant, we name the variable-sets $S_1 = \{x_1, x_2, x_3\}$,

$S_2 = \{y_1, y_2, y_3\}$, and $S_3 = \{z_1, z_2, z_3\}$, and denote the abscissa of x_i by a_i, the abscissa of y_i by b_i, and the abscissa of z_i by c_i. With this notation and using the index ordering, ESMVC is the following:

$$\begin{Vmatrix} \begin{vmatrix} a_1^1 & a_1^0 & a_1^2 x_1 \\ a_2^1 & a_2^0 & a_2^2 x_2 \\ a_3^1 & a_3^0 & a_3^2 x_3 \end{vmatrix} & \begin{vmatrix} a_1^1 & a_1^0 & a_1^1 x_1 \\ a_2^1 & a_2^0 & a_2^1 x_2 \\ a_3^1 & a_3^0 & a_3^1 x_3 \end{vmatrix} & \begin{vmatrix} a_1^1 & a_1^0 & a_1^0 x_1 \\ a_2^1 & a_2^0 & a_2^0 x_2 \\ a_3^1 & a_3^0 & a_3^0 x_3 \end{vmatrix} \\ \begin{vmatrix} b_1^1 & b_1^0 & b_1^2 y_1 \\ b_2^1 & b_2^0 & b_2^2 y_2 \\ b_3^1 & b_3^0 & b_3^2 y_3 \end{vmatrix} & \begin{vmatrix} b_1^1 & b_1^0 & b_1^1 y_1 \\ b_2^1 & b_2^0 & b_2^1 y_2 \\ b_3^1 & b_3^0 & b_3^1 y_3 \end{vmatrix} & \begin{vmatrix} b_1^1 & b_1^0 & b_1^0 y_1 \\ b_2^1 & b_2^0 & b_2^0 y_2 \\ b_3^1 & b_3^0 & b_3^0 y_3 \end{vmatrix} \\ \begin{vmatrix} c_1^1 & c_1^0 & c_1^2 z_1 \\ c_2^1 & c_2^0 & c_2^2 z_2 \\ c_3^1 & c_3^0 & c_3^2 z_3 \end{vmatrix} & \begin{vmatrix} c_1^1 & c_1^0 & c_1^1 z_1 \\ c_2^1 & c_2^0 & c_2^1 z_2 \\ c_3^1 & c_3^0 & c_3^1 z_3 \end{vmatrix} & \begin{vmatrix} c_1^1 & c_1^0 & c_1^0 z_1 \\ c_2^1 & c_2^0 & c_2^0 z_2 \\ c_3^1 & c_3^0 & c_3^0 z_3 \end{vmatrix} \end{Vmatrix}.$$

The elementary properties of EVC_l^k from Proposition 3.4 extend as follows to ESMVC_l^k.

Proposition 7.3. *For any $k, l \in \mathbb{N}$ and index sets $S_1, \ldots, S_{l+1}$ as in Definition 7.1, ESMVC_l^k is skew-symmetric with respect to the order of the sets $S_1, \ldots, S_{l+1}$, and the choice of order within each set, in that any permutation thereof changes the construction by merely multiplying by the sign of the permutation. For any order, ESMVC_l^k is nonzero, homogeneous of degree $l+1$, and set-multilinear with respect to the partition $S_1, \ldots, S_{l+1}$, and every monomial consistent with the partitions appears with a nonzero coefficient. When the sets are ordered as $S_1, \ldots, S_{l+1}$ and the variables associated with S_i are labeled and ordered as $(x_{i,1}, \ldots, x_{i,k+2})$ for $i = 1, \ldots, l+1$, the coefficient of $x_{1,1} \cdot \cdots \cdot x_{l+1,1}$ equals*

$$(-1)^{(k+1)(l+1)} \cdot \begin{vmatrix} a_{1,1}^l & \cdots & a_{1,1}^0 \\ \vdots & \ddots & \vdots \\ a_{l+1,1}^l & \cdots & a_{l+1,1}^0 \end{vmatrix} \cdot \prod_{i=1}^{l+1} \begin{vmatrix} a_{i,2}^k & \cdots & a_{i,2}^0 \\ \vdots & \ddots & \vdots \\ a_{i,k+2}^k & \cdots & a_{i,k+2}^0 \end{vmatrix}. \tag{7.1}$$

Proof. All assertions to be proved follow from elementary properties of determinants, that Vandermonde determinants are nonzero unless they have duplicate rows, and the following computation for the coefficient of $x_{1,1} \cdot \cdots \cdot x_{l+1,1}$: Plug 1 into $x_{i,1}$ for $i = 1, \ldots, l+1$ and 0 into the remaining variables, and minor expand along the last column each of the inner determinants. Due to the minor expansions, the elements in the i-th row of the outer determinant have a common factor of $(-1)^{k+1}$ times the $(k+1) \times (k+1)$ determinant for that value of i in the product on the right-hand side of (7.1). After removing those common factors from all $l+1$ rows, the remaining $(l+1) \times (l+1)$ outer determinant equals the determinant in the middle of (7.1). □

The following theorem formalizes the role ESMVC plays among the degree-$(l+1)$ polynomials with respect to $\mathrm{Van}[\mathrm{RFE}_l^k]$.

Theorem 7.4. *Let $k, l \in \mathbb{N}$ and let $X_1, \ldots, X_{l+1}$ be $l+1$ disjoint sets of indices (of any size). The linear span of $\mathrm{ESMVC}_l^k[S_1, \ldots, S_{l+1}]$, over all choices of $S_i \subseteq X_i$ with $|S_i| = k+2$, equals the set-multilinear polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$ with variable partition $(X_1, \ldots, X_{l+1})$.*

Theorem 7.4 and Proposition 7.3 imply Corollary 1.7 that there are no set-multilinear polynomials of degree $l+1$ in $\mathrm{Van}[\mathrm{RFE}_l^k]$ that have at least one partition X_i of size less than $k+2$.

The proof of Theorem 7.4 follows the same outline as the one of Theorem 1.3 in Section 3. We start by showing that all instantiations of ESMVC_l^k are contained in $\mathrm{Van}[\mathrm{RFE}_l^k]$ using a similar argument as that for EVC_l^k.

Lemma 7.5. *For every $k, l \in \mathbb{N}$ and every choice of $l+1$ disjoint sets $S_1, \ldots, S_{l+1}$ of $k+2$ indices each, $\mathrm{ESMVC}_l^k[S_1, \ldots, S_{l+1}]$ vanishes at RFE_l^k.*

Proof. Let g/h be a seed for RFE_l^k. Let A be the $(l+1) \times (l+1)$ outer matrix defining ESMVC, so that $\mathrm{ESMVC} \doteq \det(A)$. Recall that the columns of A are indexed by $d = l, \ldots, 0$. Let $\vec{h} \in \mathbb{F}^{l+1}$ be the column vector where the row indexed by d is the coefficient of α^d in $h(\alpha)$. We show that, after substituting $\mathrm{RFE}_l^k(g/h)$, the matrix-vector product $A\vec{h} \in \mathbb{F}^{l+1}$ yields the zero vector. It follows that evaluating ESMVC at $\mathrm{RFE}_l^k(g/h)$ vanishes, as it is the determinant of a singular matrix.

Fix $i \in \{1, \ldots, l+1\}$, and focus on the i-th coordinate of $A\vec{h}$. The (i,d) entry of A is a determinant; let $B_{i,d}$ be the inner matrix as in Definition 7.1. As d varies, only the last column of $B_{i,d}$ changes. Thus, by multilinearity of the determinant, the i-th entry of $A\vec{h}$ is itself a determinant. Recalling that the rows of $B_{i,l}, \ldots, B_{i,0}$ are indexed by $j \in X_i$, the j-th row of this determinant is

$$\begin{bmatrix} a_j^k & \cdots & a_j^0 & h(a_j)x_j \end{bmatrix}.$$

After substituting $\mathrm{RFE}_l^k(g/h)$, it becomes

$$\begin{bmatrix} a_j^k & \cdots & a_j^0 & g(a_j) \end{bmatrix}.$$

Since g is a degree-k polynomial, the columns of $B_{i,d}$ are linearly dependent, so the determinant is zero. □

Next, we argue that every polynomial in $\mathrm{Van}[\mathrm{RFE}_l^k]$ that is set-multilinear with respect to the variable partition $(X_1, \ldots, X_{l+1})$ is in the ideal I generated by the instantiations of ESMVC_l^k in the statement of Theorem 7.4. We use a similar two-step approach as for Theorem 1.3 in Section 3.

1. Modulo the ideal I, every polynomial p is equal to a polynomial r (depending on p) with a certain structure (Lemma 7.7).

2. Every nonzero polynomial r that has the structure and is is set-multilinear with respect to the variable partition $(X_1, \ldots, X_{l+1})$ is hit by RFE_l^k (Lemma 7.8).

For step 1, we need a suitable replacement for being (c,t)-cored. The following adaptation to the set-multilinear setting suffices.

Definition 7.6. Let $X_1, \dots, X_d \subseteq [n]$ be disjoint sets of indices. A polynomial that is set-multilinear with respect to the partition $(X_1, \dots, X_d)$ is (c,t)-*multi-cored* if there exists a set $C \doteq C_1 \sqcup \dots \sqcup C_d$, with $C_i \subseteq X_i$, $|C_i| \leq c$, such that every monomial m of the polynomial satisfies $|\mathrm{supp}(M) \setminus C| \leq t$.

We refer to the set C in Definition 7.6 as a multi-core.

Lemma 7.7. *Let $k, l \in \mathbb{N}$ and let $X_1, \dots, X_{l+1} \subseteq [n]$ be disjoint sets of indices. Suppose $C \doteq C_1 \sqcup \dots \sqcup C_{l+1}$ is a set of indices such that $C_i \subseteq X_i$ and $|C_i| = k+1$. Let I be the ideal generated by the polynomials $\mathrm{ESMVC}_l^k[C_1 \sqcup \{j_1\}, \dots, C_{l+1} \sqcup \{j_{l+1}\}]$, where $j_i \in X_i \setminus C_i$. Modulo I, every set-multilinear polynomial with respect to the variable partition $X_1, \dots, X_{l+1}$ equals a $(k+1, l)$-multi-cored polynomial with multi-core C.*

Proof. By linearity it suffices to establish the result for any monomial m that is set-multilinear with respect to the partition $(X_1, \dots, X_{l+1})$. If $\mathrm{supp}(m) \cap C$ is nonempty, then m is already $(k+1, l)$-multi-cored with multi-core C because m only has $l+1$ variables in its support. Otherwise, let $m = x_{j_1} \cdots x_{j_{l+1}}$. By Proposition 7.3, the polynomial $\mathrm{ESMVC}_l^k[C_1 \sqcup \{j_1\}, \dots, C_{l+1} \sqcup \{j_{l+1}\}]$ can be written as $c \cdot m + r$ where $c \in \mathbb{F}$ is nonzero and r is a linear combination of monomials m' that are set-multilinear with respect to the partition $(X_1, \dots, X_{l+1})$ and such that $\mathrm{supp}(m') \cap C$ is nonempty. The result for m follows by writing $m \equiv -c^{-1} \cdot r \bmod I$. ◻

Step 2 is another application of the Zoom Lemma. We make use of the version geared towards multilinear polynomials, namely Lemma 5.3.

Lemma 7.8. *Let $k, l \in \mathbb{N}$ and let $X_1, \dots, X_{l+1} \subseteq [n]$ be disjoint sets of indices. Every nonzero polynomial that is set-multilinear with respect to the partition $(X_1, \dots, X_{l+1})$ and that is $(k+1, l)$-multi-cored is hit by RFE_l^k.*

Proof. Let r satisfy the hypotheses of the lemma with multi-core C. Let m^* be a monomial in r for which $\mathrm{supp}(m^*) \setminus C$ is maximal with respect to inclusion. Such a monomial exists because r is nonzero. Let $j^* \in \mathrm{supp}(m^*) \cap C$. Such an index exists since $|\mathrm{supp}(m^*)| = l+1$ and $|\mathrm{supp}(m^*) \setminus C| \leq l$ by the multi-core property. Let $i^* \in [l+1]$ be such that $j^* \in X_{i^*}$. Set $K \doteq C \cap X_{i^*} \setminus \{j^*\}$ and $L \doteq \mathrm{supp}(m^*) \setminus \{j^*\}$. Note that $|K| \leq (k+1)-1 = k$ and $|L| \leq (l+1)-1 = l$. By set-multilinearity, monomials m in r for which $\partial_L m|_{K \leftarrow 0}$ is nonzero need to have the form $x_j \cdot x^L$ where $j \in X_{i^*} \setminus K$. The monomial m^* is of the form with $j = j^*$. By the maximality of m^*, any monomial in r of the form has to have $j \in C$. Since $C \cap (X_{i^*} \setminus K) = \{j^*\}$, it follows that m^* is the only monomial in r that contributes to $\partial_L r|_{K \leftarrow 0}$. Since $\partial_L m^*|_{K \leftarrow 0} = x_{j^*}$, it follows that $\partial_L r|_{K \leftarrow 0}$ is nonzero upon the substitution (5.1). We conclude that RFE_l^k hits r by Lemma 5.3. ◻

We now have all ingredients to establish Theorem 7.4.

Proof of Theorem 7.4. Let $\mathcal{S} \doteq (S_1, \ldots, S_{l+1})$ range as in the statement. The linear span of the polynomials $\mathrm{ESMVC}_l^k[\mathcal{S}]$ is set-multilinear with respect to the variable partition $(X_1, \ldots, X_{l+1})$ by Proposition 7.3, and in $\mathrm{Van}[\mathrm{RFE}_l^k]$ by Lemma 7.5. In the other direction, the combination of Lemma 7.7 and Lemma 7.8 imply that every polynomial $p \in \mathrm{Van}[\mathrm{RFE}_l^k]$ that is set-multilinear with respect to the variable partition $(X_1, \ldots, X_{l+1})$ falls inside the ideal I generated by the polynomials $\mathrm{ESMVC}_l^k[\mathcal{S}]$, i. e., $p = \sum_{\mathcal{S}} q_{\mathcal{S}} \mathrm{ESMVC}_l^k[\mathcal{S}]$ for some polynomials $q_{\mathcal{S}}$. As all polynomials $\mathrm{ESMVC}_l^k[\mathcal{S}]$ as well as p are homogeneous of degree $l+1$, it follows that each $q_{\mathcal{S}}$ can be replaced by its constant term. □

8 Read-Once Oblivious Algebraic Branching Programs

In this section we provide some background on ROABPs and establish Theorem 1.9.

8.1 Background

Algebraic branching programs are a syntactic model for algebraic computation. One forms a directed graph with a designated source and sink. Each edge is labeled by a polynomial that depends on at most one variable among $x_1, \ldots, x_n$. The branching program computes a polynomial in $\mathbb{F}[x_1, \ldots, x_n]$ by summing, over all source-to-sink paths, the product of the labels on the edges of each path.

A special subclass of algebraic branching programs are read-once oblivious algebraic branching programs (ROABPs). In this model, the vertices of the branching program are organized in layers. The layers are totally ordered, and edges exist only from one layer to the next. For each variable, there is at most one consecutive pair of layers between which that variable appears, and for each pair of consecutive layers, there is at most one variable that appears between them. In this way, every source-to-sink path reads each variable at most once (the branching program is *read-once*), and the order in which the variables are read is common to all paths (the branching program is *oblivious*). We can always assume that the number of layers equals one plus the number of variables under consideration.

The number of vertices comprising a layer is called its *width*. The width of an ROABP is the largest width of its layers. The minimum width of an ROABP computing a given polynomial can be characterized in terms of the rank of coefficient matrices constructed as follows.

Definition 8.1. Let $U \sqcup V = [n]$ be a partition of the variable indices, and let M_U and M_V be the sets of monomials that are supported on variables indexed by U and V, respectively. For any polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ define the matrix

$$\mathrm{CMat}_{U,V}(p) \in \mathbb{F}^{M_U \times M_V}$$

by setting the (m_U, m_V) entry to equal the coefficient of $m_U m_V$ in p.

$\mathrm{CMat}_{U,V}(p)$ is formally an infinite matrix, but it has only finitely many nonzero entries. When p has degree at most d, one can just as well truncate $\mathrm{CMat}_{U,V}(p)$ to include only rows and columns indexed by monomials of degree at most d.

Lemma 8.2 ([42]). *Let $p \in \mathbb{F}[x_1, \ldots, x_n]$ be any polynomial. There is an ROABP of width w computing p in the variable order $x_1, \ldots, x_n$ if and only if, for every $s \in \{0, \ldots, n\}$, with respect to the partition $U = \{1, \ldots, s\}$ and $V = \{s+1, \ldots, n\}$, we have*

$$\operatorname{rank}(\mathrm{CMat}_{U,V}(p)) \leq w.$$

We group the monomials in M_U and M_V by their degrees and order the groups by increasing degree. This induces a block structure on $\mathrm{CMat}_{U,V}(p)$ with one block for every choice of $r, c \in \mathbb{N}$; the (r, c) block is the submatrix with rows indexed by degree-r monomials in M_U and columns indexed by degree-c monomials in M_V. In the case where p is homogeneous, the only nonzero blocks occur for $r + c$ equal to the degree of p. In this case the rank of $\mathrm{CMat}_{U,V}(p)$ is the sum of the ranks of its blocks.

In general, the rank of $\mathrm{CMat}_{U,V}(p)$ is at least the rank of $\mathrm{CMat}_{U,V}(p^{(\min)})$, where $p^{(\min)}$ denotes the homogeneous component of p of the lowest degree, $d_{\min}$. This follows because the submatrix of $\mathrm{CMat}_{U,V}(p)$ consisting of the rows and columns indexed by monomials of degree at most $d_{\min}$ has a block structure that is triangular with the blocks of $\mathrm{CMat}_{U,V}(p^{(\min)})$ on the hypotenuse. The observation yields the following folklore consequence of Lemma 8.2.

Proposition 8.3. *Let $p \in \mathbb{F}[x_1, \ldots, x_n]$ be any nonzero polynomial, and let $p^{(\min)}$ be the nonzero homogeneous component of p of least degree. If p can be computed by an ROABP of width w, then so can $p^{(\min)}$.*

8.2 Hitting property / lower bound

We now prove the ROABP hitting property of SV given in Theorem 1.9 and the equivalent ROABP lower bound given in Theorem 1.10. Both theorems follow from the next statement in a standard way.

Theorem 8.4. *For any integer $l \geq 1$, every nonzero multilinear homogeneous polynomial of degree $l+1$ in the vanishing ideal of SV^l requires ROABP width at least $(l/3) + 1$.*

For completeness, before proving Theorem 8.4, we argue how our ROABP hitting property and lower bound follow.

Proof of Theorem 1.9 and Theorem 1.10. The theorems are equivalent by complementation. We explain how Theorem 1.9 follows from Theorem 8.4.

Fix p satisfying the hypotheses of Theorem 1.9. We show that RFE_l^{l-1} hits p; this implies SV^l hits p because RFE_l^{l-1} and SV^l are equivalent up to variable rescaling, and rescaling variables does not affect ROABP width.

If p contains a monomial depending on at most l variables, then Proposition 5.1 implies that RFE_l^{l-1} hits p. The remaining case is when the homogeneous component $p^{(\min)}$ of the least degree is multilinear of degree $l+1$. By Proposition 8.3, $p^{(\min)}$ has ROABP width less than $(l/3) + 1$. By Theorem 8.4, $p^{(\min)}$ is hit by RFE_l^{l-1}, and by Proposition 5.2 so is p. ◻

In the remainder of this section we establish Theorem 8.4. We do not try to optimize the dependence of the bound on l.

Fix a positive integer l, and fix an arbitrary variable order, say $x_1, \ldots, x_n$. We show that, for every polynomial p that is nonzero, multilinear, and homogeneous of degree $l+1$, and that belongs to the vanishing ideal of RFE_l^{l-1}, there exists some $s \in \{0, \ldots, n\}$ so that, with respect to the partition $U = \{1, \ldots, s\}$, $V = \{s+1, \ldots, n\}$, it holds that $\mathrm{rank}(\mathrm{CMat}_{U,V}(p)) \geq (l/3) + 1$. Theorem 8.4 then follows by Lemma 8.2.

Let $C \doteq \mathrm{CMat}_{U,V}(p)$. As p is homogeneous of degree $l+1$, C is block diagonal, with a block C_d for each $d \in \{0, \ldots, l+1\}$ consisting of the rows indexed by monomials of degree d and the columns indexed by monomials of degree $l+1-d$. The block diagonal structure implies $\mathrm{rank}(C) = \sum_{d=0}^{l+1} \mathrm{rank}(C_d)$.

Via condition 2 of Theorem 1.8, the hypothesis that p belongs to $\mathrm{Van}[\mathrm{RFE}_l^{l-1}]$ induces linear equations on the entries in the blocks C_d. For homogeneous polynomials like p, the condition stipulates that for all disjoint subsets $K, L \subseteq [n]$ with $|K| = k = l-1$ and $|L| = l$, $\partial_L p|_{K \leftarrow 0}$ vanishes at the point (1.7) with $z = 1$. This is a linear equation in the coefficients of $\partial_L p|_{K \leftarrow 0}$, which are entries in the blocks C_d of C. In fact, each of these equations only reads entries from two *adjacent* blocks, i. e., blocks C_d and $C_{d'}$ with $|d - d'| = 1$. This is because L has size l, one less than the degree of p, so the only monomials that contribute to $\partial_L p|_{K \leftarrow 0}$ are those that are one variable x_i times the product of the variables indexed by L. It follows that the corresponding linear equation on C reads only entries that reside in the blocks $C_{|L \cap U|+1}$ (for $i \in U$) and $C_{|L \cap U|}$ (for $i \in V$).

We exploit the structure of these equations and argue that, for an appropriate choice of the partition index s, $\mathrm{rank}(C)$ is high.

Ingredients. Our analysis has four ingredients. The first ingredient is the fact that $\mathrm{rank}(C)$ is at least the number of nonzero blocks C_d. This is because a nonzero block has rank at least 1, and $\mathrm{rank}(C)$ is the sum of the ranks of the blocks. This simple observation means we can focus on situations where relatively few of the blocks are nonzero.

The second ingredient establishes an alternative lower bound on $\mathrm{rank}(C)$ in terms of the minimum distance between a nonzero block C_d and either extreme ($d = 0$ or $d = l+1$). Another way to think about this distance is as the maximum t such that every monomial in p depends on at least t variables indexed by U and at least t variables indexed by V.

Lemma 8.5. *Let $p \in \mathrm{Van}[\mathrm{RFE}_l^{l-1}]$ be nonzero, multilinear, and homogeneous of degree $l+1$, let $U \sqcup V$ be a partition of $[n]$, and let $C \doteq \mathrm{CMat}_{U,V}(p)$. If every monomial in p depends on at least t variables indexed by U and at least t variables indexed by V, then* $\mathrm{rank}(C) \geq t+1$.

The proof involves revisiting the equations from condition 2 of Theorem 1.8 and modifying[4] the underlying instantiations of Lemma 4.3 to obtain a system of linear equations with a simple enough structure that we can analyze.

[4]This is a setting where we exploit the possibility of the sets K and L in the Zoom Lemma to overlap.

The remaining ingredients allow us to reduce to situations where either the first or second ingredient applies. The third ingredient lets us fix any two zero blocks and zero out all the blocks that are not between them.

Proposition 8.6. *Let $p \in \mathrm{Van}[\mathrm{RFE}_l^{l-1}]$ be multilinear and homogeneous of degree $l+1$. Let $U \sqcup V$ be a partition of $[n]$, and let $C \doteq \mathrm{CMat}_{U,V}(p)$. Suppose that for some $d_1, d_2 \in \{-1, \ldots, l+2\}$ with $d_1 \leq d_2$, we have $C_{d_1} = 0$ and $C_{d_2} = 0$, where $C_{-1} \doteq 0$ and $C_{l+2} \doteq 0$. Let p' be the polynomial obtained from p by zeroing out the blocks C_d with $d < d_1$ or $d > d_2$. Then p' belongs to $\mathrm{Van}[\mathrm{RFE}_l^{l-1}]$.*

As zeroing out blocks does not increase the rank of C, our lower bound for $\mathrm{rank}(C)$ reduces to the same lower bound for the rank of $\mathrm{CMat}_{U,V}(p')$. This effectively extends the scope of the second ingredient: Alone, the second ingredient requires that *all* nonzero blocks of C be far from the extremes; with the third ingredient, it suffices that there exists a subinterval of nonzero blocks that is surrounded by zero blocks and that is far from the extremes. The proof hinges on the adjacent-block property of the equations from condition 2 of Theorem 1.8.

The ingredients thus far suffice provided there exists a nonzero block far from the extremes: Such a block belongs to some subinterval of nonzero blocks that is surrounded by zero blocks, say C_{d_1} to the left and C_{d_2} to the right, and the subinterval either is large and therefore has many nonzero blocks such that the first ingredient applies, or else it is small and therefore stays far from the extremes such that the combination of the second and third ingredients applies. See Figure 1 for an illustration. The fourth and final ingredient lets us ensure there is a nonzero block far from the extremes by setting the partition index s appropriately. In fact, it lets us guarantee a zero-to-nonzero transition at a position of our choosing.

Proposition 8.7. *For every $d \in \{-1, \ldots, l\}$, there is $s \in \{0, \ldots, n\}$ such that $C_d = 0$ and $C_{d+1} \neq 0$ with respect to the partition $U = \{1, \ldots, s\}$, $V = \{s+1, \ldots, n\}$, where $C_{-1} \doteq 0$.*

Combining ingredients. Let us find out what lower bound on $\mathrm{rank}(C)$ the prior ingredients give us as a function of the position $d = d_1$ in the interval where we have a guaranteed zero-to-nonzero transition as in Proposition 8.7. Starting from position d_1, keep increasing the position index until we hit the next zero block, say at position d_2, where we use $C_{l+2} \doteq 0$ as a sentinel. See Figure 1.

1. By the first ingredient, since the middle interval consists of nonzero blocks only, $\mathrm{rank}(C) \geq d_2 - d_1 - 1$.

2. By the combination of the second and the third ingredient, we have that $\mathrm{rank}(C) \geq t+1$ where $t = \min(d_1 + 1, l + 2 - d_2)$ is the minimum length of the leftmost and rightmost intervals. Indeed, let p' be the polynomial obtained from p by zeroing out the blocks C_d with $d < d_1$ or $d > d_2$. By Proposition 8.6 $p' \in \mathrm{Van}[\mathrm{RFE}_l^{l-1}]$. The polynomial p' is nonzero as it contains the original block C_{d+1}, which is nonzero. It is homogeneous of degree $l+1$ and multilinear as all of its monomials also occur in the homogeneous multilinear polynomial p of degree $l+1$. By construction, every monomial in p' contains at least $d_1 + 1$ variables indexed by U, and at least $l + 2 - d_2$ variables indexed by V. As such,

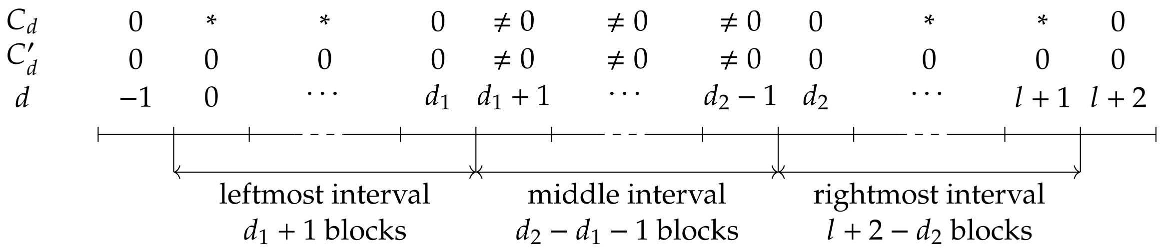

Figure 1: Rank lower bound analysis in terms of the blocks C_d of p and C'_d of p' (Proposition 8.6)

p' satisfies the conditions of Lemma 8.5 with $t = \min(d_1 + 1, l + 2 - d_2)$. It follows that $\operatorname{rank}(C) \geq \operatorname{rank}(\mathrm{CMat}_{U,V}(p')) \geq t + 1$.

If the rightmost interval has length at least the leftmost interval ($l + 2 - d_2 \geq d_1 + 1$), then item 2 yields $\operatorname{rank}(C) \geq d_1 + 2$. Otherwise, the rightmost interval is strictly shorter than the leftmost interval ($d_1 + 1 > l + 2 - d_2$); this implies that the middle interval has length at least $l - 2d_1 + 1$, which by item 1 yields $\operatorname{rank}(C) \geq l - 2d_1 + 1$. In any case, the bound $\operatorname{rank}(C) \geq \min(d_1 + 2, l - 2d_1 + 1)$ holds. Taking $d_1 = \left\lfloor \frac{l-1}{3} \right\rfloor$ optimizes this expression, achieving $\operatorname{rank}(C) \geq \left\lfloor \frac{l-1}{3} \right\rfloor + 2 \geq (l/3) + 1$. This completes the proof of Theorem 8.4 modulo the proofs of ingredients two through four.

Proofs. We conclude by proving ingredients two through four. We start with the one that requires the least specificity (ingredient 4, Proposition 8.7), then do ingredient 3 (Proposition 8.6), and end with the one that involves the most structure (ingredient 2, Lemma 8.5).

Proof of Proposition 8.7. When $s = 0$, C_0 contains all entries. As s increases by 1, some entries move from their current block $C_{d'}$ to the next block $C_{d'+1}$. Finally, when $s = n$, C_{l+1} contains all entries. For $d \geq 0$, it follows that every nonzero entry moves from C_d to C_{d+1} at some time. If we stop increasing s right after the last nonzero entry of C moves out of C_d, we have $C_d = 0$ and $C_{d+1} \neq 0$. For $d = -1$, we can pick $s = 0$ as $C_{-1} = 0$ and $C_0 = C \neq 0$. ◻

Proof of Proposition 8.6. It suffices to show that whenever p satisfies the two conditions in Theorem 1.8, then so does p'. Both p and p' are homogeneous. Condition 1 holds for p' as p' either is zero or else has the same degree as p. Regarding condition 2, as mentioned, the condition is equivalent to a system of homogeneous linear equations on $C' \doteq \mathrm{CMat}_{U,V}(p')$, each involving only an adjacent pair of blocks in C'. Those that involve only blocks C'_d with $d \leq d_1$ are met as the equations are homogeneous and the involved blocks are all zero. The same holds for the equations that involve only blocks C'_d with $d \geq d_2$. The remaining equations involve only blocks C'_d with $d \in \{d_1, \ldots, d_2\}$, on which p and p' agree. As the equations hold for C, they also hold for C'. ◻

It remains to argue Lemma 8.5. Our proof makes use of linear equations that are closely related to those given by Theorem 1.8, which in turn come from the Zoom Lemma. We revisit the application of the Zoom Lemma so as to obtain a simpler coefficient matrix—ultimately a Cauchy matrix—that enables a deeper analysis. To facilitate the discussion, we utilize the following notation. As p is multilinear, we only need to consider rows indexed by monomials of the form $\prod_{i\in I} x_i$ for $I \subseteq U$ and columns indexed by monomials of the form $\prod_{j\in J} x_j$ for $J \subseteq V$. This allows us to index rows by subsets $I \subseteq U$ and columns by subsets $J \subseteq V$. For $I \subseteq U$ and $J \subseteq V$ we denote by $C(I, J)$ the corresponding entry of C. The following proposition describes the linear equations we use.

Proposition 8.8. *Let $p \in \mathrm{Van}[\mathrm{RFE}_l^{l-1}]$ be multilinear, and homogeneous of degree $l+1$, let $U \sqcup V$ be a partition of $[n]$, and let $C \doteq \mathrm{CMat}_{U,V}(p)$. For every $I \subseteq U$ and $J \subseteq V$ with $|I| + |J| = l$, and for every $j^* \in I \cup J$,*

$$\sum_{i\in U\setminus I} \frac{C(\{i\}\cup I, J)}{a_i - a_{j^*}} + \sum_{i\in V\setminus J} \frac{C(I, \{i\}\cup J)}{a_i - a_{j^*}} = 0. \tag{8.1}$$

Proof. Set $L \doteq I \cup J$ and $K \doteq L \setminus \{j^*\}$, and note that $K \subseteq L$. Let $d^* \in \mathbb{N}^L$ be the all-1 degree pattern with domain L, and let $m^* \doteq \prod_{i\in L} x_i$ be the monomial supported on L that matches d^*. As p is multilinear, d^* is (K, L)-extremal in p. Since p is in $\mathrm{Van}[\mathrm{RFE}_l^{l-1}]$, the contrapositive of the Zoom Lemma tells us that the coefficient p_{d^*} of p vanishes at the point (4.1) with $z = 1$.

The multilinear monomials m of degree $l+1$ that match d^* have the form $m = x_i \cdot m^*$, where $i \in [n] \setminus L$. Thus, we can write the coefficient p_{d^*} as

$$p_{d^*} = \sum_{i\in U\setminus I} C(\{i\}\cup I, J)\cdot x_i + \sum_{i\in V\setminus J} C(I, \{i\}\cup J)\cdot x_i. \tag{8.2}$$

For each $i \in [n] \setminus L$, (4.1) with $z = 1$ substitutes $1/(a_i - a_{j^*})$ into x_i. Plugging this into (8.2) yields (8.1). ◻

Proof of Lemma 8.5. The proof goes by induction on t. The base case is $t = 0$, where the lemma holds because the rank of a nonzero matrix is always at least 1. For the inductive step, where $t \geq 1$, we zoom in on the contributions of the monomials that contain a particular variable. More precisely, for $j^* \in [n]$, let p_{j^*} denote the partial derivative $p_{j^*} \doteq \partial_{x_{j^*}} p$. Consider any $j^* \in [n]$ such that p_{j^*} is nonzero. As p is multilinear and homogeneous of degree $l+1$, p_{j^*} is multilinear and homogeneous of degree l. As every monomial in p depends on at least t variables indexed by U and at least t variables indexed by V, every monomial in p_{j^*} depends on at least $t-1$ variables indexed by U and at least $t-1$ variables indexed by V. In a moment, we argue that for every $j^* \in [n]$, $p_{j^*} \in \mathrm{Van}[\mathrm{RFE}_{l-1}^{l-2}]$. Then we will show the following:

Claim 8.9. *There exists $j^* \in [n]$ such that $p_{j^*} \neq 0$ and*

$$\mathrm{rank}(\mathrm{CMat}_{U,V}(p)) \geq \mathrm{rank}(\mathrm{CMat}_{U,V}(p_{j^*})) + 1. \tag{8.3}$$

Given j^* as in Claim 8.9, we conclude by induction that

$$\mathrm{rank}(\mathrm{CMat}_{U,V}(p)) \geq \mathrm{rank}(\mathrm{CMat}_{U,V}(p_{j^*})) + 1 \geq (t-1) + 1 + 1 = t + 1.$$

To see that p_{j^*} belongs to the vanishing ideal of RFE_{l-1}^{l-2}, we use Theorem 1.8. Note that p_{j^*} is homogeneous, just like p. Condition 1 of Theorem 1.8 is satisfied by p_{j^*} since it it is satisfied by p, and all of k, l, n, and the degree of p_{j^*} are one less. Given K and L as in condition 2 of Theorem 1.8, we have

$$\partial_L p_{j^*}\big|_{K\leftarrow 0} = p_{d^*} \tag{8.4}$$

where d^* is the degree pattern with domain $K \cup L \cup \{j^*\}$ that has $d_j^* = 1$ for $j \in L \cup \{j^*\}$ and $d_j^* = 0$ for $j \in K$. Since $p \in \mathrm{Van}[\mathrm{RFE}_l^{l-1}]$, the contrapositive of the Zoom Lemma applied to p with $K' = K \cup \{j^*\}$, $L' = L \cup \{j^*\}$, d^*, says that (8.4) is zero upon the substitution (1.7). So $p_{j^*} \in \mathrm{Van}[\mathrm{RFE}_{l-1}^{l-2}]$ by Theorem 1.8. This concludes the proof of Lemma 8.5 modulo the proof of Claim 8.9. □

Proof of Claim 8.9. Let $U' \subseteq U$ be the indices of variables x_i such that p depends on x_i, and similarly define $V' \subseteq V$. We first consider the possibility that (8.3) fails for every $j^* \in V'$. We show that this can only happen when $|V'| < |U'|$. A symmetric argument shows that if (8.3) fails for all $j^* \in U'$, then it must be that $|U'| < |V'|$. As both inequalities cannot simultaneously occur, this guarantees the existence of the desired j^*.

Suppose that (8.3) fails for each $j^* \in V'$. Observe that the column of $\mathrm{CMat}_{U,V}(p_{j^*})$ corresponding to a monomial m equals the column of $\mathrm{CMat}_{U,V}(x_{j^*}p_{j^*})$ corresponding to the monomial $x_{j^*}m$; all other columns of $\mathrm{CMat}_{U,V}(x_{j^*}p_{j^*})$ are zero. The matrix $\mathrm{CMat}_{U,V}(x_{j^*}p_{j^*})$ can also be formed from $\mathrm{CMat}_{U,V}(p)$ by zeroing out all the columns indexed by subsets that do not contain j^* (corresponding to multilinear monomials not involving x_{j^*}). The failure of (8.3) for j^* implies that $\mathrm{CMat}_{U,V}(p_{j^*})$ has the same rank as $\mathrm{CMat}_{U,V}(p)$, which is to say that the columns of $\mathrm{CMat}_{U,V}(p)$ indexed by subsets that contain j^* span *all* the columns of $\mathrm{CMat}_{U,V}(p)$. Going block by block, this implies that for every block C_d of $C = \mathrm{CMat}_{U,V}(p)$, the columns within C_d that are indexed by subsets containing j^* span all the columns of C_d. This goes for every $j^* \in V'$, as we are assuming that (8.3) fails for all of them.

Let d be minimal such that $C_d \neq 0$, i. e., such that p has a monomial depending on exactly d variables indexed by U. We have $d \geq t \geq 1$ and $C_{d-1} = 0$. The entries of C_d appear in the linear equations (8.1) given in Proposition 8.8, either with entries from C_{d-1} or from C_{d+1}. Since C_{d-1} is zero, the equations involving C_{d-1} and C_d simplify to equations on C_d only. Namely, for every $I \subseteq U$ with $|I| = d-1$, every $J \subseteq V$ with $|J| = l - (d-1)$, and every $j^* \in I \cup J$, equation (8.1) simplifies to

$$\sum_{i \in U \setminus I} \frac{C_d(\{i\} \cup I, J)}{a_i - a_{j^*}} = 0. \tag{8.5}$$

For any fixed $i \in U \setminus U'$, all entries of the form $C_d(\{i\} \cup I, J)$ are zero. Thus, we can restrict the range of i in (8.5) from $U \setminus I$ to $U' \setminus I$:

$$\sum_{i \in U' \setminus I} \frac{C_d(\{i\} \cup I, J)}{a_i - a_{j^*}} = 0. \tag{8.6}$$

Since $C_d \neq 0$, there is at least one fixed I for which not all entries of the form $C_d(\{i\} \cup I, J)$ are zero as i and J vary. Let I^* be such an I, and let C_d^* denote the submatrix of C_d that consists of all entries of the form $C_d(\{i\} \cup I^*, J)$ as i and J vary. For every $J \subseteq V$ with $|J| = l - (d-1)$ and every $j^* \in I^* \cup J$, we have

$$\sum_{i \in U' \setminus I^*} \frac{C_d^*(\{i\} \cup I^*, J)}{a_i - a_{j^*}} = 0. \tag{8.7}$$

For each $j^* \in V'$, consider the equations (8.7) where J ranges over all subsets of V of size $|J| = l - (d-1)$ that contain j^*. Observe that the coefficients $\frac{1}{a_i - a_{j^*}}$ in (8.7) are independent of the choice of J. We argued that the columns of C_d indexed by subsets J that contain j^* span all columns of C_d. The same holds for C_d^*, as C_d^* is obtained from C_d by removing rows. It follows that (8.7) holds for *every* subset J of V of size $l - (d-1)$ (not just the ones containing j^*).

In particular, consider any one nonzero column of C_d^*. The column represents a nontrivial solution to the homogeneous system (8.7) of $|V'|$ linear equations (one for each choice of $j^* \in V'$) in $|U' \setminus I^*|$ unknowns (one for each $i \in U' \setminus I^*$). The coefficient matrix $[\frac{1}{a_i - a_{j^*}}]$ is a Cauchy matrix, which is well-known to have full rank. In order for there to be a nontrivial solution, the number of equations must be strictly less than the number of unknowns. In other words, we have $|V'| < |U' \setminus I^*| \leq |U'|$, as desired. □

9 Alternating Algebra Representation

In this section we present in greater detail the alternating algebra-based representation of (multilinear) polynomials suited to studying the vanishing ideal of RFE. Subsection 9.1 expands the informal discussion from the overview, describing the representation and characterization for the setting when $l = 1$, $k = 0$, and degree $d = 2$. Subsection 9.2 provides a brief introduction to alternating algebra suited to our purpose. Subsection 9.3 formalizes the discussion from Subsection 9.1 and extends it to the case of multilinear polynomials for general k, l, and d.

9.1 Basic case

For the purposes of this subsection, we fix the parameters $k = 0$, $l = 1$, and $d = 2$. That is to say, we are studying which degree-2 polynomials belong to the vanishing ideal for RFE_1^0.

In Theorem 1.3, we proved that the polynomials $\mathrm{EVC}_1^0[i_1, i_2, i_3]$ as i_1, i_2, i_3 range over $[n]$ generate $\mathrm{Van}[\mathrm{RFE}_1^0]$. As these generators are all homogeneous degree-2 polynomials, a degree-2 polynomial p is in the ideal if and only if it is a linear combination of instantiations of EVC_1^0.

Consider the generator when expanded as a linear combination of monomials:

$$\mathrm{EVC}_1^0[i_1, i_2, i_3] = \begin{vmatrix} a_{i_1} & 1 \\ a_{i_2} & 1 \end{vmatrix} x_{i_1} x_{i_2} + \begin{vmatrix} a_{i_3} & 1 \\ a_{i_1} & 1 \end{vmatrix} x_{i_3} x_{i_1} + \begin{vmatrix} a_{i_2} & 1 \\ a_{i_3} & 1 \end{vmatrix} x_{i_2} x_{i_3}.$$

We represent it graphically by creating a vertex $v_i \in V$ for each variable x_i, an undirected edge for each monomial, and assigning to each edge a weight equal to the coefficient of that monomial:

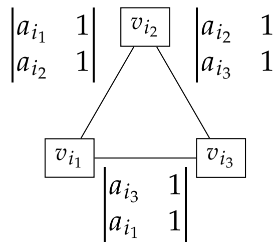

Observe that the coefficient of $x_{i_1}x_{i_2}$ has no dependence on a_{i_3}. In particular, as i_3 varies, the coefficient of $x_{i_1}x_{i_2}$ in $\mathrm{EVC}_1^0[i_1, i_2, i_3]$ does not change. In any other instantiation of EVC_1^0 involving both i_1 and i_2, the coefficient is either the same, or else differs by a sign, according to whether i_1 or i_2 precedes the other in the determinant. A similar pattern holds with respect to all other monomials. This suggests we can modify the graphical representation by rescaling the weights on edges and suppress the dependence on the abscissas. To capture the signs, we use oriented edges. More precisely, for each edge $\{v_{i_1}, v_{i_2}\}$, we consider either of its two orientations, say $v_{i_1} \to v_{i_2}$, and then divide its coefficient by $\begin{vmatrix} a_{i_1} & 1 \\ a_{i_2} & 1 \end{vmatrix}$. Note that considering the opposite orientation coincides with flipping the sign of the scaling factor. With these changes, $\mathrm{EVC}_1^0[i_1, i_2, i_3]$ may be drawn in any of the following ways (among others).

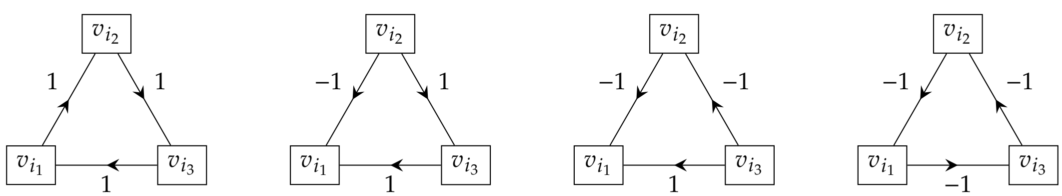

While different choices of edge orientations lead to different illustrations, any one illustration can be transformed into any other by considering edges in opposite orientations as needed, and flipping the sign of each associated coefficient. By identifying each edge in one orientation with the negative of itself in the opposite orientation, we can view all the illustrations as renditions of the same underlying object.

In general, we can represent any degree-2 homogeneous multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ in a similar way: For each monomial $x_{i_1}x_{i_2}$ create an oriented edge $v_{i_1} \to v_{i_2}$ and set the weight of the edge to be the coefficient of $x_{i_1}x_{i_2}$ in p divided by $\begin{vmatrix} a_{i_1} & 1 \\ a_{i_2} & 1 \end{vmatrix}$. The representation determines the polynomial: simply undo the scaling on each edge, and read off a linear combination of monomials. Note moreover that this graphical representation is linear in the polynomial: adding or rescaling polynomials coincides with adding or rescaling coefficients on the edges.

Observe that, in every graphical representation of $\mathrm{EVC}_1^0[i_1, i_2, i_3]$, at every vertex, the sum of the coefficients on edges oriented out of that vertex equals the sum of the coefficients on edges oriented in to that vertex. Indeed, we can interpret $\mathrm{EVC}_1^0[i_1, i_2, i_3]$ as a *circulation* in which one unit of flow travels around a simple 3-cycle $v_{i_1} \to v_{i_2} \to v_{i_3} \to v_{i_1}$. The coefficient on an oriented edge $v_1 \to v_2$ measures how much flow is traveling in the direction $v_1 \to v_2$, with negatives representing flow in the opposite direction. That the sum of coefficients on outgoing edges equals the sum of coefficients on incoming edges reflects the defining property of a circulation, namely that the *conservation law* holds at every vertex: the total flow in equals the total flow out.

Conservation is maintained under linear combinations. Since every degree-2 polynomial p in $\mathrm{Van}[\mathrm{RFE}^0_1]$ is a linear combination of instantiations of EVC^0_1, the representation of p also satisfies the conservation law at every vertex, i. e., the representation of p is a circulation. Thus, conservation is a necessary condition for membership in $\mathrm{Van}[\mathrm{RFE}^0_1]$.

Conservation is *sufficient* for ideal membership, as well. By definition, conservation at every vertex means that the representation is a circulation. By the well-known flow decomposition theorem (see, e. g., [5, p. 80-81]), every circulation can be decomposed into a superposition of circulations around simple cycles. A unit circulation around a simple cycle can be decomposed into a sum of unit circulations around 3-cycles; this is depicted for a 5-cycle below, where each edge indicates unit flow:

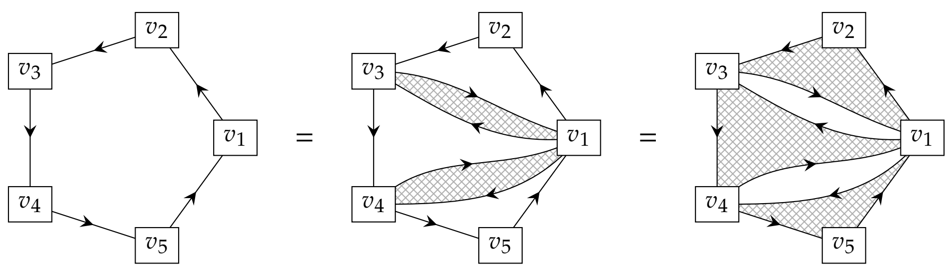

The basis of the first equality in the above figure is that a unit flow $v_1 \to v_3$ cancels with a unit flow $v_3 \to v_1$, and similar for v_4 in lieu of v_3. Thus, conservation implies that we have a linear combination of unit circulations on 3-cycles, i. e., a linear combination of instantiations of EVC^0_1.

In summary, a multilinear homogeneous degree-2 polynomial is in $\mathrm{Van}[\mathrm{RFE}^0_1]$ *if and only if* its graphical representation satisfies the conservation law at every vertex. This is the representation and ideal membership characterization in the basic setting with $k = 0$, $l = 1$, and $d = 2$ for multilinear homogeneous polynomials. Note that, in this basic setting, the multilinear homogeneous degree-2 case represents the core of the problem. The remaining cases contain a univariate monomial, and are outside of RFE^0_1 by Proposition 5.1.

9.2 Alternating algebra

In order to generalize Subsection 9.1, we need to be able to discuss higher-dimensional analogues of "flow" and "circulation", as well as appropriately-generalized notions of "conservation." Suited to this purpose is the language of *alternating algebra*. Alternating algebra was introduced in the 1800s by Hermann Grassmann [22, 23] and is the formalism underlying differential geometry and its applications to physics. We give a brief introduction to alternating algebra here, tailored toward our purposes.

For each $i \in [n]$, we create a fresh vertex $v_i \in V$, which corresponds to the variable x_i. The alternating algebra provides a multiplication, denoted $\wedge$, that can be thought of as a constructor to make *oriented simplices* out of these vertices. For example, the $\wedge$-product of v_1 with v_2, written $v_1 \wedge v_2$, encodes the simplex with vertices v_1 and v_2 in a particular orientation; $v_2 \wedge v_1$ encodes the same simplex with the opposite orientation. When $v_1 = v_2$, $v_1 \wedge v_2$ is defined to be zero. $\wedge$-multiplication is associative. Rather than being commutative, the $\wedge$-product is *anti-commutative* in the sense that $v_1 \wedge v_2 = -v_2 \wedge v_1$. In this way the order of the vertices in the product encodes an orientation. There are only ever two orientations. In a larger product such as $v_1 \wedge v_2 \wedge v_3$, we have

$$\begin{array}{rcrcr} & & v_1 \wedge v_2 \wedge v_3 & = & -v_1 \wedge v_3 \wedge v_2 \\ & = & v_3 \wedge v_1 \wedge v_2 & = & -v_3 \wedge v_2 \wedge v_1 \\ & = & v_2 \wedge v_3 \wedge v_1 & = & -v_2 \wedge v_1 \wedge v_3. \end{array}$$

In general, permuting the vertices in a $\wedge$-product by an even permutation has no effect, while permuting by an odd permutation flips the sign. Any $\wedge$-product that uses the same vertex more than once is zero.

We can formally extend $\wedge$-multiplication to linear combinations of vertices in V. Denote U to be the $\mathbb{F}$-vector space with basis V. The $\wedge$-multiplication extends to U by being *distributive*. Overall, $\wedge$-multiplication has the following defining properties, for any $u_1, u_2, u_3 \in U$:

- *Associativity:* $u_1 \wedge (u_2 \wedge u_3) = (u_1 \wedge u_2) \wedge u_3$.
- *Distributivity:* $u_1 \wedge (u_2 + u_3) = u_1 \wedge u_2 + u_1 \wedge u_3$.
- *Alternation:* $u_1 \wedge u_1 = 0$.

The alternation property implies anti-commutativity[5]: $u_1 \wedge u_2 = -u_2 \wedge u_1$. The alternating algebra consists of all formal linear combinations of $\wedge$-products of vertices from V, or equivalently, of elements from U. We denote the underlying universe as follows.

Definition 9.1 (space of oriented simplices)**.** For each $t \in \mathbb{N}$, we let

$$\Lambda^t(U) \doteq \operatorname{span}(u_1 \wedge \cdots \wedge u_t : u_1, \ldots, u_t \in U)$$

denote the space of linear combinations of t-vertex oriented simplices. For a set of indices T, we write $u^T \doteq \bigwedge_{i \in T} u_i$, with the convention that the indices are listed in increasing order.

The distributivity and anti-commutativity properties of $\wedge$ imply that

$$\Lambda^t(U) = \operatorname{span}(u^{[t]} : u_1, \ldots, u_t \text{ are distinct elements in } V),$$

which justifies the reference to t-vertex simplices. The properties also imply that changing the order of the vertices in the wedge product yields the same element up to a sign, namely the sign of the underlying permutation. This justifies the reference to orientation, where there are two possible orientations. To emphasize, the t in $\Lambda^t(U)$ counts the number of vertices in the simplices; this is one more than the usual notion of dimension of a simplex. For $t = 0$, we have a distinct simplex corresponding to the empty product, denoted 1, which is an identity for $\wedge$. Note that not every element in $\Lambda^t(U)$ can be expressed in the form $u_1 \wedge \cdots \wedge u_t$.

To connect this with Subsection 9.1, recall the graphical depiction of $\mathrm{EVC}^0_1[i_1, i_2, i_3]$:

[5]Alternation and anti-commutativity are equivalent provided the characteristic of the field differs from 2.

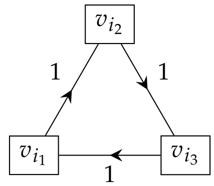

Adopting the convention that an arrow $v_1 \to v_2$ is $v_1 \wedge v_2$ (and so an arrow $v_2 \to v_1$ is $v_2 \wedge v_1 = -v_1 \wedge v_2$), we can alternatively express the above as

$$v_{i_1} \wedge v_{i_2} \;+\; v_{i_2} \wedge v_{i_3} \;+\; v_{i_3} \wedge v_{i_1}.$$

In general, the graphical representation of a homogeneous degree-2 multilinear polynomial is some linear combination of 2-vertex oriented simplices. When we go to higher-degree polynomials, we make use of oriented simplices with more vertices.

To express conservation, we introduce *boundary maps*, which are parametrized by a linear weight function $w : U \to \mathbb{F}$. The boundary map ∂_w is a linear map that sends each simplex to a linear combination of its boundary faces (and the empty simplex to zero) according to a formula reminiscent of the minor expansion of a determinant along a column consisting of the values of w.

Definition 9.2 (boundary map)**.** For any linear function $w : U \to \mathbb{F}$, the boundary map with weight function w is the linear map $\partial_w : \bigoplus_{t=0}^{n} \Lambda^t(U) \to \bigoplus_{t=0}^{n} \Lambda^t(U)$ realizing

$$u_1 \wedge \cdots \wedge u_t \mapsto \sum_{i=1}^{t} (-1)^{i+1} w(u_i)(u_1 \wedge \cdots \wedge u_{i-1} \wedge u_{i+1} \wedge \cdots \wedge u_t) \tag{9.1}$$

for all $u_1, \ldots, u_t \in U$.

The boundary map ∂_w is well-defined. To see this, note that the sign factor $(-1)^{i+1}$ in (9.1) ensures well-definedness of the restriction to vertices, i. e., for $u_1, \ldots, u_t \in V$. This is because changing the order of the vertices on the left-hand side results in the correct sign change on the right-hand side. The linearity of w then guarantees that the linear extension of the restriction to vertices coincides with (9.1). For each $t \geq 1$, $\partial_w(\Lambda^t(U)) \subseteq \Lambda^{t-1}(U)$, while $\partial_w(\Lambda^0(U)) = \{0\}$.

In the simplest case, w is the function that is 1 on every $v \in V$. In this case, the boundary of some 2-vertex simplex is given by

$$\partial_1(v_1 \wedge v_2) = v_2 - v_1.$$

In particular, $v_1 \wedge v_2$ contributes -1 toward v_1 and $+1$ toward v_2. This coincides with the contribution of the edge $v_1 \to v_2$ toward the net flow into the vertices v_1 and v_2. In exactly this way, conservation is identified with having a *vanishing boundary*. Note also that for this choice of weight function

$$\partial_1(v_1 \wedge v_2 \wedge v_3) = v_2 \wedge v_3 \;-\; v_1 \wedge v_3 \;+\; v_1 \wedge v_2 = v_1 \wedge v_2 \;+\; v_2 \wedge v_3 \;+\; v_3 \wedge v_1.$$

Thus, unit circulations on 3-cycles are in one-to-one and onto correspondence with the images under ∂_1 of oriented 3-simplices on the vertices. By the decomposition discussed in the Subsection 9.1, it follows that circulations are in one-to-one and onto correspondence with the elements of $\partial_1(\Lambda^3(U))$. This means that $\partial_1(\Lambda^3(U)) = \ker(\partial_1) \cap \Lambda^2(U)$.

In general, for every linear $w : U \to \mathbb{F}$

$$\operatorname{im}(\partial_w) = \ker(\partial_w), \tag{9.2}$$

or equivalently, $\partial_w(\Lambda^t(U)) = \ker(\partial_w) \cap \Lambda^{t-1}(U)$ for every $t \in [n]$. This key relationship implies that taking the same boundary multiple times always vanishes. That is, for any w, $\partial_w \circ \partial_w = 0$, often written as $\partial_w^2 = 0$. Another property is that for any w, w' and $\beta, \beta' \in \mathbb{F}$, $\partial_{\beta w + \beta' w'} = \beta \partial_w + \beta' \partial_{w'}$, which is to say that the boundary maps themselves are linear in w. It follows from these that, for any w, w', $\partial_w \circ \partial_{w'} = -\partial_{w'} \circ \partial_w$. This means that the boundary maps themselves behave like an alternating algebra, with $\circ$ as the multiplication rather than $\wedge$. For any $w_1, \ldots, w_{k+1}$, write $\omega = w_1 \wedge \cdots \wedge w_{k+1}$, and define $\partial_\omega = \partial_{w_{k+1}} \circ \cdots \circ \partial_{w_1}$. That is, $w_1 \wedge \cdots \wedge w_{k+1}$ means apply ∂_{w_1}, then ∂_{w_2}, and so on, up to $\partial_{w_{k+1}}$. The result is well-defined, and we borrow the shorthand notation introduced in Definition 9.1: $w^T \doteq \bigwedge_{j \in T} w_j$, where $T \subseteq [k+1]$ and the indices in the wedge product are taken in increasing order.

The image-kernel relationship (9.2) extends as follows: For any linearly independent $w_1, \ldots, w_{k+1}$,

$$\operatorname{im}(\partial_{w_1 \wedge \ldots \wedge w_{k+1}}) = \bigcap_{r=1}^{k+1} \ker(\partial_{w_r}). \tag{9.3}$$

If $w_1, \ldots w_{k+1}$ are linearly dependent, then $\partial_{w_1 \wedge \ldots \wedge w_{k+1}}$ vanishes. In fact, a further generalization holds and will be useful. We include a proof for completeness. (9.3) corresponds to the special case $\Delta = 0$.

Proposition 9.3 (generalized image-kernel relationship)**.** *For $k, \Delta \in \mathbb{N}$ and any linearly independent linear functions $w_1, \ldots, w_{k+\Delta+1} : U \to \mathbb{F}$*

$$\operatorname*{span}_{\substack{B \subseteq [k+\Delta+1] \\ |B| = k+1}} \operatorname{im}(\partial_{w^B}) = \bigcap_{\substack{B \subseteq [k+\Delta+1] \\ |B| = \Delta+1}} \ker(\partial_{w^B}). \tag{9.4}$$

Proof. Extend $w_1, \ldots, w_{k+\Delta+1}$ to a basis $w_1, \ldots, w_n$ of all linear functions $U \to \mathbb{F}$. We can interpret $w_1, \ldots, w_n$ as a basis of the dual space U^*, and the mapping $(w, u) \mapsto w(u)$ as a bilinear form $U^* \times U \to \mathbb{F}$. This means we can construct a dual basis $u_1, \ldots, u_n \in U$ such that for $i, j \in [n]$, $w_j(u_i)$ is 1 if $i = j$ and 0 if $i \neq j$.

In this particular basis $u_1, \ldots, u_n$, the boundary maps with weight functions w_j take a very simple form: The only term in (9.1) that remains for $w = w_j$ is the one with $i = j$. More generally, for $B, T \subseteq [n]$,

$$\partial_{w^B}(u^T) = \begin{cases} \pm u^{T \setminus B} & \text{if } B \subseteq T \\ 0 & \text{otherwise.} \end{cases} \tag{9.5}$$

With this characterization, we can see that both the span of the images and the intersection of the kernels coincide with

$$\mathrm{span}(u^S : S \subseteq [n] \text{ with } |S \cap [k+\Delta+1]| \leq \Delta).$$

We obtain $\pm u^S$ as $\partial_{w^B}(u^T)$ if and only if $T = B \sqcup S$. Such a choice of T and B with $B \subseteq [k+\Delta+1]$ and $|B| = k+1$ exists if and only if there are at least $k+1$ elements in $[k+\Delta+1] \setminus S$, or equivalently, $|S \cap [k+\Delta+1]| \leq (k+\Delta+1) - (k+1) = \Delta$. This proves the equality for the span of the images.

On the other hand, u^S falls within $\ker(\partial_{w^B})$ if and only if $B \not\subseteq S$. This is case for every $B \subseteq [k+\Delta+1]$ with $|B| = \Delta+1$ if and only if S contains at most Δ elements in $[k+\Delta+1]$. This proves the equality of the intersection of the kernels. □

In the following subsection, we will need an explicit formula for computing $\partial_\omega(u^T)$ for generic $u_1, \ldots, u_n \in U$ and $T \subseteq [n]$. From a concrete perspective, the effect of a single boundary map in Definition 9.2 resembles one level of determinant minor expansion, so composing boundary maps should produce a partially expanded determinant. We formalize that intuition with the following proposition, which characterizes the boundary of a t-simplex after applying k weighted boundaries as a linear combination of $(t-k)$-simplices. Each $(t-k)$-simplex is indexed by a subset J of T.

Proposition 9.4 (composed boundary maps)**.** *Let $w_1, \ldots, w_{k+1} : U \longrightarrow \mathbb{F}$ be linear functions, T a set of indices, and $u_i \in U$ for $i \in T$.*

$$\partial_{w^{[k+1]}}(u^T) = \sum_{\substack{I \sqcup J = T \\ |I| = k+1}} (-1)^{\mathrm{XInv}(I,J)} \cdot \det \left[w_r(u_i) \right]_{i \in I}^{r \in [k+1]} \cdot u^J, \tag{9.6}$$

where in the determinant, the rows from top to bottom and the columns from left to right are in increasing order of index i and r, respectively.

Proof. Observe that $\partial_{w^{[k+1]}}(u^T) \in \Lambda^{t-k-1}(U)$ and, by Definition 9.2, can be written as a linear combination of u^J over all $J \subseteq T$ with $|J| = |T| - k - 1$. It suffices to show that the coefficients of each u^J match the ones given above.

Without loss of generality, let $T = [t]$, as this does not change the relative order of any determinants or $\wedge$-products. Consider the terms formed by iteratively expanding $\partial_{w^{[k+1]}}(u^T) \doteq \partial_{w_1 \wedge \ldots \wedge w_{k+1}}(u^T)$ by Definition 9.2. Each term is in one-to-one correspondence with the choices of i we make in the expansions of Definition 9.2. In particular, the terms that yield u^J correspond to the bijections $\sigma : [k+1] \longrightarrow I$, where $I = T \setminus J$. For a given σ, the corresponding coefficient is equal to

$$(-1)^{(\sum_{i \in I} i) + k + 1 - |\{r, r' \in [k+1] : r' < r, \sigma(r') < \sigma(r)\}|} \prod_{r \in [k+1]} w_r(u_{\sigma(r)}).$$

The $|\{r' < r, \sigma(r') < \sigma(r)\}|$ term accounts for the fact that, when each r is selected, some terms of T may have been previously removed, shifting the relative rank of r. Since for any distinct $r, r' \in [k+1]$, either $\sigma(r') > \sigma(r)$ or $\sigma(r') < \sigma(r)$, we can rewrite $|\{r' < r, \sigma(r') < \sigma(r)\}| = \binom{k+1}{2} - |\{r' < r, \sigma(r') > \sigma(r)\}|$.

As for the term $\sum_{i\in I} i$, writing i as $i = |\{i' \in I : i' \leq i\}| + |\{j \in J : j < i\}|$ and summing over all $i \in I$, we have that $\sum_{i\in I} i = \sum_{r=1}^{k+1} r + \mathrm{XInv}(I, J) = \binom{k+2}{2} + \mathrm{XInv}(I, J)$.

As $\binom{k+2}{2} = \binom{k+1}{2} + k + 1$, we get that the coefficient of u^J equals

$$\sum_{\sigma:[k+1]\to I} (-1)^{\mathrm{XInv}(I,J)+|\{r'<r,\sigma(r')>\sigma(r)\}|+2k+2} \prod_{r\in[k+1]} w_r(u_{\sigma(r)}),$$

and by definition of the determinant and simplifying, this is equal to

$$(-1)^{\mathrm{XInv}(I,J)} \det \left[w_r(u_i)\right]_{i\in I}^{r\in[k+1]} . \qquad \square$$

In the next subsection we will apply Proposition 9.4 with $u_i = v_i$. For the choice of u_i in the proof of Proposition 9.3, the matrix in (9.6) is the identity matrix and thus has determinant 1, which results in (9.5).

9.3 General case

With the notation of alternating algebra in hand, we turn now to generalizing the characterization of $\mathrm{Van}[\mathrm{RFE}_l^k]$ based on the representation of polynomials that we introduced in Subsection 9.1, henceforth the *simplicial representation*. We focus on multilinear polynomials, but the parameters $k, l \in \mathbb{N}$ may be arbitrary. For starters, we still restrict to degree $d = l + 1$. We then generalize to multilinear polynomials of arbitrary degree and present an alternate proof to Theorem 1.8. We end with some thoughts about the non-multilinear case.

As before, we associate each variable x_i with a distinct vertex $v_i \in V$, where $U \doteq \mathrm{span}(V)$ denotes an underlying vector space over $\mathbb{F}$. We view a polynomial as a linear combination of monomials and represent each degree-t multilinear monomial as an oriented simplex with t vertices. The representation makes use of the Vandermonde determinants $\det(A_T)$ for $T \subseteq [n]$, where A_T refers to the notation that we introduced in (3.2) for the Vandermonde matrix built from the abscissas a_i for $i \in T$ in increasing order. The Vandermonde determinant $\det(A_T)$ can be written as the product of pairwise differences:

$$\det(A_T) = \prod_{i,j\in T, i<j} (a_i - a_j). \tag{9.7}$$

In particular, as the abscissas are distinct, $\det(A_T)$ is always nonzero.

Let $v^T \doteq \bigwedge_{i\in T} v_i$, where the indices are listed in increasing order. We represent the monomial $x^T \doteq \prod_{i\in T} x_i$ for $T \subseteq [n]$ by the element $v^T/\det(A_T)$. Formally, we define the following "decoder map," which maps a simplicial representation to the polynomial it represents.

Definition 9.5 (representation). $\rho : \bigoplus_{t=0}^{n} \Lambda^t(U) \to \mathbb{F}[x_1, \ldots, x_n]$ is the linear map extending

$$v^T \;\mapsto\; \det(A_T) \cdot x^T \tag{9.8}$$

for every $T \subseteq [n]$.

Note that (9.8) holds irrespective of the order of the indices, as long as the same order is used for both v^T and $\det(A_T)$. This is because exchanging any two indices changes the sign of both the left-hand side and the determinant on the right-hand side. The mapping ρ induces a vector space isomorphism between $\Lambda^{l+1}(U)$ and the space of multilinear homogeneous degree-$(l+1)$ polynomials.

The strategy for our membership test in Van[RFE] consists of two steps: First express Van[RFE] in terms of ρ and the image of the boundary maps ∂_w, and then apply the (generalized) image-kernel relationship from alternating algebra. In Definition 9.2, w is taken to be a linear function from U to $\mathbb{F}$. As a linear function, w is completely defined by its values on the basis V. By Lagrange interpolation, every function w from V to $\mathbb{F}$ can be viewed as a univariate polynomial of degree less than $n \doteq |V|$ restricted to the abscissas, namely the polynomial interpolating $a_i \mapsto w(v_i)$ for $i \in [n]$.

Definition 9.6 (degree of boundary map)**.** Let $w : U \to \mathbb{F}$ be linear and $a_1, \ldots, a_n$ be distinct elements of $\mathbb{F}$. We say that w is *interpolated* by $q \in \mathbb{F}[\alpha]$ if $w(v_i) = q(a_i)$ for $i \in [n]$. We say that w is of degree d if w is interpolated by a degree-d polynomial q.

Furthermore, given fixed $a_1, \ldots, a_n$, the correspondence between a weight function w and its interpolating polynomial q forms an isomorphism; if w_1, w_2 are interpolated by q_1, q_2, then $w_1 + w_2$ is interpolated by $q_1 + q_2$, and cw_1 is interpolated by cq_1. From now on, we directly refer to a weight function by the polynomial in $\mathbb{F}[\alpha]$ that interpolates it. We will be interested in the boundaries that are weighted by low-degree polynomials.

Multilinear case for degree $d = l + 1$. In the case of degree $d = l + 1$, the first step of our approach boils down to finding a simplicial representation for the generators EVC_l^k. We do so using composed boundary maps of degree at most k.

Lemma 9.7. *For any $k, l \in \mathbb{N}$ and $S \subseteq [n]$, $|S| = k + l + 2$,*

$$\mathrm{EVC}_l^k[S] = \rho\left(\partial_{\alpha^k \wedge \cdots \wedge \alpha^0}\left(v^S\right)\right). \tag{9.9}$$

That is, EVC_l^k is the polynomial formed from a given $(k + l + 2)$-vertex simplex by iteratively applying to it the $k + 1$ boundaries weighted by $\alpha^k, \alpha^{k-1}, \ldots, \alpha^0$ respectively, where α^r stands for the weight function interpolated by the polynomial α^r.

Proof. Using our notation, the explicit expression (3.3) in Proposition 3.4 can be rewritten as

$$\mathrm{EVC}_l^k[S] = \sum_{\substack{K \sqcup L = S \\ |K| = k+1}} (-1)^{\mathrm{XInv}(K,L)} \cdot \det(A_K) \cdot \det(A_L) \cdot x^L. \tag{9.10}$$

For the right-hand side, we use Proposition 9.4 to get:

$$\begin{aligned}\rho\left(\partial_{\alpha^k\wedge\cdots\wedge\alpha^0}\left(v^S\right)\right) &= \sum_{\substack{K\sqcup L=S\\|K|=k+1}}(-1)^{\mathrm{XInv}(K,L)}\cdot\det(A_K)\cdot\rho\left(v^L\right)\\ &= \sum_{\substack{K\sqcup L=S\\|K|=k+1}}(-1)^{\mathrm{XInv}(K,L)}\cdot\det(A_K)\cdot\det(A_L)\cdot x^L.\end{aligned}$$

The sum is identical to (9.10). □

Lemma 9.7 yields the following characterization of the part of $\mathrm{Van}[\mathrm{RFE}_l^k]$ of degree $l+1$. We state it in a format to which we can directly apply the image-kernel relationship (9.3).

Corollary 9.8. *For any $k, l \in \mathbb{N}$, the set of polynomials of degree $l+1$ in $\mathrm{Van}[\mathrm{RFE}_l^k]$ is given by*

$$\rho(\partial_{\alpha^k\wedge\cdots\wedge\alpha^0}(\Lambda^{k+l+2}(U))).$$

Proof. Since every degree-$(l+1)$ polynomial p in $\mathrm{Van}[\mathrm{RFE}_l^k]$ is a linear combination of instantiations of EVC_l^k, Lemma 9.7 allows to us to express the subset in $\mathrm{Van}[\mathrm{RFE}_l^k]$ as

$$\operatorname*{span}_{\substack{S\subseteq[n]\\|S|=k+l+2}}\ \rho(\partial_{\alpha^k\wedge\cdots\wedge\alpha^0}(v^S)).$$

The result follows by linearity and the fact that $U = \mathrm{span}(V)$. □

The image-kernel relationship (9.3) then leads to the following membership test. Recall that ρ induces an isomorphism from the space of $(l+1)$-vertex oriented simplices $\Lambda^{l+1}(U)$ to the set of multilinear polynomials of degree $l+1$, so ρ^{-1} is well-defined on multilinear polynomials.

Theorem 9.9. *Let $k, l \in \mathbb{N}$. For any multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$ of degree $l+1$, $p(\mathrm{RFE}_l^k) = 0$ if and only if p is homogeneous of degree $l+1$ and*

$$\partial_w(\rho^{-1}(p)) = 0$$

for every weight function w of degree at most k.

Proof. The criterion in Corollary 9.8 can be rewritten as

$$\rho^{-1}(p) \in \partial_{\alpha^k\wedge\cdots\wedge\alpha^0}(\Lambda^{k+l+2}(U)).$$

By Proposition 9.3, this is equivalent to

$$\rho^{-1}(p) \in \left(\bigcap_{r=0}^{k}\ker(\partial_{\alpha^r})\right)\cap\Lambda^{l+1}(U).$$

The intersection with $\Lambda^{l+1}(U)$ means that p is homogeneous of degree $l+1$. For such polynomials p, we have that $p(\mathrm{RFE}_l^k) = 0$ if and only if $\partial_{\alpha^r}(\rho^{-1}(p)) = 0$ for $r = 0, \ldots, k$, which by linearity is equivalent to $\partial_w(\rho^{-1}(p)) = 0$ for all weight functions w of degree at most k. □

Theorem 9.9 states that a multilinear polynomial p of degree $l+1$ is in the vanishing ideal of RFE_l^k if and only if it is homogeneous of degree $l+1$ and the simplicial representation of p satisfies conservation with respect to all degree-k boundaries. This is the representation and ideal membership characterization for such polynomials for general k and l in the special case of degree $d = l+1$. As we will argue in Proposition 9.14, the characterization coincides with the membership test from Theorem 1.8 for multilinear polynomials of degree $l+1$.

In Section 9.1 we considered the special case with $k = 0$ and $l = 1$. In that basic setting, the only weight functions of degree k are the constant functions, and only $w \equiv 1$ needs to be considered in Theorem 9.9. The resulting criterion is exactly the conservation criterion that we developed in Section 9.1.

Note that the restriction in Theorem 9.9 to *multilinear* polynomials p is just to ensure that $\rho^{-1}(p)$ is well-defined. For polynomials of degree $l+1$ that are not multilinear, one could interpret the non-existence of $\rho^{-1}(p)$ as not satisfying the criterion. This is consistent with Proposition 5.1, which implies that polynomials of degree $l+1$ that are not multilinear are automatically outside $\mathrm{Van}[\mathrm{RFE}_l^k]$ since they necessarily have a monomial supported on l or fewer variables.

Through Lemma 9.7, the property that $\partial_w(\rho^{-1}(\mathrm{EVC}_l^k)) = 0$ for every weight function w of degree at most k can be viewed as an application of $\partial_{w\wedge\alpha^k\wedge\cdots\wedge\alpha^0} = 0$ to $\Lambda^t(U)$ with $t = k+l+2$. The equations (1.6) follow in a similar way from an application with $t = k+l+3$.

Multilinear case of arbitrary degree. The two-step approach underlying Theorem 9.9 extends to multilinear polynomials of higher degrees. Whereas in the special case of degree $d = l+1$ we only needed simplicial representations for $\mathrm{EVC}_l^k[S]$ in the first step, we now need them for polynomials of the more general form $\mathrm{EVC}_l^k[S]\cdot x^M$ where $M \subseteq [n]$ is disjoint from S. We can handle the additional term x^M in Lemma 9.7 by including a multiplicative factor

$$\mu_M(\alpha) \doteq \prod_{j\in M}(\alpha - a_j) \tag{9.11}$$

in each of the weight functions. The extra factor acts as a masking term and ensures that in the expansions of (9.1) the terms with $i \in M$ vanish, so under ρ the factor x^M remains.

Lemma 9.10. *For any $k, l \in \mathbb{N}$, $S \sqcup M \subseteq [n]$ with $|S| = k+l+2$, and $\mu_M(\alpha) \doteq \prod_{j\in M}(\alpha - a_j)$,*

$$\mathrm{EVC}_l^k[S]\cdot x^M = \frac{\det(A_S)}{\det(A_{S\sqcup M})}\cdot \rho(\partial_{\mu_M(\alpha)\alpha^k\wedge\cdots\wedge\mu_M(\alpha)\alpha^0}(v^{S\sqcup M})). \tag{9.12}$$

Proof. Expand $\partial_{\mu_M(\alpha)\alpha^k\wedge\cdots\wedge\mu_M(\alpha)\alpha^0}(v^{S\sqcup M})$ by Proposition 9.4. Notice that the only nonzero terms in the expansion correspond to subsets J that contain M. Substituting $I \leftarrow K$ and $J \leftarrow L \sqcup M$, and factoring out the $\mu_M(a_i)$ terms from the determinant, we can write

$$\partial_{\mu_M(\alpha)\alpha^k\wedge\cdots\wedge\mu_M(\alpha)\alpha^0}(v^{S\sqcup M}) = \sum_{\substack{K\sqcup L=S\\|K|=k+1}} (-1)^{\mathrm{XInv}(K,L\sqcup M)}\cdot\left(\prod_{i\in K}\mu_M(a_i)\right)\cdot\det(A_K)\cdot v^{L\sqcup M}.$$

Applying ρ yields

$$\rho(\partial_{\mu_M(\alpha)\alpha^k \wedge \cdots \wedge \mu_M(\alpha)\alpha^0}(v^{S \sqcup M})) = \sum_{\substack{K \sqcup L = S \\ |K| = k+1}} (-1)^{\mathrm{XInv}(K, L \sqcup M)} \cdot \left(\prod_{i \in K} \mu_M(a_i) \right) \cdot \det(A_K) \cdot \det(A_{L \sqcup M}) \cdot x^{L \sqcup M}. \tag{9.13}$$

Applying (9.7) to $T = L \sqcup M$, $T = L$, and $T = M$, rearranging terms, and remembering that A_T takes rows in increasing index, we obtain

$$\det(A_{L \sqcup M}) = (-1)^{\mathrm{XInv}(L,M)} \cdot \left(\prod_{i \in L, j \in M} (a_i - a_j) \right) \cdot \det(A_L) \cdot \det(A_M). \tag{9.14}$$

We can expand $(-1)^{\mathrm{XInv}(K, L \sqcup M)}$ as the product $(-1)^{\mathrm{XInv}(K,L)}(-1)^{\mathrm{XInv}(K,M)}$ because $\mathrm{XInv}(K, L \sqcup M)$ equals the sum $\mathrm{XInv}(K, L) + \mathrm{XInv}(K, M)$. By the definition of μ_M, we can expand $\prod_{i \in K} \mu_M(a_i)$ as $\prod_{i \in K, j \in M}(a_i - a_j)$. Those expansions and (9.14) allow us to write the summand on the right-hand side of (9.13) as

$$(-1)^{\mathrm{XInv}(K,L)}(-1)^{\mathrm{XInv}(K,M)}(-1)^{\mathrm{XInv}(L,M)} \cdot \left(\prod_{i \in K \sqcup L, j \in M} (a_i - a_j) \right) \cdot \det(A_K) \cdot \det(A_L) \cdot \det(A_M) \cdot x^{L \sqcup M}$$

Using the similar fact as above that $(-1)^{\mathrm{XInv}(K \sqcup L, M)} = (-1)^{\mathrm{XInv}(K,M)}(-1)^{\mathrm{XInv}(L,M)}$, recalling that $K \sqcup L = S$, and pulling out the terms independent of the choice of K, we obtain

$$\begin{aligned}
&\rho(\partial_{\mu_M(\alpha)\alpha^k \wedge \cdots \wedge \mu_M(\alpha)\alpha^0}(v^{S \sqcup M})) \\
&= (-1)^{\mathrm{XInv}(S,M)} \cdot \left(\prod_{i \in S, j \in M} (a_i - a_j) \right) \cdot \det(A_M) \cdot x^M \sum_{\substack{K \sqcup L = S \\ |K| = k+1}} (-1)^{\mathrm{XInv}(K,L)} \cdot \det(A_K) \cdot \det(A_L) \cdot x^L \\
&= \frac{\det(A_{S \sqcup M})}{\det(A_S)} \cdot x^M \sum_{\substack{K \sqcup L = S \\ |K| = k+1}} (-1)^{\mathrm{XInv}(K,L)} \cdot \det(A_K) \cdot \det(A_L) \cdot x^L,
\end{aligned}$$

where the last step applies (9.14) with $L \leftarrow S$. By Proposition 3.4, this establishes the result. □

The multilinear elements in $\mathrm{Van}[\mathrm{RFE}_l^k]$ are exactly the linear combinations of terms of the form (9.12) where $S \subseteq [n]$ ranges over subsets of size $k + l + 2$ and $M \subseteq [n]$ over subsets disjoint with S. In order to obtain a simpler characterization of the same type, as well as one to which we can apply the generalized image-kernel relationship, we show that we can replace the weight functions on the right-hand side of (9.12) by generic weight functions of the same degree or by Lagrange interpolants with respect to a subset of abscissas of size one more.

Proposition 9.11. *Let $k + 1, m, t \in \mathbb{N}$ with $t \geq k + 1$, $v \in \Lambda^t(U)$, and $N \subseteq [n]$ with $|N| = k + m + 1$. Let $L_{N,j}$ for $j \in N$ denote the Lagrange interpolants for the subset of abscissas $\{a_i\}_{i \in N}$, i.e., $L_{N,j}$*

denotes the unique univariate polynomial of degree at most $|N|-1$ *satisfying* $L_{N,j}(a_i) = 1$ *for* $i = j$ *and* $L_{N,j}(a_i) = 0$ *for* $i \in N \setminus \{j\}$. *For all weight functions* $w_1, \ldots, w_{k+1}$ *of degree at most* $k+m$,

$$\operatorname*{span}_{\substack{M \subseteq N \\ |M|=m}} \partial_{\mu_M(\alpha)\alpha^k \wedge \cdots \wedge \mu_M(\alpha)\alpha^0}(\nu) = \operatorname*{span}_{\substack{w_1,\ldots,w_{k+1} \in \mathbb{F}[\alpha] \\ \deg(w_1),\ldots \deg(w_{k+1}) \leq k+m}} \partial_{w^{[k+1]}}(\nu) = \operatorname*{span}_{\substack{B \subseteq N \\ |B|=k+1}} \partial_{L_N^B}(\nu). \tag{9.15}$$

Some explanation of the compact notation on the right-hand side of (9.15) is in order. First, we use $L_{N,j}$ to differentiate with the notation L_j for Lagrange interpolants that we introduced in Definition 1.1, where L_j corresponds to $L_{[n],j}$. Second, for a subset $B \subseteq N$, we write L_N^B as a shorthand for $\bigwedge_{j \in B} L_{N,j}$, where the indices in the wedge product are taken in increasing order. Finally, in the composed boundary operator $\partial_{L_N^B}$, the Lagrange interpolant $L_{N,j}$ represents the weight function interpolated by $L_{N,j}$ as in Definition 9.6.

Proof. The inclusion $\subseteq$ of the first equality in (9.15) follows because the weight functions $\mu_M(\alpha)\alpha^r$ for $r \in \{0, \ldots, k\}$ have degree at most $k + |M| = k + m$.

To argue the inclusion $\subseteq$ of the second equality in (9.15), note that the Lagrange interpolants $L_{N,j}$ for $j \in N$ are linearly independent and that there are as many of them as the dimension of the space of polynomials of degree at most $|N| - 1 = k + m$, so they form a basis for that space. In particular, we can write all weight functions $w_1, \ldots, w_{k+1}$ of degree at most $k+m$ as linear combinations of the Lagrange interpolants $L_{N,j}$, $j \in N$. By the distributivity and antisymmetry of the wedge product, this implies that

$$\partial_{w^{[k+1]}}(\nu) \in \operatorname*{span}_{\substack{B \subseteq N \\ |B|=k+1}} \partial_{L_N^B}(\nu).$$

It remains to argue that the right-most side of (9.15) is included in the left-most side. Fix a subset $B \subseteq N$ of size $|B| = k+1$. Since the polynomials $L_{N,j}$ for $j \in B$ individually have roots in all but one element of $\{a_i\}_{i \in N}$, they collectively have common roots among exactly $|N| - |B| = m$ of these abscissas, which form a set $M \subseteq N$. Each $L_{N,j}$ can therefore be written as the product of μ_M and a polynomial of degree at most k, or equivalently, as a linear combination of $\mu_M(\alpha)\alpha^k, \ldots, \mu_M(\alpha)\alpha^0$. Once again, by the distributivity and antisymmetry of the wedge product, we have that

$$\partial_{L_N^B}(\nu) \in \operatorname*{span}_{\substack{M \subseteq N \\ |M|=m}} \partial_{\mu_M(\alpha)\alpha^k \wedge \cdots \wedge \mu_M(\alpha)\alpha^0}(\nu). \qquad \square$$

The first equality in (9.15) connects weight functions as on the right-hand side of (9.12) with generic ones of the same degree. This leads to the following simple characterization of the multilinear part of $\mathrm{Van}[\mathrm{RFE}_l^k]$ in terms of ρ and the image of composed boundary maps. The characterization naturally decomposes into separate ones for the homogeneous components of the various degrees d.

Corollary 9.12. *For any* $k, l \in \mathbb{N}$, *the set of multilinear polynomials* $p \in \mathbb{F}[x_1, \ldots, x_n]$ *in* $\mathrm{Van}[\mathrm{RFE}_l^k]$ *is given by the direct sum* $\oplus_{d=0}^{n-k-1} H_d$ *of homogeneous components of degree* $d \in \{0, \ldots, n-k-1\}$ *given by*

$$H_d \doteq \operatorname{span} \rho(\partial_{w^{[k+1]}}(\Lambda^{k+d+1}(U))), \tag{9.16}$$

where $w_1, \ldots, w_{k+1}$ range over all weight functions of degree at most $k + d - l - 1$.

For $d \leq l$, the only possible choices for the weight functions $w_1, \ldots w_{k+1}$ in Corollary 9.12 are linearly dependent, which implies that $\partial_{w^{[k+1]}}$ vanishes and therefore H_d only contains the zero polynomial. This is consistent with Proposition 5.1, as is the restriction $d \leq n - k - 1$.

Proof. By Theorem 1.3 and the fact that all the instantiations EVC_l^k are homogeneous of degree $l+1$, the multilinear elements in $\mathrm{Van}[\mathrm{RFE}_l^k]$ are exactly the linear combinations of terms of the form (9.12) where $S \subseteq [n]$ ranges over subsets of size $k + l + 2$ and $M \subseteq [n]$ over subsets disjoint with S. The homogeneous component of degree d equals the contributions of the combinations (S, M) where $|M| = m \doteq d - l - 1$. Since $S \sqcup M \subseteq [n]$ and $|S| + |M| = k + d + 1$, it follows that $d \leq n - k - 1$.

Since the weight functions on the right-hand side of (9.12) are of degree at most $|M| + k = k + d - l - 1$, the homogeneous component of degree d falls inside H_d. For the other inclusion, consider $\nu = v^T$ for $T \subseteq [n]$ with $|T| = t \doteq k + d + 1$. The first equality in (9.15) applies for any $N \subseteq [n]$ with $|N| = k + m + 1 = t - l - 1$. If we pick $N \subseteq T$, we have that $M \subseteq N \subseteq T$ and we can write T as $T = S \sqcup M$ where $|S| = |T| - |M| = k + l + 2$. Thus, each term on the left-most side of (9.15) is of the form of the boundary expression on the right-hand side of (9.12). By Lemma 9.10 and linearity, it follows that all of H_d can be realized as homogeneous components of degree d of polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$. □

For $d = l + 1$, up to constant factors, there is only one nontrivial composed boundary map $\partial_{w^{[k+1]}}$ up to scalar multiplication, namely the map $\partial_{\alpha^k \wedge \cdots \wedge \alpha^0}$ from Corollary 9.8. Thus, Corollary 9.8 represents the special case of Corollary 9.12 for degree $d = l + 1$.

The second equality in (9.15) from Proposition 9.11 leads to another characterization of the multilinear part of $\mathrm{Van}[\mathrm{RFE}_l^k]$ in terms of ρ and composed boundary maps, one that is more technical but to which we can directly apply the generalized image-kernel relationship. This leads to the following test for membership of multilinear polynomials in $\mathrm{Van}[\mathrm{RFE}_l^k]$. Consistent with the characterization in Corollary 9.12 and with Proposition 5.2, the test decomposes into independent ones for each of the homogeneous components.

Theorem 9.13. *Let $k, l \in \mathbb{N}$. For any multilinear polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$, $p(\mathrm{RFE}_l^k) = 0$ if and only if the homogeneous components $p^{(d)}$ of p for all degrees d satisfy the following requirements:*

1. *$p^{(d)} = 0$ if $d \leq l$ or $d \geq n - k$.*

2. *For all $d = l + \Delta + 1$ with $\Delta \in \{0, \ldots, n - k - l - 2\}$ and all weight functions $w_1, \ldots, w_{\Delta+1}$ of degree at most $k + \Delta$,*
$$\partial_{w_1 \wedge \cdots \wedge w_{\Delta+1}}(\rho^{-1}(p^{(d)})) = 0. \tag{9.17}$$

Proof. Consider the characterization (9.16) of the homogeneous components H_d in Corollary 9.12. We already argued that H_d only contains the zero polynomial for $d \leq l$ and that there are no terms for $d \geq n - k - 1$. This gives us condition 1.

In the remainder of the proof we consider the requirements for $d \in \{l+1, \ldots, n-k-1\}$. For multilinear p, $\rho^{-1}(p)$ is well-defined. Applying ρ^{-1} and the second equality in (9.15), we can alternately write (9.16) as

$$\rho^{-1}(H_d) = \operatorname*{span}_{\substack{B \subseteq N \\ |B|=k+1}} \partial_{L_N^B}(\Lambda^{k+d+1}(U)), \tag{9.18}$$

where $N \subseteq [n]$ can be any fixed subset of size $|N| = k+d-l$, namely by setting $m = d-l-1$, which we know is non-negative. For easier notation, we pick $N = [k+d-l]$. By Proposition 9.3 with $w_j \doteq L_{N,j}$ and $\Delta \doteq d-l-1$, we can further rewrite the right-hand side of (9.18) as

$$\rho^{-1}(H_d) = \bigcap_{\substack{B \subseteq [k+\Delta+1] \\ |B|=\Delta+1}} \ker(\partial_{L_N^B}) \cap \Lambda^d(U).$$

Thus $p^{(d)} \in H_d$ if and only if

$$(\forall B \subseteq [k+\Delta+1] \text{ with } |B| = \Delta+1) \;\; \partial_{L_N^B}(\rho^{-1}(p^{(d)})) = 0. \tag{9.19}$$

Another application of the second part of Proposition 9.11, this time with $k \leftarrow \Delta$, $m \leftarrow k$, $t \leftarrow d$, $\nu = \rho^{-1}(p^{(d)})$, and $N = [k+\Delta+1]$, shows that if (9.19) holds for the particular choice of weight functions $w_j = L_{N,j}$, then (9.19) holds for all choices of weight functions w_j of degree at most $k+\Delta$. The statement follows. □

As we will argue in more detail below, by another application of the first part of Proposition 9.11, it suffices in condition 2 of Theorem 9.13 to consider weight functions of the form $w_j(\alpha) = \mu_K(\alpha)\alpha^{\Delta-j+1}$ for $j \in [\Delta+1]$, where K ranges over all subsets of size k of some fixed $N \subseteq [n]$ with $|N| = k+\Delta+1$. In this case, (9.17) becomes

$$\partial_{\mu_K(\alpha)\alpha^\Delta \wedge \cdots \wedge \mu_K(\alpha)\alpha^0}(\rho^{-1}(p^{(d)})) = 0. \tag{9.20}$$

The left-hand side of (9.20) lives in $\Lambda^l(U)$, and the condition is equivalent to requiring that the coefficient of v^L vanishes for every subset $L \subseteq [n] \setminus K$ of size $|L| = l$. Those coefficients can be expressed in terms of evaluations of $\partial_L p^{(d)}\big|_{K \leftarrow 0}$, where we take the partial derivative with respect to the variables x_i for $i \in L$ and set the variables x_i for $i \in K$ to zero. Intuitively, whereas in Lemma 9.10 the effect of the masking factors μ_M was to retain only contributions of monomials that contain x_i for every $i \in M$, in this dual setting the effect of μ_K is to cancel the contributions of monomials that contain x_i for at least one $i \in K$.

Proposition 9.14. *Let $p \in \mathbb{F}[x_1, \ldots, x_n]$ be a multilinear polynomial, let $K, L \subseteq [n]$ be disjoint subsets with $|K| = k$ and $|L| = l$, and $\Delta \in \mathbb{N}$. Let $c_{K,L}$ denote the coefficient of v^L in $\partial_{\mu_K(\alpha)\alpha^\Delta \wedge \cdots \wedge \mu_K(\alpha)\alpha^0}(\rho^{-1}(p))$, and $e_{K,L}$ denote the value of $\partial_L p|_{K \leftarrow 0}$ upon the substitution $x_i \leftarrow \mu_K(a_i)/\mu_L(a_i)$ for $i \in [n] \setminus (K \sqcup L)$. Then $c_{K,L} = e_{K,L}/\det(A_L)$.*

Proof. By linearity, it suffices to establish the result for monomials $p = x^T$ where $T \subseteq [n]$. In such a case $\rho^{-1}(p) = v^T/\det(A_T)$.

If $L \not\subseteq T$, then $c_{K,L}$ vanishes because boundary maps can only remove components from a wedge product, not insert new components (see (9.5)). On the other hand, $\partial_L x^T\big|_{K\leftarrow 0}$ is identically zero because we are taking a partial derivative with respect to a variable that does not appear, so $e_{K,L}$ vanishes and the equality holds.

If $L \subseteq T$, then by applying Proposition 9.4 to v^T and scaling,

$$c_{K,L} = (-1)^{\mathrm{XInv}(M,L)} \prod_{i\in M} \mu_K(a_i) \cdot \det(A_M)/\det(A_T),$$

where $M \doteq T \setminus L$. Note that if $K \cap M \neq \emptyset$ then the term $\prod_{i\in M} \mu_K(a_i)$ vanishes, hence c_L vanishes. On the other hand, $\partial_L x^T\big|_{K\leftarrow 0}$ is identically zero because we are setting a variable to zero that appears in the monomial x^T. So, $e_{K,L}$ vanishes and the equality holds.

The remaining cases are those where $L \subseteq T$ and $K \cap M = \emptyset$. By (9.14)

$$\det(A_T) = \det(A_{M\sqcup L}) = (-1)^{\mathrm{XInv}(M,L)} \cdot \left(\prod_{i\in M}\prod_{j\in L}(a_i - a_j)\right) \cdot \det(A_M) \cdot \det(A_L).$$

Combining this with the notation $\mu_L(a_i) \doteq \prod_{j\in L}(a_i - a_j)$, we can rewrite the expression for $c_{K,L}$ as

$$c_{K,L} = \left(\prod_{i\in M} \frac{\mu_K(a_i)}{\mu_L(a_i)}\right) \cdot \frac{1}{\det(A_L)}. \tag{9.21}$$

On the other hand, we have that $\partial_L x^T\big|_{K\leftarrow 0} = x^M$, and the value upon the substitution $x_i \leftarrow \mu_K(a_i)/\mu_L(a_i)$ for $i \in [n] \setminus (K \sqcup L)$ equals

$$e_{K,L} = \prod_{i\in M} \frac{\mu_K(a_i)}{\mu_L(a_i)}. \tag{9.22}$$

The result follows by comparing (9.21) and (9.22). □

In combination with Theorem 9.13, the connection in Proposition 9.14 yields an alternate proof of Theorem 1.8. It provides a membership test for the ideal generated by the instantiations of EVC^k_l that, beyond the machinery of alternating algebra developed in this section, only requires the elementary properties of EVC^k_l stated in Proposition 3.4. In particular, it does not make use of the Zoom Lemma, which we developed as a tool to obviate the need for alternating algebra after we had obtained our results. Note that the alternate approach to Theorem 1.8 still relies on the Zoom Lemma for the connection to RFE, namely in the argument that the ideal generated by the instantiations of EVC^k_l includes all of $\mathrm{Van}[\mathrm{RFE}^k_l]$.

Alternate proof of Theorem 1.8. Consider the membership test given by Theorem 9.13. Condition 1 is equivalent to condition 1 in Theorem 1.8. It remains to argue that condition 2 is equivalent to condition 2 in Theorem 1.8.

Fix $\Delta \in \{0, \ldots, n-k-l-2\}$ and consider Proposition 9.11 with $k \leftarrow \Delta$, $m \leftarrow k$, $t \leftarrow d \doteq l+\Delta+1$, and $\nu = \rho^{-1}(p^{(d)})$. Set $N \subseteq [n]$ to be an arbitrary subset of size $N = k + \Delta + 1$ and rename the set M as K. The application of the first equality in Proposition 9.11 tells us that the combined requirements (9.17) over all choices of weight functions $w_1, \ldots, w_{\Delta+1}$ of degree at most $k + \Delta$ are equivalent to the combined requirements (9.20) over all subsets $K \subseteq N$ of size k, or, because of the arbitrariness of N, over all subsets $K \subseteq [n]$ of size k. The left-hand side of (9.20) is a linear combination of terms of the form v^L, where $L \subseteq [n]$ is a subset of size $|L| = d - \Delta - 1 = l$ and is disjoint from K because of the masking factor μ_K in all weight functions. Thus, (9.20) holds if and only if the coefficient $c_{K,L,d}$ of v^L on the left-hand side vanishes for every such L. By Proposition 9.14, $c_{K,L,d} = 0$ is equivalent to $e_{K,L,d} = 0$, where $e_{K,L,d}$ denotes the value of $\partial_L p^{(d)}\big|_{K\leftarrow 0}$ upon the substitution $x_i \leftarrow \mu_K(a_i)/\mu_L(a_i)$ for $i \in [n] \setminus (K \sqcup L)$.

In summary, condition 2 in Theorem 9.13 stipulates that for all disjoint subsets $K, L \subseteq [n]$ with $|K| = k$ and $|L| = l$,

$$(\forall d \in \{l+1, \ldots, n-k-1\})\ e_{K,L,d} = 0. \tag{9.23}$$

The value $e_{K,L,d}$ is also the coefficient of degree $d - l$ of the univariate polynomial in z obtained from $\partial_L p|_{K\leftarrow 0}$ after the substitution (1.7) from condition 2 in Theorem 1.8. Since the range of d in (9.23) covers all terms of this univariate polynomial in z, (9.23) is equivalent to the polynomial being zero, which is exactly condition 2 in Theorem 1.8. $\square$

Beyond multilinearity. Theorem 9.13 does well for understanding the multilinear elements of the vanishing ideal. For non-multilinear elements, one may do the following. Let $\widehat{\Lambda}^t(U)$ be $\Lambda^t(U)$ except that coefficients may be arbitrary polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ rather than just scalars in $\mathbb{F}$. The decoder map ρ and boundary maps ∂_w carry over to $\widehat{\Lambda}^t(U)$ directly, though now ρ is no longer injective. The following variation of Theorem 9.9 characterizes ideal membership for arbitrary polynomials.

Proposition 9.15. *Let $k, l \in \mathbb{N}$. For any polynomial $p \in \mathbb{F}[x_1, \ldots, x_n]$, $p(\mathrm{RFE}_l^k) = 0$ if and only if there exists $\eta \in \widehat{\Lambda}^{l+1}(U)$ with $\rho(\eta) = p$ such that, for every weight function w of degree at most k,*

$$\partial_w(\eta) = 0.$$

Proof. For the forward direction, we consider polynomials of the form $p = \mathrm{EVC}_l^k[S] \cdot m$, where $S \subseteq [n]$ with $|S| = k + l + 2$ and m is a (not necessarily multilinear) monomial in $\mathbb{F}[x_1, \ldots, x_n]$. One choice of $\eta \in \widehat{\Lambda}^{l+1}(U)$ for which $\rho(\eta) = p$ is $\eta = (-1)^{(k+1)(l+1)}\partial_{\alpha^k \wedge \cdots \wedge \alpha^0}(v^S)\cdot m$, by Lemma 9.7. For this choice of η and any weight function w of degree at most k, $\partial_w(\eta) = 0$. The forward direction follows since every polynomial p for which $p(\mathrm{RFE}_l^k) = 0$ can be expressed as a linear combination of polynomials of the described form.

For the backward direction, suppose there exists $\eta \in \widehat{\Lambda}^{l+1}(U)$ such that $\rho(\eta) = p$ and $\partial_w(\eta) = 0$ for all weight functions w of degree at most k. We can write $\eta = \sum_m \omega_m \cdot m$ as a linear combination of monomials $m \in \mathbb{F}[x_1, \ldots, x_n]$, each with coefficient $\omega_m \in \Lambda^{l+1}(U)$. Since ∂_w does not affect polynomial coefficients by nonconstant factors, we have that for each m, $\partial_w(\omega_m)$ vanishes for all w of degree at most k. Theorem 9.9 implies that $\rho(\partial_w(\omega_m))$ is not hit by RFE_l^k. By linearity, $\rho(\eta)$ is not hit by RFE_l^k. $\square$

While Proposition 9.15 applies to a broader class of polynomials, it has the drawback that representing polynomials with $\widehat{\Lambda}^{l+1}(U)$ is too redundant. Specifically, whenever p has a representation in $\widehat{\Lambda}^{l+1}(U)$, there are many $\eta \in \widehat{\Lambda}^{l+1}(U)$ that represent p, and most of them do *not* satisfy the boundary conditions, even when p belongs to the vanishing ideal. This erodes the utility of the characterization. Theorems 9.9 and 9.13 yield straightforward tests: Given p, form the unique η with $\rho(\eta) = p$, and then check whether the boundary conditions hold for η. Proposition 9.15, on the other hand, leaves η underspecified.

Acknowledgements

We are grateful to Hervé Fournier and Arpita Korwar for their presentation at WACT'18 in Paris [20]. We are indebted to Gautam Prakriya for helpful discussions and detailed feedback. We also thank Amir Shpilka and Michael Forbes for comments and encouragement, the anonymous reviewers for their careful proofreading and interesting suggestions, and the ToC editors for their thorough work. Finally, we appreciate the partial support for this research by the U.S. National Science Foundation under Grants No. 1838434, 2137424, and 2312540. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] WILLIAM ADAMS AND PHILIPPE LOUSTAUNAU: *An Introduction to Gröbner Bases*. Amer. Math. Soc., 1994. [doi:10.1090/gsm/003] 6, 21

[2] MANINDRA AGRAWAL: Proving lower bounds via pseudo-random generators. In *Proc. 25th Found. Softw. Techn. Theoret. Comp. Sci. Conf. (FSTTCS'05)*, pp. 92–105. Springer, 2005. [doi:10.1007/11590156_6] 2

[3] MANINDRA AGRAWAL, ROHIT GURJAR, ARPITA KORWAR, AND NITIN SAXENA: Hitting-sets for ROABP and sum of set-multilinear circuits. *SIAM J. Comput.*, 44(3):669–697, 2015. [doi:10.1137/140975103] 3, 10

[4] MANINDRA AGRAWAL, CHANDAN SAHA, AND NITIN SAXENA: Quasi-polynomial hitting-set for set-depth-Δ formulas. In *Proc. 45th STOC*, pp. 321–330. ACM Press, 2013. [doi:10.1145/2488608.2488649] 3

[5] RAVINDRA AHUJA, THOMAS MAGNANTI, AND JAMES ORLIN: *Network Flows: Theory, Algorithms, and Applications*. Pearson, 1993. [doi:10.5555/137406] 50

[6] MATTHEW ANDERSON, MICHAEL A. FORBES, RAMPRASAD SAPTHARISHI, AMIR SHPILKA, AND BEN LEE VOLK: Identity testing and lower bounds for read-k oblivious algebraic branching programs. *ACM Trans. Comput. Theory*, 10(1/3):1–30, 2018. [doi:10.1145/3170709] 10

[7] MATTHEW ANDERSON, DIETER VAN MELKEBEEK, AND ILYA VOLKOVICH: Deterministic polynomial identity tests for multilinear bounded-read formulae. *Comput. Complexity*, 24(4):695–776, 2015. [doi:10.1007/s00037-015-0097-4] 3, 7, 8, 10

[8] ROBERT ANDREWS AND MICHAEL A. FORBES: Ideals, determinants, and straightening: proving and using lower bounds for polynomial ideals. In *Proc. 54th STOC*, pp. 389–402. ACM Press, 2022. [doi:10.1145/3519935.3520025] 14

[9] VIKRAMAN ARVIND, PUSHKAR S. JOGLEKAR, PARTHA MUKHOPADHYAY, AND S. RAJA: Randomized polynomial-time identity testing for noncommutative circuits. *Theory of Computing*, 15(7):1–36, 2019. [doi:10.4086/toc.2019.v015a007] 16

[10] VISHWAS BHARGAVA AND SUMANTA GHOSH: Improved hitting set for orbit of ROABPs. *Comput. Complexity*, 31(15), 2022. [doi:10.1007/s00037-022-00230-9] 3, 10

[11] PRERONA CHATTERJEE AND ANAMAY TENGSE: On annihilators of explicit polynomial maps, 2023. [arXiv:2309.07612] 14

[12] DAVID A. COX, JOHN LITTLE, AND DONAL O'SHEA: *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer, 2013. [doi:10.1007/978-3-319-16721-3] 6, 21

[13] RICHARD A. DEMILLO AND RICHARD J. LIPTON: A probabilistic remark on algebraic program testing. *Inform. Process. Lett.*, 7(4):193–195, 1978. [doi:10.1016/0020-0190(78)90067-4] 16

[14] MICHAEL A. FORBES: Deterministic divisibility testing via shifted partial derivatives. In *Proc. 56th FOCS*, pp. 451–465. IEEE Comp. Soc., 2015. [doi:10.1109/FOCS.2015.35] 3, 7, 11

[15] MICHAEL A. FORBES, RAMPRASAD SAPTHARISHI, AND AMIR SHPILKA: Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Proc. 46th STOC*, pp. 867–875. ACM Press, 2014. Full version arXiv:1309.5668. [doi:10.1145/2591796.2591816] 3, 10, 14

[16] MICHAEL A. FORBES AND AMIR SHPILKA: On identity testing of tensors, low-rank recovery and compressed sensing. In *Proc. 44th STOC*, p. 163–172. ACM Press, 2012. [doi:10.1145/2213977.2213995] 14

[17] MICHAEL A. FORBES AND AMIR SHPILKA: Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. In *Proc. 54th FOCS*, pp. 243–252. IEEE Comp. Soc., 2013. [doi:10.1109/FOCS.2013.34] 10

[18] MICHAEL A. FORBES, AMIR SHPILKA, IDDO TZAMERET, AND AVI WIGDERSON: Proof complexity lower bounds from algebraic circuit complexity. *Theory of Computing*, 17(10):1–88, 2021. [doi:10.4086/toc.2021.v017a010] 14

[19] MICHAEL A. FORBES, AMIR SHPILKA, AND BEN LEE VOLK: Succinct hitting sets and barriers to proving lower bounds for algebraic circuits. *Theory of Computing*, 14(18):1–45, 2018. Preliminary version in STOC'17. [doi:10.4086/toc.2018.v014a018] 3, 7, 11, 14

[20] HERVÉ FOURNIER AND ARPITA KORWAR: Limitations of the Shpilka–Volkovich generator. In *Workshop on Algebraic Complexity Theory (WACT), Paris*, 2018. CONF and SLIDES. 7, 65

[21] ARIEL GABIZON AND RAN RAZ: Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008. [doi:10.1007/s00493-008-2259-3] 14

[22] HERMANN GRASSMANN: *Die lineale Ausdehnungslehre, ein neuer Zweig der Mathematik: dargestellt und durch Anwendungen auf die übrigen Zweige der Mathematik, wie auch auf die Statik, Mechanik, die Lehre vom Magnetismus und die Krystallonomie erläutert*. Otto Wigand, Leipzig, 1844. Available at HathiTrust. 50, 67

[23] HERMANN GRASSMANN: *Extension theory*. Amer. Math. Soc., 2000. Translation of [22], translated by LLOYD C. KANNENBERG. 50

[24] ZEYU GUO AND ROHIT GURJAR: Improved explicit hitting-sets for ROABPs. In *Proc. 24th Internat. Conf. on Randomization and Computation (RANDOM'20)*, pp. 4:1–16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020. [doi:10.4230/LIPIcs.APPROX/RANDOM.2020.4] 10

[25] ROHIT GURJAR, ARPITA KORWAR, AND NITIN SAXENA: Identity testing for constant-width, and any-order, read-once oblivious arithmetic branching programs. *Theory of Computing*, 13(2):1–21, 2017. [doi:10.4086/toc.2017.v013a002] 3, 10

[26] ROHIT GURJAR, ARPITA KORWAR, NITIN SAXENA, AND THOMAS THIERAUF: Deterministic identity testing for sum of read-once oblivious arithmetic branching programs. *Comput. Complexity*, 26(4):835–880, 2017. [doi:10.1007/s00037-016-0141-z] 3, 7, 10

[27] JOOS HEINTZ AND CLAUS-PETER SCHNORR: Testing polynomials which are easy to compute. In *Proc. 12th STOC*, pp. 262–272. ACM Press, 1980. [doi:10.1145/800141.804674] 2

[28] RUSSELL IMPAGLIAZZO AND AVI WIGDERSON: P=BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Proc. 29th STOC*, pp. 220–229. ACM Press, 1997. [doi:10.1145/258533.258590] 2

[29] MAURICE JANSEN, YOUMING QIAO, AND JAYALAL SARMA M. N.: Deterministic identity testing of read-once algebraic branching programs, 2009. [arXiv:0912.2565] 10

[30] MAURICE JANSEN, YOUMING QIAO, AND JAYALAL SARMA M. N.: Deterministic black-box identity testing π-ordered algebraic branching programs. In *Proc. 30th Found. Softw. Techn. Theoret. Comp. Sci. Conf. (FSTTCS'10)*, pp. 296–307. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2010. [doi:10.4230/LIPIcs.FSTTCS.2010.296] 10

[31] VALENTINE KABANETS AND RUSSELL IMPAGLIAZZO: Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complexity*, 13(1–2):1–46, 2004. [doi:10.1007/s00037-004-0182-6] 2

[32] ZOHAR S. KARNIN, PARTHA MUKHOPADHYAY, AMIR SHPILKA, AND ILYA VOLKOVICH: Deterministic identity testing of depth-4 multilinear circuits with bounded top fan-in. *SIAM J. Comput.*, 42(6):2114–2131, 2013. [doi:10.1137/110824516] 3, 8, 14

[33] ZOHAR S. KARNIN AND AMIR SHPILKA: Black box polynomial identity testing of generalized depth-3 arithmetic circuits with bounded top fan-in. *Combinatorica*, 31(3):333–364, 2011. [doi:10.1007/s00493-011-2537-3] 14

[34] ADAM R. KLIVANS AND DANIEL SPIELMAN: Randomness efficient identity testing of multivariate polynomials. In *Proc. 33rd STOC*, p. 216–223. ACM Press, 2001. [doi:10.1145/380752.380801] 14

[35] ARPITA KORWAR: Personal communication, 2021. 7

[36] KLAUS KÜHNLE AND ERNST W. MAYR: Exponential space computation of Gröbner bases. In *Proc. 21st Internat. Symp. Symbolic and Algebraic Computation (ISSAC'96)*, pp. 63–71. ACM Press, 1996. [doi:10.1145/236869.236900] 6

[37] MRINAL KUMAR AND SHUBHANGI SARAF: Arithmetic circuits with locally low algebraic rank. *Theory of Computing*, 13(6):1–33, 2017. [doi:10.4086/toc.2017.v013a006] 3, 11

[38] ERNST W. MAYR: Some complexity results for polynomial ideals. *J. Complexity*, 13(3):303–325, 1997. [doi:10.1006/jcom.1997.0447] 6

[39] DORI MEDINI AND AMIR SHPILKA: Hitting sets and reconstruction for dense orbits in VP_e and $\Sigma\Pi\Sigma$ circuits. In *Proc. 36th Comput. Complexity Conf. (CCC'21)*, pp. 19:1–27. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2021. [doi:10.4230/LIPIcs.CCC.2021.19] 3, 14, 34

[40] DIETER VAN MELKEBEEK AND ANDREW MORGAN: Polynomial identity testing via evaluation of rational functions. In *Proc. 13th Innovations in Theoret. Comp. Sci. Conf. (ITCS'22)*, pp. 119:1–24. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2022. [doi:10.4230/LIPIcs.ITCS.2022.119] 1

[41] DANIEL MINAHAN AND ILYA VOLKOVICH: Complete derandomization of identity testing and reconstruction of read-once formulas. *ACM Trans. Comput. Theory*, 10(3/10):1–11, 2018. [doi:10.1145/3196836] 3, 9, 10, 33

[42] NOAM NISAN: Lower bounds for non-commutative computation. In *Proc. 23rd STOC*, pp. 410–418. ACM Press, 1991. [doi:10.1145/103418.103462] 11, 42

[43] NOAM NISAN AND AVI WIGDERSON: Hardness vs randomness. *J. Comput. System Sci.*, 49(2):149–167, 1994. Preliminary version in FOCS'88. [doi:10.1016/S0022-0000(05)80043-1] 2

[44] ØYSTEIN ORE: Über höhere Kongruenzen. *Norsk Mat. Forenings Skrifter, Ser. I*, 1(7):1–15, 1922. 16

[45] RAN RAZ AND AMIR SHPILKA: Deterministic polynomial identity testing in non-commutative models. *Comput. Complexity*, 14(1):1–19, 2005. [doi:10.1007/s00037-005-0188-8] 10

[46] CHANDAN SAHA AND BHARGAV THANKEY: Hitting sets for orbits of circuit classes and polynomial families. In *Proc. 25th Internat. Conf. on Randomization and Computation (RANDOM'21)*, pp. 50:1–26. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2021. [doi:10.4230/LIPIcs.APPROX/RANDOM.2021.50, ECCC:TR21–015] 3, 10

[47] JACOB T. SCHWARTZ: Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980. [doi:10.1145/322217.322225] 16

[48] AMIR SHPILKA AND ILYA VOLKOVICH: Read-once polynomial identity testing. *Comput. Complexity*, 24(3):477–532, 2015. [doi:10.1007/s00037-015-0105-8] 2, 3, 7, 8, 9

[49] AMIR SHPILKA AND AMIR YEHUDAYOFF: Arithmetic circuits: A survey of recent results and open questions. *Found. Trends Theor. Comp. Sci.*, 5(3–4), 2010. [doi:10.1561/0400000039] 2

[50] RICHARD ZIPPEL: Probabilistic algorithms for sparse polynomials. In *Proc. Internat. Symp. Symbolic and Algebraic Computation (EUROSAM'79)*, pp. 216–226. ACM Press, 1979. [doi:10.1007/3-540-09519-5_73] 16

AUTHORS

Ivan Hu
Ph. D. student
Department of Computer Science
University of Wisconsin – Madison
Madison, Wisconsin, USA
ilhu@wisc.edu
hhttps://pages.cs.wisc.edu/~ihu/

Dieter van Melkebeek
Professor
Department of Computer Science
University of Wisconsin – Madison
Madison, Wisconsin, USA
dieter@cs.wisc.edu
https://pages.cs.wisc.edu/~dieter/

Andrew Morgan
Software engineer
Google
amorgan@cs.wisc.edu
https://pages.cs.wisc.edu/~amorgan/

ABOUT THE AUTHORS

IVAN HU is a second year Ph. D. student at the University of Wisconsin–Madison under the supervision of Dieter van Melkebeek. He is studying complexity theory, with interests in algebraic complexity and pseudorandomness. He is currently an NSF Graduate Research Fellow.

DIETER VAN MELKEBEEK received his Ph. D. from the University of Chicago, under the supervision of Lance Fortnow. His thesis was awarded the ACM Doctoral Dissertation Award. After postdocs at DIMACS and the Institute for Advanced Study, he joined the faculty at the University of Wisconsin-Madison, where he currently is a full professor. His research interests include the power of randomness, lower bounds for NP-complete problems, and connections between derandomization and lower bounds.

ANDREW MORGAN received his Ph. D. in 2022 from the University of Wisconsin–Madison under the supervision of Dieter van Melkebeek. After graduating, he became a software engineer at Google.