The complexity of the annihilating polynomial.

Neeraj Kayal [*]
kayaln@dimacs.rutgers.edu

December 10, 2007

Abstract

Let $\mathbb{F}$ be a field and $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be a set of k polynomials of degree d in n variables over the field $\mathbb{F}$. These polynomials are said to be *algebraically dependent* if there exists a nonzero k-variate polynomial $A(t_1, \ldots, t_k) \in \mathbb{F}[t_1, \ldots, t_k]$ such that $A(f_1, \ldots, f_k) = 0$. A is then called an $(f_1, \ldots, f_k)$-annihilating polynomial.

Given $(f_1, \ldots, f_k)$ can we determine the existence of an $(f_1, \ldots, f_k)$-annihilating polynomial? Even for the very concise representation of polynomials as arithmetic circuits, there exists an efficient randomized algorithm for doing this. What is the degree of $A(t_1, \ldots, t_k)$? We give closely matching upper and lower bounds for the degree of the annihilating polynomial. The degree bounds also provide a PSPACE algorithm for computing $A(t_1, \ldots, t_k)$. Can $A(t_1, \ldots, t_k)$ be computed efficiently? We show that it is NP-hard to decide if $A(0, \ldots, 0)$ equals zero and #P-hard to evaluate $A(0, \ldots, 0) \pmod p$ for a given prime p. Even if $A(t_1, \ldots, t_k)$ cannot be computed efficiently, does there at least exist a small circuit representation of $A(t_1, \ldots, t_k)$? We show that $A(t_1, \ldots, t_k)$ does not admit a small circuit representation unless the polynomial hierarchy collapses.

In summary, this means that testing for algebraic dependence is one of those extremely rare problems where determining the existence of an object (the annihilating polynomial in our case) can be done efficiently but the actual computation of the object is *provably hard.* These results also answer some questions posed by Dvir, Gabizon and Wigderson [DGW07].

1 Introduction

Motivation

The notion of algebraic dependence between a set of polynomials is defined as follows.

Definition 1. Algebraic Dependence. Let $\mathbf{f} = (f_1, \ldots, f_k)$ be a vector of k polynomials (of degree $\leq d$) where each $f_i \in \mathbb{F}[x_1, \ldots, x_n]$ is an n-variate polynomial over the field $\mathbb{F}$. A nonzero polynomial $A(t_1, \ldots, t_k) \in \mathbb{F}[t_1, \ldots, t_k]$ is said to be an $\mathbf{f}$-*annihilating polynomial* if $A(f_1, \ldots, f_k) = 0$. The polynomials $f_1, \ldots, f_k$ are said to be *algebraically dependent* if there exists an $\mathbf{f}$-annihilating polynomial.

Example: The polynomials $f_1(x, y) := (x^2 + y)^2$ and $f_2(x, y) := (x^2 + y)^3 + 1$ are algebraically dependent for they satisfy the equation $f_1^3 = (f_2 - 1)^2$. Thus $A(t_1, t_2) = t_1^3 - (t_2 - 1)^2$ is a (f_1, f_2)-annihilating polynomial. On the other hand the monomials x and y in $\mathbb{F}[x, y]$ are algebraically independent.

[*] DIMACS Center. This research was supported by a fellowship from D. E. Shaw & Co., L.P while the author was visiting the Institute for Advanced Study, Princeton.

The concept of algebraic dependence is a basic concept in algebra and algebraic geometry (cf. the texts by Schinzel [Sch82] and by Hartshorne [Har77]). Within computational complexity, this notion of algebraic independence has been used for example in proving that $n/2$ multiplications are necessary for evaluating a specific fixed polynomial at a given input point [**?**]. More recently it was used by Dvir, Gabizon and Wigderson [DGW07] to construct deterministic extractors for sources which are polynomial maps over finite fields. In this paper we study this relationship between polynomials from a computational perspective. This kind of study has been done before by L'vov [L'v84]. L'vov was motivated to work on this problem for it amounts to computing the 'invariant relationships' that exist between the values computed at an intermediate stage of execution of an arithmetic straight-line program. An invariant relationship is an algebraic equation satisfied by the intermediate values *which holds true for any choice of input values to the program*, and thus its the same thing as an annihilating polynomial.

Discussion and statement of results

Let us quickly review some useful combinatorial properties of this relationship. Observe that when the polynomials all have degree $d = 1$ then algebraic independence coincides with the notion of linear independence. From the definition of algebraic dependence it quickly follows that, as with linear dependence, this relationship gives rise to a *matroid* over the set of polynomials. That is, it satisfies the following combinatorial properties.

- If $f_1, \ldots, f_k$ are algebraically independent, then any subset of polynomials in $\{f_1, \ldots, f_k\}$ are also algebraically independent.
- If a polynomial g is algebraically dependent on $f_1, \ldots, f_k$ but not on $f_1, \ldots, f_{k-1}$ then f_k is algebraically dependent upon $f_1, \ldots, f_{k-1}, g$.

These properties imply that all maximal algebraically independent subsets of $f_1, \ldots, f_k$ have the same cardinality. This number is defined to be the *algebraic rank* of the set of polynomials $f_1, \ldots, f_k$. As we shall see, algebraic dependence also gives rise to a number of very interesting algebraic properties. Now suppose that $\mathbf{f} = (f_1, \ldots, f_k)$ are algebraically dependent polynomials, no proper subset of which are algebraically dependent. We will see (Lemma 7) that the minimal degree $\mathbf{f}$-annihilating polynomial $A(t_1, \ldots, t_k)$ is unique upto constant factors. In this case a *monic* minimal degree $\mathbf{f}$-annihilating polynomial $A(\mathbf{t})$ is uniquely defined and this polynomial $A(\mathbf{t})$ we shall refer to as *the* $\mathbf{f}$-annihilating polynomial. The above discussion motivates the following two questions:

- What is the complexity of testing the algebraic dependence of a given set of polynomials?
- What is the complexity of computing *the smallest annihilating polynomial* of a given set of algebraically dependent polynomials?

For the sake of concreteness, we shall fix the underlying field $\mathbb{F}$ to be the field $\mathbb{Q}$ of rational numbers. The questions posed above make sense for any field and many of the results contained here also hold for any field $\mathbb{F}$, especially if the field has a large enough characteristic. At first sight, the two questions posed above appear to have the same complexity but it turns out that the decision problem concerning the existence of an annihilating polynomial turns out to be far easier.

Let $J_{\mathbf{f}}(\mathbf{x})$ be the partial derivative matrix,

$$J_{\mathbf{f}}(\mathbf{x}) \stackrel{\text{def}}{=} \left(\left(\frac{\partial f_i}{\partial x_j} \right) \right)_{k \times n}.$$

This matrix is known as the jacobian of the set of polynomials in $\mathbf{f}$. The following is a classical theorem (cf. Ehrenborg and Rota [ER93] for a proof).

Theorem 2. The Jacobian Criterion for algebraic independence: *Let $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be a set of k polynomials in n variables over the field $\mathbb{F}$. Then these polynomials are algebraically independent if and only if the Jacobian matrix, $J_{\mathbf{f}}(\mathbf{x})$, matrix has rank k.*

From this result follows easily an efficient randomized algorithm due to [DGW07] for testing algebraic dependence.

Corollary 3. *There exists a randomized polynomial time algorithm that on input a set of k arithmetic circuits over a field $\mathbb{F}$, determines if the polynomials computed by these arithmetic circuits are algebraically dependent or not.*

Proof. Let the polynomial computed by the i-th circuit be $f_i(x_1, \ldots, x_n)$. By the result of Baur-Strassen-Morgenstern [BS83, Mor85] we can efficiently construct another circuit that computes $\frac{\partial f_i}{\partial x_j}$ for all $j \in [n]$. With these circuits for partial derivatives in hand, we can determine the rank of the partial derivative matrix by plugging in random values. As in the randomized Schwarz-Zippel identity testing algorithm [Sch80, Zip90] the rank of the jacobian matrix with these random values of the variables plugged in will equal the algebraic rank of the jacobian matrix with high probability. □

The above theorem gave rise to the hope that we may also be able to compute the $\mathbf{f}$-annihilating polynomial (perhaps by examining the null space and the range space of $J_{\mathbf{f}}(\mathbf{x})$ more closely). As we shall see, however, the task of computing the annihilating polynomial turns out to be much more difficult. In studying the complexity of various computational problems, it is rather unusual to come across a computational problem for which we can determine the existence of a solution to a computational problem without being able to compute the solution explicitly. So let us take a moment to understand what is going on here. A key observation that is used in the proof of theorem 2 is the following:

Lemma 4. *Over any field and for any set of polynomials in n variables, their algebraic rank is at most n. In particular, a set of $(n+1)$ polynomials in n variables is algebraically dependent.*

The proof of this lemma is via a dimension counting argument and it is the chief non-constructive ingredient in the proof of Theorem 2. Our main technical contribution is is to analyse this situation and to give, in Lemma 9, a much more insightful description of the minimal annihilating polynomial for a set of $n+1$ polynomials in n variables. This description is then used to prove the degree and complexity lower bounds.

Let us first ask the question - 'If the input polynomials have degree at most d, then what is the maximum possible degree of the annihilating polynomial'? We show that if the f_i's are dependent then there exists an annihilating polynomial of degree at most $(r+1) \cdot d^r$, where r is the algebraic rank of the input polynomials. We also give a very explicit family of polynomials whose minimal annihilating polynomial has degree at least d^r. These bounds give a satisfactory answer to the last question. They also imply a PSPACE-algorithm for computing the annihilating polynomial.

Let us now look a the problem of computing $A(\mathbf{t})$. Let us assume that we are given as input a set of polynomials $f_1, \ldots, f_k$ in the usual dense representation of polynomials, wherein a polynomial of degree d in n variables is specified by specifying the coefficients of all the $\binom{d+n}{n}$ possible monomials. The degree lower bound means that it may take exponential time to write down the polynomial $A(t_1, \ldots, t_k)$ in the usual dense representation of polynomials. But perhaps it is feasible to compute the coefficients of some specific monomials of $A(\mathbf{t})$ or to evaluate $A(\mathbf{t})$ at some specific points of interest? So let us look at the complexity of computing just the constant term, or equivalently, of evaluating $A(\mathbf{t})$ at the point $(0, \ldots, 0)$. We show that it is NP-hard to decide if $A(0, \ldots, 0)$ equals zero and #P-hard to evaluate $A(0, \ldots, 0)$ modulo a given prime p. Even if $A(t_1, \ldots, t_k)$ cannot be computed efficiently, does there at least exist a small circuit representation of $A(t_1, \ldots, t_k)$? We show that $A(t_1, \ldots, t_k)$ does not admit a small circuit representation unless the polynomial hierarchy collapses

Comparison with previous work. As mentioned earlier these questions were investigated earlier by [L'v84]. For the degree an upper bound of $(n+1)d^n$ was previously established. Note that the algebraic rank r is always less than or equal to the number of variables n (Lemma 4) and thus our degree upper bound of $(r+1)d^r$ is an improvement on the previous bound. In cases where r is significantly smaller than n, this is much better. For example, when $r = 1$, our bound is linear whereas the previous bound is exponential. To the best of our knowledge no lower bounds were known previously.

Remark. For many of the usual fields $\mathbb{F}$ such as finte fields, real numbers and rational function fields over $\mathbb{C}$, the converse problem of deciding whether a given polynomial $A(t_1, \ldots, t_k) \in \mathbb{F}[t_1, \ldots, t_k]$ admits a solution in the ring $\mathbb{F}[x]$ of polynomials over $\mathbb{F}$ is known to be undecidable [KR92, Vid94, Den78].

The rest of this paper is organized as follows: after reviewing some preliminaries in section 2, we prove the unqueness and other fundamental properties of the minimal annihilating polynomial in section 3. These are then used to give degree bounds for the annihilating polynomial in section 4 and computational complexity bounds in section 5. For lack of space, we push some of the proofs to the appendix.

2 Preliminaries

We will use $[n]$ to denote the set of integers $\{1, 2, \ldots, n\}$. $\mathbb{F}$ will denote a field and $\overline{\mathbb{F}}$ its algebraic closure. $\mathbb{F}[x_1, \ldots, x_n]$ shall denote the ring of polynomials in variables $\{x_i | i \in [n]\}$ over $\mathbb{F}$ and $\mathbb{F}(x_1, \ldots, x_n)$ the field of rational functions in these variables. For $f \in \mathbb{F}[x_1, \ldots, x_n]$, $\deg(f)$ will denote the total degree of f and $\deg_{x_i}(f)$ will denote the degree of f with respect to the variable x_i. We shall say that $f(\mathbf{x})$ is monic if the coefficient of the largest monomial in antilexicographic order occuring in $f(\mathbf{x})$ is 1. For $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$, $\mathbb{F}[f_1, \ldots, f_k] \subseteq \mathbb{F}[x_1, \ldots, x_n]$ shall denote the subalgebra of $\mathbb{F}[x_1, \ldots, x_n]$ generated by $f_1, \ldots, f_k$. That is,

$$\mathbb{F}[f_1, \ldots, f_k] \overset{\mathrm{def}}{=} \{B(f_1, \ldots, f_k) \ \mid B(t_1, \ldots, t_k) \in \mathbb{F}[t_1, \ldots, t_k]\}.$$

We shall use bold letters such as $\mathbf{x}$ to denote a vector with the number of elements in the vector being understood from the context in which it appears.

2.1 Algebraic Preliminaries

Absolute Irreducibility.

Definition 5. A polynomial $f(x_1, \ldots, x_n) \in \mathbb{F}[x_1, \ldots, x_n]$ is said to be *absolutely irreducible* if it is irreducible over the algebraic closure of $\mathbb{F}$.

Example: For example, $(y^2 - x^3) \in \mathbb{Q}[x, y]$ is absolutely irreducible whereas $(y^2 + x^2) \in \mathbb{Q}[x, y]$ is irreducible over $\mathbb{Q}$ but factors into $(y + \sqrt{-1}x)(y - \sqrt{-1}x)$ over the extension $\mathbb{Q}(\sqrt{-1})$ and hence is not absolutely irreducible.

We shall require the following special case of the Hilbert Nullstellensatz.

Lemma 6. *Let $\overline{\mathbb{F}}$ be an algebraically closed field and $f, g \in \overline{\mathbb{F}}[x_1, \ldots, x_n]$ be (multivariate) polynomials over $\overline{\mathbb{F}}$. Let $g(\mathbf{x})$ be $\overline{\mathbb{F}}$-irreducible. If every zero $\mathbf{a} \in \overline{\mathbb{F}}^n$ of $f(\mathbf{x})$ is also a zero of $g(\mathbf{x})$ then the polynomial $f(\mathbf{x})$ is a power of $g(\mathbf{x})$. That is, there exists an integer $t \in \mathbb{Z}_{\geq 0}$ such that $f(\mathbf{x}) = g(\mathbf{x})^t$.*

Rings, characteristic polynomials and resultants.

Let $R \supseteq \mathbb{F}$ be a ring whose elements constitute a finite dimensional vector space over the field $\mathbb{F}$. Then for any $\alpha \in R$, the map $\phi_\alpha : R \mapsto R$, $\phi_\alpha(r) = r \cdot \alpha$ is a linear transformation on this vector space. We will refer to the characteristic polynomial (in the indeterminate z) of this linear tranformation ϕ_α as the characteristic polynomial of α and denote it by $\text{charpoly}_{\alpha \in R/\mathbb{F}}(z) \in \mathbb{F}[z]$. Recall that $\text{charpoly}_{\alpha \in R/\mathbb{F}}(z)$ has a zero constant term (i.e. $\text{charpoly}_{\alpha \in R/\mathbb{F}}(0) = 0$) if and only if the element $\alpha \in R$ is a zero divisor in the ring R.

Now suppose that R is of the form $R := \mathbb{F}[x]/\langle g(x) \rangle$, where $g(x) \in \mathbb{F}[x]$ is a univariate polynomial. Viewing an arbitary polynomial $f(x) \in \mathbb{F}[x]$ as an element of R via the natural homomorphism $f(x) \mapsto f(x) \bmod g(x)$, we will denote by $\text{charpoly}_{f(x) \bmod g(x)}(z) \in \mathbb{F}[z]$ the characteristic polynomial of this element of R. Then $\text{charpoly}_{f(x) \text{mod } g(x)}(z)$ has a zero constant term if and only if $f(x)$ and $g(x)$ have a common zero in the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$. Finally, the resultant with respect to the variable x_i of two polynomials $f(\mathbf{x})$ and $g(\mathbf{x})$ shall be denoted by $\text{Res}_x(f, g)$.

3 Properties of the annihilating polynomial.

Let $\mathbb{F}$ be a field and $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be a set of k polynomials in n variables over the field $\mathbb{F}$. Suppose that these polynomials are algebraically dependent. Then the set of $\mathbf{f}$-annihilating polynomials forms an ideal $\mathfrak{U}$ of the polynomial ring $\mathbb{F}[t_1, \ldots, t_k]$ (easy verification). In this section we investigate the properties of this ideal $\mathfrak{U}$ and the minimal degree polynomials in $\mathfrak{U}$. We shall establish that when $\mathbf{f} := (f_1, \ldots, f_k)$ forms a minimal set of algebraically dependent polynomials then there exists a unique (upto constant factors) minimal $\mathbf{f}$-annihilating polynomial $A(t_1, \ldots, t_k)$. Moreover this polynomial $A(\mathbf{t})$ is absolutely irreducible and it generates the ideal $\mathfrak{U}$.

Lemma 7. *Let $f_1, f_2, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be a set of* algebraically dependent *polynomials in n variables over the field $\mathbb{F}$, no proper subset of which is algebraically dependent. Then the ideal $\mathfrak{U}$ of $\mathbf{f}$-annhilating polynomials is generated by a single absolutely irreducible polynomial $A(\mathbf{t})$. Moreover, $A(\mathbf{t})$ remains the minimal annihilating polynomial of $\{f_1, f_2, \ldots, f_k\}$ when they are viewed as polynomials over the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$.*

Proof. See appendix. □

We now investigate the annihilating polynomial more closely and show its relationship to a certain characteristic polynomial. We begin with an easy observation.

Lemma 8. *Let $f_1, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be polynomials such that the system of equations*

$$f_1 = \ldots = f_k = 0$$

has a common solution $P \in \overline{\mathbb{F}}^n$. If $f_1, \ldots, f_k$ happen to be algebraically dependent with an annihilating polynomial $A(t_1, \ldots, t_k)$ then $A(0, \ldots, 0) = 0$.

Proof. Since $A(\mathbf{t})$ is $\mathbf{f}$-annihilating, we have

$$\begin{array}{rcl} A(f_1(\mathbf{x}), \ldots, f_k(\mathbf{x})) & = & 0 \\ \Rightarrow \ A(f_1(P), \ldots, f_k(P)) & = & 0 \\ \Rightarrow \ A(0, \ldots, 0) & = & 0. \end{array}$$

□

At this point, it is natural to wonder if the converse of the above lemma is true. If the converse were true, it would have immediately implied the NP-hardness of determining if $A(0, \ldots, 0) = 0$. As it turns out, the converse is not true as the following example shows.

Example: Consider the following polynomials over $\mathbb{C}$:

$$f_1 \stackrel{\text{def}}{=} x_1, \quad f_2 \stackrel{\text{def}}{=} x_1x_2 - 1, \quad f_3 \stackrel{\text{def}}{=} x_1^2x_2$$

The polynomial

$$A(t_1, t_2, t_3) \stackrel{\text{def}}{=} t_3 - t_1 \cdot (t_2 + 1)$$

is an (f_1, f_2, f_3)-annihilating polynomial (easy verification), is absolutely irreducible (since it is linear with respect to t_3) and is therefore the unique minimal $\mathbf{f}$-annihilating polynomial. Note that $A(0, 0, 0) = 0$. However, the system of equations

$$f_1 = f_2 = f_3 = 0$$

has no common solution in $\mathbb{C}$ because $x_2f_1 - f_2 = 1$ (Hilbert Nullstellensatz).

Nevertheless, we shall show that a partial converse to Lemma 8 does hold true. This partial converse will fall out of a nice characterization of the annihilating polynomial as a certain characteristic polynomial.

Lemma 9. *For each $i \in [n]$, let $f_i(x_1, \ldots, x_i) \in \mathbb{F}[x_1, \ldots, x_i]$ be a polynomial in the variables $x_1, \ldots, x_i$. Suppose further that f_i when viewed as a univariate polynomial in x_i with coefficients from $\mathbb{F}[x_1, \ldots, x_{i-1}]$ is monic. Let $g(x_1, \ldots, x_n) \in \mathbb{F}[x_1, \ldots, x_n]$ be some n-variate polynomial over $\mathbb{F}$. Define $r(v_1, \ldots, v_n, u) \in \mathbb{F}(v_1, \ldots, v_n)[u]$ to be the polynomial*

$$r(v_1, \ldots, v_n, u) \stackrel{\text{def}}{=} \mathit{charpoly}_{g(x) \ \mathit{mod} \ (f_1 - v_1, \ldots, f_n - v_n)}(u).$$

Then the set of polynomials $\mathbf{f} := (f_1, \ldots, f_n, g)$ has algebraic rank n and $r(v_1, \ldots, v_n, u)$ is a power of the minimal $\mathbf{f}$-annihilating polynomial.

Proof.

Claim 9.1. *The set of polynomials $f_1, \ldots, f_n, g$ has algebraic rank n.*
Proof of Claim 9.1: Observe that the polynomial f_i depends only on the variables $x_1, \ldots, x_i$ and by being monic in x_i we also have that $\frac{\partial f_i}{\partial x_i}$ is nonzero. Thus the $n \times n$ partial derivative matrix

$$J \stackrel{\text{def}}{=} \left(\left(\frac{\partial f_i}{\partial x_j} \right) \right)_{n \times n}$$

is a lower triangular matrix with non-zero entries on the diagonal. This means that J is non-singular and by Theorem 2 we get that $f_1, \ldots, f_n$ are algebraically independent. [1] Thus the rank

[1] This proof fails when $\mathbb{F}$ has a small characteristic but nevertheless it can be shown that $f_1, \ldots, f_n$ are algebraically independent over any field $\mathbb{F}$.

of $f_1, \ldots, f_n, g$ is at least n. Moreover $f_1, \ldots, f_n, g$ being polynomials in n variables, their rank is most n. Thus this set of polynomials has rank exactly n. □

Since the minimal annihilating polynomial remains unchanged when we move from $\mathbb{F}$ to $\overline{\mathbb{F}}$ (Lemma 7), we can assume without loss of generality that $\mathbb{F}$ itself is algebraically closed. We now claim that:

Claim 9.2. *If $(b_1, \ldots, b_n, a) \in \overline{\mathbb{F}}^{n+1}$ is any zero of $r(v_1, \ldots, v_n, u)$ then it is also a zero of $A(v_1, \ldots, v_n, u)$.*

Assuming this claim, and using the absolute irreducibility of $A(v_1, \ldots, v_n, u)$ (lemma 7), we get the desired result immediately from an application of the Nullstellensatz (Lemma 6).

Proof of Claim 9.2: Note that $A(v_1, \ldots, v_n, u)$ is the minimal $(f_1, \ldots, f_n, g)$-annihilating polynomial if and only if $A(v_1 + b_1, \ldots, v_n + b_n, u + a)$ is the minimal $(f_1 - b_1, \ldots, f_n - b_n, f_{n+1} - a)$-annihilating polynomial. Replacing f_i by $f_i - b_i$ and g by $g - a$ throughout, we can assume without loss of generality that $(b_1, \ldots, b_n, a) = (0, \ldots, 0, 0)$. Let

$$\mathtt{V} \stackrel{\text{def}}{=} \{\mathbf{a} := (\alpha_1, \ldots, \alpha_n) \in \overline{\mathbb{F}}^n \mid f_1(\mathbf{a}) = \ldots = f_n(\mathbf{a}) = 0\}$$

be the set of common zeroes to the system of equations $f_1 = \ldots = f_n = 0$ in $\overline{\mathbb{F}}$. Note that by assumption, this system of equations is a diagonal system of equations with each f_i being monic with respect to x_i. Thus the number of points in $\mathtt{V}$ is precisely $\prod_{i \in [n]} d_i$. In particular $\mathtt{V}$ is finite and non-empty. Let R be the ring

$$R \stackrel{\text{def}}{=} \mathbb{F}[x_1, \ldots, x_n]/\langle f_1(x_1), \ldots, f_n(x_1, \ldots, x_n)\rangle.$$

Then the ring R is isomorphic to the ring

$$\bigoplus_{(\alpha_1, \ldots, \alpha_n) \in \mathtt{V}} \mathbb{F}[x_1, \ldots, x_n]/\langle x_1 - \alpha_1, \ldots, x_n - \alpha_n\rangle.$$

Let θ be the isomorphism mapping elements of R to the direct-sum ring above. Viewing $g(x_1, \ldots, x_n)$ as an element of R via the canonical map $g(\mathbf{x}) \mapsto g(\mathbf{x}) \pmod{f_1(\mathbf{x}), \ldots, f_n(\mathbf{x})}$, $\theta(g)$ equals

$$\bigoplus_{(\alpha_1, \ldots, \alpha_n) \in \mathtt{V}} g(x_1, \ldots, x_n) \pmod{x_1 - \alpha_1, \ldots, x_n - \alpha_n} \;=\; \bigoplus_{(\alpha_1, \ldots, \alpha_n) \in \mathtt{V}} g(\alpha_1, \ldots, \alpha_n).$$

Thus the set of eigenvalues of the linear transformation corresponding to multiplication by g in the ring R is precisely the set

$$\{g(\alpha_1, \ldots, \alpha_n) \mid (\alpha_1, \ldots, \alpha_n) \in \mathtt{V}\}.$$

This means that the characteristic polynomial of $g(x_1, \ldots, x_n)$ in the ring R is

$$\prod_{(\alpha_1, \ldots, \alpha_n) \in \mathtt{V}} (u - g(\alpha_1, \ldots, \alpha_n)).$$

That is

$$r(0, \ldots, 0, u) = \prod_{(\alpha_1, \ldots, \alpha_n) \in \mathtt{V}} (u - g(\alpha_1, \ldots, \alpha_n)).$$

By assumption of the claim $r(0, \ldots, 0, 0) = 0$ and therefore there exists

$$(\alpha_1, \ldots, \alpha_n) \in \mathtt{V} \text{ such that } g(\alpha_1, \ldots, \alpha_n) = 0.$$

This means that the point $P \stackrel{\text{def}}{=} (\alpha_1, \ldots, \alpha_n) \in \overline{\mathbb{F}}^n$ is a common zero to the system of equations

$$f_1 = \ldots = f_n = g = 0.$$

By lemma 8 we get $A(0, \ldots, 0, 0) = 0$ as required. □

This completes the proof of the theorem.

□

4 Degree bounds for the annihilating polynomial.

4.1 Upper bound for the degree of the annihilating polynomial.

By a dimension counting argument followed by induction, [DGW07] showed that

Theorem 10. *[DGW07] Let $\mathbf{f} = (f_1, \ldots, f_k))$ be a set of algebraically dependent polynomials of degree d in n variables over the field $\mathbb{F}$. Then there exists an $\mathbf{f}$-annihilating polynomial of degree at most $D = (n+1) \cdot d^n$.*

We will improve the bound on the degree D to $(r+1) \cdot d^r$, where r is the algebraic rank of $\mathbf{f}$. Note that $r \leq n$ and thus our bound is an improvement over the previous bound.

Theorem 11. *Let $\mathbb{F}$ be a field and let $\mathbf{f} = (f_1, \ldots, f_{k+1})$ be a set of polynomials of degree d in n variables over the field $\mathbb{F}$ having algebraic rank r. Then there exists an $\mathbf{f}$-annihilating polynomial of degree at most $D = (r+1) \cdot d^r$.*

Proof. See appendix. □

4.2 Lower bound for the degree of the annihilating polynomial.

Theorem 12. *Let $\mathbb{F}$ be any field of characteristic zero. For any $d \in \mathbb{Z}_{\geq 1}$, there exists a set of polynomials $f_1, \ldots, f_n, g \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most d and algebraic rank n such that the minimal $(f_1, \ldots, f_n, g)$-annihilating polynomial has degree at least d^n.*

Proof. Define the polynomials f_i's as follows. For each $i \in [n]$,

$$f_i \stackrel{\text{def}}{=} x_i^d - 1 \quad \text{and} \quad g \stackrel{\text{def}}{=} x_1 + \ldots + x_n - n.$$

We shall denote by ω a primitive d-th root of unity. Then for each $i \in [n]$, we have

$$f_i(x_i) = \prod_{j \in [d]} (x_i - \omega^j).$$

Let $A(v_1, \ldots, v_n, u)$ be the minimal $(f_1, \ldots, f_n, g)$-annihilating polynomial. We claim that $A(v_1, \ldots, v_n, u)$ is precisely the charecterictic polynomial

$$r(v_1, \ldots, v_n, u) \stackrel{\text{def}}{=} \text{charpoly}_{g(x) \bmod (f_1(x_1) - v_1, \ldots, f_n(x_n) - v_n)}(u) \in \mathbb{F}[v_1, \ldots, v_n][u].$$

By lemma 9 we get that

$$r(v_1,\ldots,v_n,u) = A(v_1,\ldots,v_n,u)^t \text{ for some integer } t \geq 1.$$

Thus we need to show that $t = 1$, or equivalently that r is not a proper power. We show this by showing that $r(0,\ldots,0,u)$ is not a proper power. This is in turn shown by showing that u divides $r(0,\ldots,0,u)$ but u^2 does not divide it. To see this note that

$$\begin{aligned} r(0,\ldots,0,u) &= \text{charpoly}_{g \ (\text{mod } f_1,\ldots,f_n)}(u) & (1)\\ &= \prod_{(i_1,\ldots,i_n)\in[d]^n} (u - (\omega^{i_1} + \ldots + \omega^{i_n} - n)) & (2)\end{aligned}$$

Also

$$\omega^{i_1} + \ldots + \omega^{i_n} = n \text{ if and only if } i_1 = \ldots = i_n = 0 \ (\text{mod } d).$$

and thus u exactly divides $r(0,\ldots,0,u)$. Consequently, as argued above we get $t = 1$ and

$$r(v_1,\ldots,v_n,u) = A(v_1,\ldots,v_n,u).$$

Thus we get

$$deg(A) = deg(r) \geq deg_u(r) \geq d^n \text{ (By eqn (2))}.$$

Thus any annihilating polynomial for the above set of polynomials as degree at least d^n. □

Thus our lower bound on the degree of the annihilating polynomial of a set of algebraically dependent polynomials of algebraic rank n and degree d closely (upto a factor of n) matches the upper bound.

5 Complexity bounds for the annihilating polynomial.

Suppose that some set of polynomials $f_1,\ldots,f_k \in \mathbb{F}[\mathbf{x}]$ is algebraically dependent of rank $k-1$ and we wish to compute their unique minimal monic annihilating polynomial $A(t_1,\ldots,t_k)$. We will work over the field of rational numbers. We can then assume that the input polynomials all have having integer coefficients. Our degree lower bounds show that the annihilating polynomial itself may have exponential degree. This means that for an integer point $\mathbf{a} \in \mathbb{Z}^n$, one may require an exponential number of digits to write down the number $A(\mathbf{a})$. So we need to define the problem of computing the minimal annihilating polynomial more carefully. So we shall define the computational problem of computing the annihilating polynomial as computing the value $A(\mathbf{a})(\text{mod } p)$ for a given integer $p \in \mathbb{Z}$. The motivation for defining it this way is that if $A(\mathbf{t})$ admitted a small (polynomial in the input size) arithemtic circuit representation that could be computed efficiently then we would also have an efficient algorithm for computing $A(\mathbf{a})(\text{mod } p)$. Let us make this more precise by defining two concrete computational problems.

Definition 13. ANNIHILATING-AT-ZERO shall denote the following decision problem: Given a set of polynomials $f_1,\ldots,f_k \in \mathbb{F}[x_1,\ldots,x_n]$ of algebraic rank $(k-1)$, having minimal annihilating polynomial $A(t_1,\ldots,t_k)$, determine if $A(0,\ldots,0)$ equals zero or not.

Definition 14. ANNIHILATING-EVALUATION is the functional problem of evaluating the annihilating polynomial at a given point. The input consists of

- A set of polynomials $f_1,\ldots,f_k \in \mathbb{Z}[x_1,\ldots,x_n]$ with integer coefficients having the minimal monic annihilating polynomial $A(t_1,\ldots,t_k) \in \mathbb{Z}[t_1,\ldots,t_k]$.

- A prime p.

The output is the integer $A(0, \ldots, 0) \pmod{p}$.

For both these problems the input polynomials are assumed to be given using the dense representation. The degree bounds for the annihilating polynomial imply that both these problems lie in the complexity class PSPACE. We will now show that over the field $\mathbb{Q}$, the ANNIHILATING-AT-ZERO problem is NP-hard and ANNIHILATING-EVALUATION is #P-hard. The proof can be easily adapted to work for any finite field, including fields of small characteristic.

Theorem 15. *ANNIHILATING-AT-ZERO is NP-hard and ANNIHILATING-EVALUATION is #P-hard.*

Proof. See appendix. □

Having shown that it is not possible to evaluate minimum annihilating polynomial at any point unless P = NP, let us examine if the annihilating polynomial admits polynomial sized circuit, even if these circuits may be difficult to compute. We will show this by observing that if this happens then coNP $\subseteq$ AM and therefore the polynomial hierarchy collapses.

Theorem 16. *Unless coNP $\subseteq$ AM , there exist a set of algebraically dependent cubic polynomials whose minmal annihilating polynomial has superpolynomial circuit complexity.*

The proof of this theorem is by a very similar reduction.

These complexity-theoretic lower bounds improve the gap in our understanding of the computational complexity of computing the annihilating polynomial. The best upper bounds that we have for both ANNIHILATING-AT-ZERO as well as for the general problem ANNIHILATING-EVALUATION is PSPACE. This then invites work on the following questions.

Question 1. The NP-hardness result of ANNIHILATING-AT-ZERO was via the problem of determining if a given system of polynomial equations over integers have a common solution in $\mathbb{C}$ or not. This later problem, famously known as the Nullstellensatz problem was shown to be in AM under the generalized Riemann Hypothesis by Koiran [Koi96]. Is it then true that the ANNIHILATING-AT-ZERO problem is also in AM , assuming the generalized Riemann Hypothesis?

Question 2. Is ANNIHILATING-EVALUATION in #P? ... or is it PSPACE-complete?

The proofs of theorem 2 and its algorithmic Corollary 3 are valid only over fields of characteristic $p > (r+1) \cdot d^r$.

Question 3. ([DGW07]:) Does there exist an RP algorithm for testing algebraic dependence over any finite field $\mathbb{F}_q$?

Acknowledgements

We would like to thank Avi Wigderson for posing the problems that led to this research, for suggesting avenues of investigation and for many helpful conversations. We also want to thank Vladimir Trifonov and Jonathan Kelner for helpful conversations and Zeev Dvir for commenting on an earlier draft of this work.

References

[BS83] Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22:317–330, 1983.

[Den78] J. Denef. The Diophantine problem for polynomial rings and fields of rational functions. *Transactions of the American Mathematical Society*, 242:391–399, Aug 1978.

[DGW07] Z. Dvir, A. Gabizon, and A. Wigderson. Extractors and rank extractors for polynomial sources. In *Proceesings of FOCS 2007*, 2007.

[ER93] Richard Ehrenborg and Gian-Carlo Rota. Apolarity and canonical forms for homogeneous polynomials. *European Journal of Combinatorics*, 14(3):157–181, 1993.

[Har77] Robin Hartshorne. *Algebraic Geometry*, volume 52 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1977.

[Koi96] Pascal Koiran. Hilbert's nullstellensatz is in the polynomial hierarchy. *Journal of Complexity*, 12(4):273–286, 1996.

[KR92] K. H. Kim and F. W. Roush. Diophantine undecidability of $\mathbf{C}(t_1, t_2)$. *Journal of Algebra*, 150(1):35–44, 1992.

[L'v84] M.S. L'vov. Calculation of invariants of programs interpreted over an integrality domain. *Kibernetika*, 4:23–28, 1984.

[Mor85] Jacques Morgenstern. How to compute fast a function and all its derivatives, A variation on the theorem of Baur-Strassen. *SIGACT News*, 16:60–62, 1985.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980.

[Sch82] A. Schinzel. *Selected Topics on Polynomials*. University of Michigan Press, Ann Arbor, MI., 1982.

[Vid94] Carlos R Videla. Hilbert's tenth problem for rational function fields in characteristic 2. *Proceedings of the American Mathematical Society*, 120(1):249–253, 1994.

[Zip90] R. Zippel. Interpolating polynomials from their values. *JSC*, 9(3):375–403, March 1990.

Appendix

Here we give the proofs that were omitted above due to lack of space.

Lemma 7. Let $f_1, f_2, \ldots, f_k \in \mathbb{F}[x_1, \ldots, x_n]$ be a set of *algebraically dependent* polynomials in n variables over the field $\mathbb{F}$, no proper subset of which is algebraically dependent. Then the ideal $\mathfrak{U}$ of $\mathbf{f}$-annhilating polynomials is generated by a single absolutely irreducible polynomial $A(\mathbf{t})$. Moreover, $A(\mathbf{t})$ remains the minimal annihilating polynomial of $\{f_1, f_2, \ldots, f_k\}$ when they are viewed as polynomials over the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$.

Proof of Lemma 7:

Let $A(t_1, \ldots, t_k) \in \mathfrak{U} \subseteq \mathbb{F}[t_1, \ldots, t_k]$ be a minimal degree $(f_1, \ldots, f_k)$-annihilating polynomial. We shall prove that this polynomial $A(\mathbf{t})$ is unique and it satisfies the properties claimed in the theorem. We carry out this proof through a sequence of observations. We begin with the simple observation that $A(\mathbf{t})$ must be $\mathbb{F}$-irreducible.

Claim 7.1. *$A(\mathbf{t})$ is $\mathbb{F}$-irreducible.*

Proof of Claim 7.1: If $A(t_1, \ldots, t_k)$ is $\mathbb{F}$-reducible then

$$A(t_1, \ldots, t_k) = a_1(t_1, \ldots, t_k) \cdot A_2(t_1, \ldots, t_k),$$

where A_1 and A_2 both have degrees smaller than A. Now since

$$A(f_1, \ldots, f_k) = 0$$

therefore either

$$A_1(f_1, \ldots, f_k) = 0 \text{ or } A_2(f_1, \ldots, f_k) = 0.$$

In either case, we get an annihilating polynomial of smaller degree, a contradiction. □

Now suppose that $\deg(A(\mathbf{t})) = d$. Can there exist an $\mathbf{f}$-annihilating polynomial with coefficients from the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$ which has degree smaller than d? We next observe that this is not possible.

Claim 7.2. *Any minimal degree $\mathbf{f}$-annihilating polynomial $B(t_1, \ldots, t_k)$ over the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$ has degree at least d.*

Proof of Claim 7.2: Suppose to the contrary that there exists a $B(t_1, \ldots, t_k) \in \overline{\mathbb{F}}[t_1, \ldots, t_k]$ of total degree less than d such that

$$B(f_1, \ldots, f_k) = 0. \tag{3}$$

Then the coefficients of various monomials appearing in B all come from some finite extension field $\mathbb{K} \supseteq \mathbb{F}$ of dimension say ℓ over $\mathbb{F}$. Any finite extension field of $\mathbb{F}$ is generated by some single primitive element and so suppose that $\mathbb{K} = \mathbb{F}(\gamma)$ where γ is algebraic over $\mathbb{F}$ with minimal polynomial of degree ℓ.

Then $\{1, \gamma, \gamma^2, \ldots, \gamma^{\ell-1}\}$ form a basis of $\mathbb{K}$ over $\mathbb{F}$ and consequently B can be expressed as

$$B(\mathbf{t}) = B_0(\mathbf{t}) + \gamma \cdot B_1(\mathbf{t}) + \ldots + \gamma^{\ell-1} B_{\ell-1}(\mathbf{t}),$$

where the $B_i(\mathbf{t})$'s are polynomials in $\mathbb{F}[\mathbf{t}]$. Since

$$B(\mathbf{f}) = 0 \text{ (by equation (3))}$$

therfore we get that for all $i \in [\ell]$:

$$B_i(\mathbf{f}) = 0$$

By assumption the degree of $B(\mathbf{t})$ is less than d and therefore the degree of every $B_i(\mathbf{t})$ is also less than d. By the assumption that $A(\mathbf{t})$ is a minimal degree $\mathbf{f}$-annihilating polynomial we get that any annihilating polynomial over $\mathbb{F}$ has degree at least d and therefore every $B_i(\mathbf{t})$ is zero and thus $B(\mathbf{t}) = 0$, a contradiction. □

Note that this means that $A(\mathbf{t})$ is also absolutely irreducible for if it were to factor over $\overline{\mathbb{F}}$ then we would get an $\mathbf{f}$-annihilating polynomial over $\overline{\mathbb{F}}$ of smaller degree (as in Claim 7.1). The next observation gives a characterization of the ideal $\mathfrak{U}$ of annihilating polynomials. It says that the ideal $\mathfrak{U}$ is a *principal ideal* is of the form $\mathfrak{U} = \langle A(\mathbf{t}) \rangle$ and thus any minimal degree polynomial in $\mathfrak{U}$ is a constant multiple of $A(\mathbf{t})$.

Claim 7.3. *If $B(t_1, \ldots, t_k) \in \mathbb{F}[t_1, \ldots, t_k]$ is any $\mathbf{f}$-annihilating polynomial then $A(\mathbf{t})$ divides $B(\mathbf{t})$.*

Proof of Claim 7.3: Note that $A(\mathbf{t})$ does depend on t_1 for it not then $\{f_2, \ldots, f_k\}$ which is a proper subset of $\{f_1, \ldots, f_k\}$ forms an algebraically dependent set contrary to the assumption of the theorem. Let

$$\rho(t_2, \ldots, t_k) \overset{\text{def}}{=} \textsc{Res}_{t_1}(A(\mathbf{t}), B(\mathbf{t})).$$

Suppose if possible that $\rho(t_2, \ldots, t_k) \neq 0$. By the resultant properties, there exist polynomials $\hat{A}(\mathbf{t})$ and $\hat{B}(\mathbf{t})$ such that

$$\rho(t_2, \ldots, t_k) \;=\; \hat{A}(\mathbf{t}) \cdot A(\mathbf{t}) + \hat{B}(\mathbf{t}) \cdot B(\mathbf{t}).$$

Making the substitution $(t_2, \ldots, t_k) := (f_2, \ldots, f_k)$, we get

$$\rho(f_2, \ldots, f_k) = 0,$$

which contradicts the algebraic independence of $f_2, \ldots, f_k$. This means that our assumption $\rho(t_2, \ldots, t_k)$ is untenable and it must be that $\rho = 0$. But this happens if and only if $A(\mathbf{t})$ and $B(\mathbf{t})$ share a common factor. By the $\mathbb{F}$-irreducibility of $A(\mathbf{t})$ (Claim 7.1) we get that $A(\mathbf{t})$ divides $B(\mathbf{t})$. □

This completes the proof of the lemma. □

Theorem 11. Let $\mathbb{F}$ be a field and let $\mathbf{f} = (f_1, \ldots, f_{k+1})$ be a set of polynomials of degree d in n variables over the field $\mathbb{F}$ having algebraic rank r. Then there exists an $\mathbf{f}$-annihilating polynomial of degree at most $D = (r+1) \cdot d^r$.

Proof of Theorem 11:

We begin with a claim.

Claim 11.1. *Let $\mathbb{F}$ be an algebraically closed field and let $\mathbf{f} = (f_1, \ldots, f_k)$ be a set polynomials of degree d in n variables over the field $\mathbb{F}$ having algebraic rank r. Then there exists a homomorphism $\sigma : \mathbb{F}[x_1, \ldots, x_n] \mapsto \mathbb{F}[y_1, \ldots, y_r]$ of the form*

$$\sigma : x_i \mapsto a_{i1} \cdot y_1 + a_{i2} \cdot y_2 + \ldots + a_{ir} \cdot y_r + b_i$$

such that the set of polynomials $\{\sigma(f_1), \ldots, \sigma(f_k)\} \subset \mathbb{F}[y_1, \ldots, y_r]$ has algebraic rank r.

Proof of Claim 11.1:

We can assume without loss of generality that $f_1, \ldots, f_r$ are algebraically independent polynomials. We will choose the homomorphism σ such that $\sigma(f_1), \ldots, \sigma(f_r)$ are algebraically independent. This will ensure that $\{\sigma(f_1), \ldots, \sigma(f_k)\} \subset \mathbb{F}[y_1, \ldots, y_r]$ has algebraic rank r. Since $\{f_1, \ldots, f_r\}$ has rank r, therefore there must exist indices $\{j_1, j_2, \ldots, j_r\} \subset [n]$ such that

$$J(x_1, \ldots, x_n) \stackrel{\text{def}}{=} \textsc{Det} \begin{pmatrix} \frac{\partial f_1}{\partial x_{j_1}} & \cdots & \frac{\partial f_r}{\partial x_{j_1}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1}{\partial x_{j_r}} & \cdots & \frac{\partial f_r}{\partial x_{j_r}} \end{pmatrix}$$

is a nonzero polynomial in the variables $x_1, \ldots, x_n$. By renaming the variables if necessary we can assume without loss of generality that $(j_1, \ldots, j_r) = (1, \ldots, r)$. Since $J(\mathbf{x})$ is a nonzero polynomial in $\mathbb{F}[\mathbf{x}]$, there exist elements $b_1, \ldots, b_{n-r} \in \mathbb{F}$ such that $J(x_1, \ldots, x_r, b_1, \ldots, b_{n-r}) \neq 0$. Fix these elements $b_1, \ldots, b_{n-r}$. Consider the homomorphism $\sigma : \mathbb{F}[x_1, \ldots, x_n] \mapsto \mathbb{F}[y_1, \ldots y_r]$ defined by

$$\sigma(x_j) = \begin{cases} y_j & \text{if } 1 \leq j \leq r, \\ b_{j-r} & \text{if } j > r. \end{cases} \tag{4}$$

For $i \in [r]$, let

$$g_i(y_1, \ldots, y_r) \stackrel{\text{def}}{=} \sigma(f_i(x_1, \ldots, x_n)).$$

Its easy to verify that for all $i, j \in [r]$,

$$\frac{\partial g_i}{\partial y_j} = \sigma\left(\frac{\partial f_i}{\partial x_j}\right).$$

Consequently,

$$\begin{aligned} \textsc{Det} \begin{pmatrix} \frac{\partial g_1}{\partial y_1} & \cdots & \frac{\partial g_r}{\partial y_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial g_1}{\partial y_r} & \cdots & \frac{\partial g_r}{\partial y_r} \end{pmatrix} &= \textsc{Det} \begin{pmatrix} \sigma\left(\frac{\partial f_1}{\partial x_1}\right) & \cdots & \sigma\left(\frac{\partial f_r}{\partial x_1}\right) \\ \vdots & \ddots & \vdots \\ \sigma\left(\frac{\partial f_1}{\partial x_r}\right) & \cdots & \sigma\left(\frac{\partial f_r}{\partial x_r}\right) \end{pmatrix} \\ &= \sigma(\textsc{Det} \begin{pmatrix} \left(\frac{\partial f_1}{\partial x_1}\right) & \cdots & \left(\frac{\partial f_r}{\partial x_1}\right) \\ \vdots & \ddots & \vdots \\ \left(\frac{\partial f_1}{\partial x_r}\right) & \cdots & \left(\frac{\partial f_r}{\partial x_r}\right)) \end{pmatrix}) \\ &= \sigma(J(\mathbf{x})) \\ &= J(y_1, \ldots, y_r, b_1, \ldots, b_{n-r}) \\ &\neq 0. \end{aligned}$$

Thus the σ chosen above ensures that $\sigma(f_1), \ldots, \sigma(f_r)$ are algebraically independent and consequently that $\sigma(f_1), \ldots, \sigma(f_k)$ have algebraic rank r. □

By induction on k. If the rank r is less than $k-1$ then there exists an $(r+1)$-sized subset of polynomials in $\{f_1, \ldots, f_k\}$ which are algebraically dependent and the theorem would be proved by the inductive assumption. Henceforth we shall assume $r = (k-1)$. Let $A(t_1, \ldots, t_k)$ be the minimal $\mathbf{f}$-annihilating polynomial. By lemma 7, $A(\mathbf{t})$ is absolutely irreducible. By lemma 11.1 there exists a rank-preserving homomorphism $\sigma : \mathbb{F}[x_1, \ldots, x_n] \mapsto \mathbb{F}[y_1, \ldots, y_r]$ in which every variable x_i is mapped to an affine combination of the y_j's. For $i \in [k]$, let

$$g_i(y_1, \ldots, y_r) \overset{\text{def}}{=} \sigma(f_i) \in \mathbb{F}[y_1, \ldots, y_k]$$

and

$$\mathbf{g} \overset{\text{def}}{=} (g_1, \ldots, g_k).$$

This implies that $deg(\sigma(f_i)) \leq d$ for all $i \in [k]$. By theorem 10 there exists an absolutely irreducible, $\mathbf{g}$-annihilating polynomial $B(t_1, \ldots t_k)$ of degree at most D. That is,

$$\deg(B) \leq (r+1) \cdot d^r. \tag{5}$$

Now note that

$$A(f_1, \ldots, f_k) = 0$$
$$\Rightarrow \sigma(A(f_1, \ldots, f_k)) = 0$$

and since σ is a homomorphism that is fixed everywhere on $\mathbb{F}$, we have

$$A(\sigma(f_1), \ldots, \sigma(f_k)) = 0.$$

This means that $A(\mathbf{t})$ is also a $\mathbf{g}$-annihilating polynomial. Since the ideal of $\mathbf{g}$-annihilating polynomials is a principal ideal (Lemma 7), $B(\mathbf{t})$ divides $A(\mathbf{t})$. But $A(\mathbf{t})$ is absolutely irreducible and therefore $B(\mathbf{t})$ equals $A(\mathbf{t})$ up to constant factors. This means that

$$\begin{aligned} \textsc{Deg}(A(\mathbf{t})) &= \textsc{Deg}(B(\mathbf{t})) \\ &\leq (r+1) \cdot d^r \quad \text{(by (5))}, \end{aligned}$$

as was required. □

Theorem 15. ANNIHILATING-AT-ZERO is NP-hard and ANNIHILATING-EVALUATION is #P-hard.

Proof of Theorem 15: Let us start out with an observation. Note that for any $(a_1, \ldots, a_k) \in \mathbb{Z}^k$, $A(t_1, \ldots, t_k)$ is the minimal $(f_1, \ldots, f_k)$-annihilating polynomial if and only if $A(t_1+a_1, \ldots, t_k+a_k)$ is the minimal $(f_1 - a_1, \ldots, f_k - a_k)$-annihilating polynomial. Thus the problem of computing $A(a_1, \ldots, a_k)(\text{mod } p)$ for a given point $(a_1, \ldots, a_k) \in \mathbb{Z}^k$ is equivalent to ANNIHILATING-EVALUATION.

For clarity of presentation we shall first describe these reductions using arithmetic formula representation of polynomials and then make remarks towards the end on how to do this for the usual dense representation of polynomials. We will give complexity lower bounds for these problems via reduction from 3SAT and #3SAT respectively. Let us therefore define a very usual and natural "algebraization" of a 3CNF formula. Let $\Phi(b_1, \ldots, b_n) := \phi_1 \wedge \ldots \wedge \phi_m$ be a boolean 3CNF formula in the boolean variables $b_1, \ldots, b_n$ having m clauses. Corresponding to Φ we shall construct a polynomial $G_\Phi(x_1, \ldots, x_n)$ such that $\overline{\Phi(\mathbf{b})} = G_\Phi(\mathbf{b})$ for all $\mathbf{b} \in \{0,1\}^n$. The polynomial $G_\Phi(\mathbf{x})$ shall equal $1 - \prod_{i \in [m]} g_i(\mathbf{x})$ where each $g_i(\mathbf{x})$ is a multilinear cubic polynomial polynomial with the property that for any boolean setting $\mathbf{b} \in \{0,1\}^n$ of the variables $\mathbf{x}$, $g_i(\mathbf{b}) = \phi_i(\mathbf{b}) \in \{0,1\}$. For a clause ϕ_i of the form $\phi_i = x_1 \vee x_2 \vee x_3$, the corresponding polynomial g_i is given by

$$g_i(\mathbf{x}) := x_1x_2x_3 - (x_1x_2 + x_2x_3 + x_3x_1) + (x_1 + x_2 + x_3).$$

For a negated literal $\neg x_j$ occuring in ϕ_i, we simply replace x_j by $1 - x_j$ in g_i. This construction gives the following property of the polynomial $G_\Phi(\mathbf{x})$:

$$\forall \mathbf{a} \in \{0,1\}^n, \quad G_\Phi(\mathbf{a}) = \begin{cases} 0 & \text{if} \Phi(\mathbf{a}) = 1, \\ 1 & \text{if} \Phi(\mathbf{a}) = 0. \end{cases} \tag{6}$$

Let the number of Φ-satisfying assignments be

$$N \overset{\text{def}}{=} |\{\mathbf{b} \in \{0,1\}^n \mid \Phi(\mathbf{b}) = 1\}|.$$

From the property (6), it immediately follows that

$$N = |\{\mathbf{b} \in \{0,1\}^n \mid G_\Phi(\mathbf{b}) = 0\}|. \tag{7}$$

For all $i \in [n]$, let us also define auxilliary polynomials $f_i(x_i) \overset{\text{def}}{=} x_i^2 - x_i$. Let $A(v_1, \ldots, v_n, u)$ be the minimal degree monic $(f_1, \ldots, f_n, G_\Phi)$-annihilating polynomial. Also let

$$r(v_1, \ldots, v_n, u) \overset{\text{def}}{=} \text{charpoly}_{G_\Phi(\mathbf{x}) \bmod (f_1(x_1) - v_1, \ldots, f_n(x_n) - v_n)}(u) \in \mathbb{F}(v_1, \ldots, v_n)[u]$$

Then $f_1(x_1), \ldots, f_n(x_n)$ form a diagonal system so that by lemma 9 we have

$$r(v_1, \ldots, v_n, u) = A(v_1, \ldots, v_n, u)^t \quad \text{for some } t \geq 1 \tag{8}$$

The reduction of 3SAT to ANNIHILATING-AT-ZERO simply involves outputting $f_1, \ldots, f_n, G_\Phi$. The correctness of the reduction involves the following claim.

Claim 15.1. *$A(0, \ldots, 0, 0) = 0$ if and only if Φ is satisfiable.*

Proof of Claim 15.1: By equation 6 we get that Φ is satisfiable if and only if the system of equations

$$f_1(x_1) = f_2(x_2) = \ldots = f_n(x_n) = G_\Phi(x_1, \ldots, x_n) = 0 \tag{9}$$

has a common solution. From the proof of lemma 9 we also get that the above system of equations (9) has a solution if and only if $r(0, \ldots 0, 0) = 0$. By equation (8) we have $r(0, \ldots, 0, 0) = 0$ if and only if $A(0, \ldots, 0, 0) = 0$. Putting these equivalences together we have that Φ is satisfiable if and only if $A(0, \ldots, 0, 0) = 0$. □

This shows that ANNIHILATING-AT-ZERO is NP-hard. Let us go further and see how evaluation of $A(v_1, \ldots, v_n, u)$ at various points can allow us to count the number of Φ-satisfying assignments. For this, we need to show that in this reduction the annihilating polynomial $A(\mathbf{v}, u)$ in fact equals the characteristic polynomial $r(\mathbf{v}, u)$. Towards that direction we first make an enabling observation. Since the polynomial $x_i^2 - x_i$ has roots $x_i = 0$ and 1, by the definition of the characteristic polynomial r, we get that

$$\begin{aligned} r(0, \ldots, 0, u) &= \prod_{\mathbf{a} \in \{0,1\}^n} (u - G_\Phi(\mathbf{a})) \\ \text{or,} \;\; A(0, \ldots, 0, u)^t &= \prod_{\mathbf{a} \in \{0,1\}^n} (u - G_\Phi(\mathbf{a})) \;\; \text{(By equation 8)} \\ &= u^N \cdot (u-1)^{2^n - N} \;\; \text{(By equation (6))} \end{aligned}$$

This implies that t divides N as well as $(2^n - N)$ and also that

$$A(0, \dots, 0, 2)^t = 2^N \quad \text{and} \quad A(0, \dots, 0, -1)^t = 2^{2^n - N}. \tag{10}$$

Let $M_1 \stackrel{\text{def}}{=} \frac{N}{t}$ and $M_2 \stackrel{\text{def}}{=} \frac{2^n - N}{t}$.

Claim 15.2. *Given an oracle for evaluating $A(0, \dots, 0, 2)(\mathit{mod}\ p)$ for a few small primes p, we can compute the integer M_1.*

Proof of Claim 15.2: Our strategy to compute M_1 using an oracle for ANNIHILATING-EVALUATION is by evaluating $M_1(\text{mod}\ \ell)$ for several small $(\Theta(n \log n))$ values of ℓ and thereby obtaining M_1 through chinese remaindering, knowing that M_1 must lie in the range $0 \leq M_1 \leq 2^n$. Let p be an odd integer. ℓ shall be the order of $2(\text{mod}\ p)$. In time $\Theta(p)$ we can compute ℓ and the smallest integer L such that $A(0, \dots, 0, 2) = 2^L \pmod{p}$. From this we can deduce that $M_1 = L \pmod{\ell}$. Doing this for all odd p between 3 to $\Theta(n \log n)$, we collect enough information about M_1 modulo various integers ℓ that we can compute the actual value of M_1. The details of this last calculation are routine and we leave it to the reader to verify. □

Similarly, using the values of $A(0, \dots, 0, -1)(\text{mod}\ p)$ for a few p, we can compute the integer M_2. Then the desired number of Φ-satisfying assignments N is simply given by

$$N = \frac{M_1}{M_1 + M_2} 2^n.$$

This completes the complexity lower bound proof for ANNIHILATING-AT-ZERO and ANNIHILATING-EVALUATION for polynomials represented as arithmetic formulas. Finally, we note that all the polynomials involved in the above reduction were quadratic except for the polynomial $G_\Phi(\mathbf{x})$ which had degree $3m$. We can use usual trick of introducing additional variables for denoting the 'intermediate steps in the computation of $G_\Phi(\mathbf{x})$' and then adjoining these additional polynomials to the input of the oracle call. This makes all the polynomials involved in the reduction cubic and gives the same complexity lower bound for densely represented polynomials.

This completes the proof of the theorem. □