Computational Complexity of the Nullstellensatz over Finite Fields

Sahil Goyat

220937

Project Supervisor: Prof. Nitin Saxena

Undergraduate Project Report (CS496)

Department of Computer Science and Engineering
Indian Institute of Technology, Kanpur
Uttar Pradesh, India

sahilkg22@iitk.ac.in

Abstract

The computational problem of Hilbert's Nullstellensatz (HN), has a well-defined upper bound of PSPACE, derived from classical effective Nullstellensatz. A significant refinement by Koiran (1996), assuming the Generalized Riemann Hypothesis (GRH), places HN for integer-coefficient polynomials within Arthur-Merlin (AM) complexity class. We investigate the feasibility of extending Koiran's complexity-theoretic advancements to the analogous problem over finite fields. Our analysis attempts to place the positive-dimensional case into AM by establishing the existence of a "short" solution. We also discuss why this method fails for the zero-dimensional case, necessitating alternative approaches for providing a short certificate of satisfiability.

Contents

1 Background and Effective Nullstellensatz

1.1 Basic Definitions

We briefly recall some basic notions from affine algebraic geometry which will be used throughout the report.

Definition 1.1 (Affine algebraic variety)**.** Let k be a field. A subset $X \subseteq k^n$ is called an *affine algebraic variety* (or simply an *affine variety*) if there exist polynomials $f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$ such that

$$X = V(f_1, \ldots, f_m) := \{x \in k^n : f_i(x) = 0 \text{ for all } i\}.$$

In other words, affine varieties are precisely the common zero sets of finitely many polynomials.

Definition 1.2 (Irreducible variety)**.** A variety $X \subseteq k^n$ is said to be *irreducible* if it cannot be written as a union $X = X_1 \cup X_2$ of two proper subvarieties $X_1, X_2 \subsetneq X$.

1.2 Dimension and Degree

Definition 1.3 (Dimension)**.** Let X be an irreducible affine variety. The *dimension* $\dim X$ of X is defined to be the largest integer r for which there exists a strictly increasing chain of irreducible closed subsets

$$X_0 \subsetneq X_1 \subsetneq \cdots \subsetneq X_r = X.$$

For a general (not necessarily irreducible) affine variety X, its dimension is defined as the maximum of the dimensions of its irreducible components.

Definition 1.4 (Degree)**.** Let $X \subseteq \mathbb{A}^n_k$ be an affine variety over an algebraically closed field k. Embed X into projective space $\mathbb{P}^n_k$ via the standard embedding, and intersect its projective closure with a general linear subspace L of complementary dimension (i.e., $\dim L + \dim X = n$). The *degree* $\deg X$ is defined to be the number of intersection points $X \cap L$ counted with multiplicity. Equivalently, $\deg X$ is the sum of the degrees of the irreducible components of X.

Intuitively, the dimension measures how many independent parameters are needed to describe a generic point of X, while the degree controls how X behaves under intersections with hypersurfaces.

1.3 Hilbert's Nullstellensatz and Effective Bounds

We now recall the classical Nullstellensatz and its effective versions which play a central role in the complexity analysis.

Theorem 1.1 (Weak Nullstellensatz)**.** Let k be an algebraically closed field and $I \subseteq k[x_1, \ldots, x_n]$ an ideal. Then

$$V(I) = \emptyset \quad \iff \quad I = (1).$$

Equivalently, if a set of polynomials $\{f_1, \ldots, f_m\} \subseteq k[x_1, \ldots, x_n]$ has no common zero in k^n, then there exist polynomials $g_1, \ldots, g_m \in k[x_1, \ldots, x_n]$ such that

$$f_1 g_1 + f_2 g_2 + \cdots + f_m g_m = 1.$$

Given polynomials $f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$, the *Hilbert Nullstellensatz problem (HN)* is:

- **Question:** Do the equations $f_1 = \cdots = f_m = 0$ have a common solution in $\mathbb{C}^n$?

Equivalently, does there exist $g_1, \dots, g_m$ with

$$f_1 g_1 + \cdots + f_m g_m = 1?$$

Thus, the computational task is to either find such a representation or prove that no solution exists.

We now recall results from Jelonek (2005) which provide explicit bounds on the degrees of the g_i in the Nullstellensatz identity.

Theorem 1.2 (Generalized Perron Theorem)**.** Let L be a field and let $Q_1, \dots, Q_{n+1} \in L[x_1, \dots, x_m]$ be non-constant polynomials with $\deg Q_j = d_j$. Let $X \subseteq L^m$ be an affine variety of dimension n and degree D. Assume the map

$$Q = (Q_1, \dots, Q_{n+1}) : X \longrightarrow L^{n+1}$$

is generically finite onto its image. Then there exists a non-zero polynomial

$$W(T_1, \dots, T_{n+1}) \in L[T_1, \dots, T_{n+1}]$$

such that

(a) $W(Q_1, \dots, Q_{n+1}) = 0$ identically on X, and

(b) writing $T_j^{d_j}$ for the monomial substitution, one has the degree estimate

$$\deg\big(W(T_1^{d_1}, T_2^{d_2}, \dots, T_{n+1}^{d_{n+1}})\big) \;\le\; D \prod_{j=1}^{n+1} d_j.$$

Proof. Replace the map Q by the auxiliary variety

$$\widetilde{X} = \{(x, w) \in X \times L^{n+1} : \; Q_j(x) = w_j^{d_j} + w_j \;\; (j = 1, \dots, n+1)\},$$

and consider the projection to the w-coordinates. By Corollary 3.2 and Bezout-type degree estimates one bounds the degree of the image hypersurface; an irreducible defining polynomial of that hypersurface provides the required W, and the degree bound follows from the Bézout estimate $\deg \widetilde{X} \le D \prod_j d_j$. □

To state the effective Nullstellensatz one convenient notation is:

$$N(d_1, \dots, d_k;\, n) \;=\; \begin{cases} \displaystyle\prod_{i=1}^{k} d_i & \text{if } 1 \le k \le n, \\ \Big(\displaystyle\prod_{i=1}^{n-1} d_i \Big)\, d_k & \text{if } k > n > 1, \\ d_1 & \text{if } n = 1. \end{cases}$$

Theorem 1.3 (Effective Nullstellensatz)**.** Let K be an algebraically closed field and let $X \subset K^m$ be an affine variety of dimension n and degree D. Let $f_1, \dots, f_k \in K[X]$ be non-constant polynomials with $\deg f_i = d_i$, arranged so that $d_1 \ge d_2 \ge \cdots \ge d_k$. Assume the f_i have no common zero on X. Then there exist $g_1, \dots, g_k \in K[X]$ such that

$$\sum_{i=1}^{k} f_i g_i = 1$$

on X, and the products $f_i g_i$ satisfy the degree bounds

$$\deg(f_i g_i) \;\leq\; \begin{cases} D\,N(d_1,\ldots,d_k;\,n), & \text{if } k \leq n, \\ 2D\,N(d_1,\ldots,d_k;\,n) - 1, & \text{if } k > n. \end{cases}$$

Since the existence of g_i with explicit degree and coefficient bounds is guaranteed, the problem of deciding whether $1 \in (f_1,\ldots,f_m)$ reduces to solving a linear system over $\mathbb{Z}$ whose size is bounded by $\text{poly}(d^n, \log H)$. Such linear algebra computations can be performed in PSPACE.

Corollary 1.4. Hilbert's Nullstellensatz problem (HN) lies in PSPACE.

2 Over characteristic zero fields

Theorem 2.1. There exist constants c_1, c_2, $c_3 \in \mathbb{N}$ such that if $A = d^{c_1 n} s(\lceil \log s \rceil + L)$ and $x_0 \geq L^{c_2} 2^{(n \log \sigma)^{c_3}}$ the following two properties hold:

- If the system of equations has no solution in $\mathbb{C}$, then $\pi_S(x_0) \leq A$.
- If the system of equations is satisfiable in $\mathbb{C}$, then $\pi_S(x_0) \geq B = 8A(\log A + 3)$.

This first theorem establishes a large gap in the number of prime moduli p for which a system of equations S is satisfiable. If the system has no solution over $\mathbb{C}$, there are very few such primes; if it does have a solution, there are many. This gap is the key to placing the problem within the polynomial hierarchy.

Theorem 2.2. Hilbert's Nullstellensatz for fields of characteristic zero, is in AM under GRH.

2.1 The Arthur–Merlin Class (AM)

The Arthur–Merlin (AM) complexity class is a class of decision problems that can be solved by a specific type of two-participant protocol called an *interactive proof system*. The two participants are:

- **Arthur**: A verifier with the power of a probabilistic polynomial-time Turing machine (BPP). He is skeptical and must be convinced.
- **Merlin**: An all-powerful prover with infinite computational resources. He wants to convince Arthur that a given input string belongs to a certain language.

The protocol for a language L proceeds as follows:

1. Arthur receives an input string x.
2. Arthur generates a random string y of length polynomial in the size of x and sends it to Merlin. This is Arthur's "challenge".
3. Merlin, seeing both x and y, computes a response (a "proof") z and sends it back to Arthur.
4. Arthur runs a deterministic polynomial-time algorithm on the triplet (x, y, z) and decides whether to **accept** or **reject**.

Definition 2.1. A language L is in **AM** if there exists a probabilistic polynomial-time verifier (Arthur) such that:

1. **Completeness**: If $x \in L$, then there exists a proof from Merlin that makes Arthur accept with high probability.
$$\forall x \in L, \quad \Pr_y[\exists z : \text{Arthur accepts}(x, y, z)] \geq \frac{2}{3}$$
2. **Soundness**: If $x \notin L$, then for any proof Merlin sends, Arthur will reject with high probability.
$$\forall x \notin L, \quad \Pr_y[\forall z : \text{Arthur rejects}(x, y, z)] \geq \frac{2}{3}$$

The probabilities 2/3 and 1/3 can be amplified to be arbitrarily close to 1 and 0 by repeating the protocol.

2.2 Overview of the Proof Strategy

The central thesis for establishing Hilbert's Nullstellensatz (HN) in the polynomial hierarchy is to create a reduction from a problem over the complex numbers $\mathbb{C}$ to a counting problem over finite fields $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$. This is achieved by demonstrating a large, computationally recognizable gap in the number of prime moduli p for which a system S is satisfiable.

Let S be the system $f_1(x) = 0, \ldots, f_s(x) = 0$ with $f_i \in \mathbb{Z}[X_1, \ldots, X_n]$. Let R_S be the set of primes p such that S has a solution in $\mathbb{F}_p$, and let $\pi_S(x) = |R_S \cap \{1, \ldots, x\}|$.

Theorem 1 in the paper establishes this gap:

- **Unsatisfiable Case:** If S has no solution in $\mathbb{C}$, then $\pi_S(x_0)$ is *small* (bounded by a polynomial A in the input size).
- **Satisfiable Case:** If S has a solution in $\mathbb{C}$, then $\pi_S(x_0)$ is *large* (bounded below by B, which is super-polynomially larger than A).

The algorithm then consists of counting $\pi_S(x_0)$ for a suitably large x_0 and checking whether this count falls below or above a threshold between A and B. This section develops the (conditional) bounds required to prove this gap exists.

2.3 The Unsatisfiable Case: An Upper Bound on Prime Solutions

The bound for the unsatisfiable case stems from the fact that if 1 is in the ideal $\langle f_1, \ldots, f_s \rangle$ over $\mathbb{C}$, it must also be in the ideal over $\mathbb{Z}$ up to a multiplicative integer constant.

Theorem 2.3 (Bound for Unsatisfiable Systems)**.** Let S be a system with degrees at most d and coefficient bit size at most L. If the system S has no solution in $\mathbb{C}$, then R_S is finite and
$$|R_S| \leq d^{O(n)} s(\log s + L).$$

Proof. By the Effective Nullstellensatz (as cited from Kollár, 1988; Krick and Pardo, 1994), if S is unsatisfiable, there exists a non-zero integer $a \in \mathbb{Z}$ and polynomials $g_1, \ldots, g_s \in \mathbb{Z}[X_1, \ldots, X_n]$ such that
$$a = \sum_{i=1}^{s} g_i f_i.$$
The size of the integer a is bounded by $\log |a| \leq d^{O(n)} s(\log s + L)$.

Now, consider this identity modulo a prime p. If S *has* a solution $\mathbf{x} \in \mathbb{F}_p^n$, then $f_i(\mathbf{x}) \equiv 0 \pmod{p}$ for all i. Substituting this into the identity gives:

$$a \equiv \sum_{i=1}^{s} g_i(\mathbf{x}) f_i(\mathbf{x}) \equiv 0 \pmod{p}.$$

This implies that S can only have a solution in $\mathbb{F}_p$ if p divides a. Since $a \neq 0$, the number of such prime divisors is at most $\log_2 |a|$. The theorem follows by applying the bound on $\log |a|$. □

2.4 The Satisfiable Case: A Lower Bound on Prime Solutions

The argument for the satisfiable case is more involved. The strategy is to (1) find a "low-complexity" algebraic solution, (2) represent this n-variate solution using a single univariate primitive element r, and (3) use (conditional) number-theoretic results to show that the minimal polynomial $R(X)$ of this primitive element r must have roots modulo "many" primes p. Finally, (4) show that a root of $R(X)$ in $\mathbb{F}_p$ "lifts" to a solution of the full system S in $\mathbb{F}_p$.

Theorem 2.4 (Low-Complexity Algebraic Solution)**.** There are absolute constants c_1, c_2 such that if S has a solution over $\mathbb{C}$, there exists a solution $\mathbf{x} = (x_1, \ldots, x_n)$ where each x_i is an algebraic number which is a root of a polynomial $A_i \in \mathbb{Z}[X]$ of degree at most $2^{(n \log \sigma)^{c_1}}$ with coefficients of bit size at most $L \cdot 2^{(n \log \sigma)^{c_2}}$. (Here σ is the total degree).

Proof. This result follows from effective quantifier elimination over $\mathbb{C}$ (e.g., Fichtas et al., 1990). The solution set $S \subset \mathbb{C}^n$ is projected onto each coordinate axis i to get $S_i \subset \mathbb{C}$. This projection S_i is definable by a formula derived from S. By quantifier elimination, S_i can be defined by a quantifier-free formula involving polynomials P_{ij}. If S_i is finite, the x_i must be roots of these P_{ij}. If S_i is infinite (i.e., $\mathbb{C}$ minus a finite set), a "small" integer solution $\alpha \in \mathbb{Z}$ can be chosen for x_i, and the process is repeated inductively on the remaining $n-1$ variables. The bounds on the degrees and coefficient sizes from the quantifier-elimination theorem are carried through this induction. □

This provides the A_i polynomials which serve as the input for constructing a primitive element.

Theorem 2.5 (Low-Complexity Primitive Element)**.** Let $x_1, \ldots, x_n$ be algebraic numbers which are roots of polynomials $A_i \in \mathbb{Z}[X]$ of degree at most d with coefficients of size at most L. There exists a primitive element r for $x_1, \ldots, x_n$ (i.e., $x_i = Q_i(r)/a_i$) which is a root of an irreducible polynomial $B \in \mathbb{Z}[X]$ of degree at most d^n. The coefficients of B are of size at most $L \cdot d^{n^{O(1)}}$. Moreover, each x_i can be represented as $x_i = Q_i(r)/a$ for a common denominator $a \in \mathbb{Z}$, where $Q_i \in \mathbb{Z}[X]$ and $\log |a| = L \cdot d^{n^{O(1)}}$.

Proof. The primitive element is constructed from the given A_i in three steps.

1. Make the polynomials A_i square-free by computing $P_i = A_i / \gcd(A_i, A_i')$. Standard subresultant gcd algorithms (as cited in Mignotte, 1982) show that the coefficients of P_i remain polynomially bounded in the size of A_i. The degrees are unchanged or reduced.

2. Apply the inductive primitive element construction (Lemma 1 in the paper, which itself is built by repeated application of the 2-variable case) to the roots of $P_1, \ldots, P_n$. This construction yields a polynomial $R \in \mathbb{Z}[X]$ of degree at most $\prod \deg(P_i) \leq d^n$ and a representation $x_i = Q_i(r)/a$ where r is a root of R. The bounds on the norms of R and the size of a and coefficients of Q_i are shown in the paper's Lemma 1 to be $N(R) \leq N^{d^{O(n^2)}}$

and $\log |a| = N^{d^{O(n^2)}}$, where N is the norm of the P_i. In our context (using bit size L instead of norm N), this translates to the $L \cdot d^{n^{O(1)}}$ bounds.

3. The primitive element r is a root of R. The polynomial B is the minimal polynomial of r, so it must be an irreducible factor of R.

The bounds for polynomial factorization (e.g., Mignotte, 1982) show that the coefficients of B are polynomially bounded by the size of R. Therefore, the dominant cost comes from step 2, and the bounds stated in the theorem hold for the minimal polynomial B and the representations. □

The following lemma connects the roots of $R(X)$ modulo p to the solutions of S modulo p.

Lemma 2.6 (Modular Solution Lifting)**.** Let $\mathbf{x} = (x_1, \ldots, x_n)$ be a vector of algebraic numbers solution of S. Let r be a primitive element for $x_1, \ldots, x_n$: there exist polynomials $Q_1, \ldots, Q_n \in \mathbb{Z}[X]$ and $a \in \mathbb{Z}, a \neq 0$ such that $x_i = Q_i(r)/a$. Let $R \in \mathbb{Z}[X]$ be the irreducible minimal polynomial of r. If R has a root in $\mathbb{F}_p$ and $a \not\equiv 0 \pmod{p}$, then S is satisfiable in $\mathbb{F}_p$.

Proof. For $j \in \{1, \ldots, s\}$, let $d_j = \deg(f_j)$. Define a new polynomial $g_j \in \mathbb{Z}[X]$ by clearing the denominator a:

$$g_j(X) = a^{d_j} f_j(Q_1(X)/a, \ldots, Q_n(X)/a). \tag{1}$$

Since f_j has integer coefficients, $g_j(X)$ is a polynomial with integer coefficients. We know r is a root of $g_j(X)$, because:

$$g_j(r) = a^{d_j} f_j(Q_1(r)/a, \ldots, Q_n(r)/a) = a^{d_j} f_j(x_1, \ldots, x_n) = a^{d_j} \cdot 0 = 0.$$

Since $g_j(r) = 0$ and $R(X)$ is the minimal polynomial of r (and is irreducible), $R(X)$ must divide $g_j(X)$ over the rationals, and by Gauss's lemma, over the integers. Thus, there exist polynomials $A_j \in \mathbb{Z}[X]$ such that

$$g_j(X) = R(X)A_j(X). \tag{2}$$

Now, consider these identities modulo p. If $a \not\equiv 0 \pmod{p}$, then a is invertible in $\mathbb{F}_p$, and both (1) and (2) hold in $\mathbb{F}_p[X]$. Let $x_0 \in \mathbb{F}_p$ be a root of R modulo p. From (2), this means $g_j(x_0) \equiv 0 \pmod{p}$ for all j. From (1), this implies:

$$a^{d_j} f_j(Q_1(x_0)/a, \ldots, Q_n(x_0)/a) \equiv 0 \pmod{p}.$$

Since $a \not\equiv 0 \pmod{p}$, we can multiply by $(a^{-1})^{d_j}$ to conclude:

$$f_j(Q_1(x_0)/a, \ldots, Q_n(x_0)/a) \equiv 0 \pmod{p}$$

for all $j = 1, \ldots, s$. Therefore, the vector $\mathbf{y} = (Q_1(x_0)/a, \ldots, Q_n(x_0)/a)$ is a solution of S in $\mathbb{F}_p$. □

The final step is to show that $R(X)$ has "many" roots mod p. This relies on the (conditional) effective Chebotarev density theorem.

Theorem 2.7 (Effective Chebotarev Bound)**.** Let $f \in \mathbb{Z}[X]$ be an irreducible polynomial of degree m with discriminant Δ. Let $W(p)$ be the number of roots of f in $\mathbb{F}_p$. Let $S(x) = \sum_{p \leq x, p \nmid \Delta} (1 - W(p))$. Assuming GRH,

$$|S(x)| = O(x^{1/2} \log(\Delta x^m)).$$

Proof Sketch. This result, cited from Adleman and Odlyzko (1983) and Weinberger (1984), is a consequence of the effective Chebotarev density theorem (Lagarias and Odlyzko, 1977) which is conditional on the Generalized Riemann Hypothesis (GRH). It states that the number of roots $W(p)$ is, on average, 1. The theorem provides an effective bound on the error term of this average. □

Corollary 2.8 (Lower Bound on Roots)**.** Assuming GRH, for an irreducible $f \in \mathbb{Z}[X]$ of degree m, there exists an absolute constant c such that

$$\pi_f(x) \geq \frac{1}{m}\left[\pi(x) - \log \Delta - c \cdot x^{1/2} \log(\Delta x^m)\right]$$

where $\pi_f(x)$ is the number of primes $\leq x$ for which f has at least one root in $\mathbb{F}_p$.

Proof. From Theorem 2.7, we have $\sum'_{p \leq x}(1 - W(p)) \leq |S(x)|$. This gives $\sum'_{p \leq x} W(p) \geq \sum'_{p \leq x} 1 - |S(x)|$. The number of primes $p \leq x$ not dividing Δ is $\pi(x)$ minus the number of prime factors of Δ, which is at most $\log \Delta$. So, $\sum'_{p \leq x} 1 \geq \pi(x) - \log \Delta$. Plugging this in and using the bound for $|S(x)|$:

$$\sum_{p \leq x}{}' W(p) \geq \pi(x) - \log \Delta - c \cdot x^{1/2} \log(\Delta x^m).$$

The total number of roots $W(p)$ is at most $m \cdot r_f(p)$, where $r_f(p) = 1$ if f has a root mod p and 0 otherwise. Thus, $\sum W(p) \leq m \sum r_f(p) = m \cdot \pi_f(x)$. Dividing by m gives the result. □

Theorem 2.9 (Bound for Satisfiable Systems)**.** If S is satisfiable, then (assuming GRH) for $x \geq L^{c_2} 2^{(n \log \sigma)^{c_3}}$ (from Theorem 1):

$$\pi_S(x) \geq \frac{\pi(x)}{L \cdot 2^{(n \log \sigma)^{O(1)}}} - L \cdot 2^{(n \log \sigma)^{O(1)}} x^{1/2}.$$

Proof. 1. Apply Theorem 2.4 to get a low-complexity solution $\mathbf{x}$.

2. Apply Theorem 2.5 to get its minimal polynomial $R(X)$ and denominator a. Let $\deg(R) = m$ and its discriminant be Δ. The bounds are:

   - $m \leq D^n = (2^{(n \log \sigma)^{c_1}})^n = 2^{n(n \log \sigma)^{c_1}}$;
   - $\log \Delta$ is polynomial in m and the coefficient size of R;
   - $\log |a| = L \cdot d^{n^{O(1)}}$.

   All these quantities are bounded by $L \cdot 2^{(n \log \sigma)^{O(1)}}$.

3. Apply Lemma 2.6: $\pi_S(x) \geq \pi_R(x) - (\text{primes dividing } a)$.

4. The number of primes dividing a is at most $\log |a|$.

5. Apply Corollary 2.8 to $R(X)$:

$$\pi_R(x) \geq \frac{1}{m}\left[\pi(x) - \log \Delta - c \cdot x^{1/2} \log(\Delta x^m)\right].$$

6. Combine these:

$$\pi_S(x) \geq \frac{\pi(x)}{m} - \frac{\log \Delta}{m} - \frac{c x^{1/2} \log(\Delta x^m)}{m} - \log |a|.$$

7. For x large enough (as specified in Theorem 1), the $\pi(x)/m$ term dominates. The error terms involving $\log \Delta$, $\log |a|$, and $x^{1/2}$ are all absorbed into the negative term. Plugging in the bounds for m, $\log \Delta$, and $\log |a|$ gives the final expression.

□

Remark 2.1. The full argument for Theorem 1 in the paper combines these two bounds.

1. **If S is unsatisfiable in $\mathbb{C}$:** By Theorem 2.3, $\pi_S(x) \leq A$ for $A = d^{O(n)} s(\log s + L)$. This bound holds for any x.

2. **If S is satisfiable in $\mathbb{C}$:** By Theorem 2.9, for a sufficiently large x_0, $\pi_S(x_0) \geq B$.

The crucial point is that B (which is roughly $\pi(x_0)/D^n$) is super-polynomially larger than A, creating the decidable gap.

3 Over Finite Fields

We now turn to Nullstellensatz-type questions over finite fields and analyse the complexity in the positive-dimensional and zero-dimensional cases separately.

3.1 Problem Setup over $\mathbb{F}_q$

Let $\mathbb{F}_q$ be a finite field with q elements and let $\overline{\mathbb{F}}_q$ denote its algebraic closure. Consider a system of polynomial equations

$$S = \{f_1 = 0, \ldots, f_m = 0\}, \qquad f_i \in \mathbb{F}_q[x_1, \ldots, x_n],$$

with degrees $\deg(f_i) \leq d$. We define the affine variety

$$V(S) \;=\; \{x \in \overline{\mathbb{F}}_q^{\,n} : f_i(x) = 0 \text{ for all } i\}.$$

We are especially interested in the case where $V(S)$ has an irreducible component of positive dimension.

Definition 3.1 (Positive- and zero-dimensional cases)**.** Let $S \subseteq \mathbb{F}_q[x_1, \ldots, x_n]$ be as above and let $V(S)$ be its zero set in $\overline{\mathbb{F}}_q^{\,n}$.

- We say that the *positive-dimensional case* occurs if $V(S)$ has an irreducible component X of dimension $r > 0$ defined over $\mathbb{F}_q$.

- We say that we are in the *zero-dimensional case* if every irreducible component of $V(S)$ has dimension 0, i.e., $V(S)$ is a finite set (over $\overline{\mathbb{F}}_q$).

A central question is whether one can bound the "size" of a solution in terms of the input parameters. Over finite fields, size is conveniently measured by the degree of the finite extension $\mathbb{F}_{q^t}$ that contains the coordinates of a solution.

3.2 Deligne's Theorem and Point Counting

Our analysis in the positive-dimensional case ultimately rests on Deligne's proof of the Weil conjectures, which we state here in a very special form.

Theorem 3.1 (Deligne's Riemann Hypothesis, special case)**.** Let X be a smooth projective variety of dimension r over $\mathbb{F}_q$. For each $k \geq 0$, the eigenvalues of the geometric Frobenius on the ℓ-adic cohomology $H_c^k(X, \mathbb{Q}_\ell)$ have absolute value $q^{k/2}$.

This deep theorem admits standard consequences for point counting on varieties over finite fields. In particular we will use the following quantitative estimate.

Corollary 3.2 (Point-counting estimate)**.** Let X be a geometrically irreducible variety of dimension r defined over $\mathbb{F}_q$, with degree bounded by D. Then there exists a constant $C_X > 0$, depending polynomially on D and the ambient dimension, such that for all integers $t \geq 1$ one has

$$\left|\#X(\mathbb{F}_{q^t}) - q^{rt}\right| \;\leq\; C_X\, q^{(r-1/2)t}.$$

In particular, if t is large enough so that $q^{t/2} > C_X$, then $X(\mathbb{F}_{q^t}) \neq \emptyset$.

Proof. **Step 1: reduction to the projective case.** Choose an embedding $X \hookrightarrow \mathbb{P}^N_{\mathbb{F}_q}$ for some N, and let $\overline{X} \subset \mathbb{P}^N_{\mathbb{F}_q}$ be the projective closure. Then $\overline{X}$ is geometrically irreducible of dimension r and $\deg \overline{X} = \deg X \leq D$. Let $Z := \overline{X} \setminus X$ be the boundary; then $\dim Z \leq r-1$.

The desired estimate for X will follow once we know it for $\overline{X}$ and for each irreducible component of Z (with r replaced by $\dim Z \leq r-1$), because

$$\#X(\mathbb{F}_{q^t}) = \#\overline{X}(\mathbb{F}_{q^t}) - \#Z(\mathbb{F}_{q^t})$$

and hence

$$\left|\#X(\mathbb{F}_{q^t}) - q^{rt}\right| \leq \left|\#\overline{X}(\mathbb{F}_{q^t}) - q^{rt}\right| + \#Z(\mathbb{F}_{q^t}).$$

The contribution of Z can be absorbed into the same type of bound because every irreducible component of Z has dimension at most $r-1$, so its point count is $\ll q^{(r-1)t}$, which is $\ll q^{(r-1/2)t}$ after enlarging the constant. Thus it is enough to prove the corollary when X itself is projective.

So assume from now on that $X \subset \mathbb{P}^N_{\mathbb{F}_q}$ is projective, geometrically irreducible of dimension r and degree $d \leq D$.

Step 2: Lang–Weil inequality in the form of Ghorpade–Lachaud. Let k be a finite field and write $\pi_r(k) := \#\mathbb{P}^r(k) = 1 + |k| + \cdots + |k|^r$. Ghorpade and Lachaud prove the following version of the classical Lang–Weil inequality; see (Ghorpade and Lachaud, 2002, Thm. 11.1). It is stated there with $k = \mathbb{F}_q$:

> *Let X be a projective algebraic subvariety of $\mathbb{P}^N_k$ of dimension r and degree d. Then*
>
> $$\left|\#X(k) - \pi_r(k)\right| \;\leq\; (d-1)(d-2)\,|k|^{r-1/2} + C_+(X)\,|k|^{r-1},$$
>
> *where $C_+(X) \geq 0$ depends only on X and admits an explicit bound in terms of the ambient dimension N and the degrees of defining equations of X, cf. the inequality following* (Ghorpade and Lachaud, 2002, Thm. 11.1)

Now fix $t \geq 1$. Consider X as a variety over the larger finite field $\mathbb{F}_{q^t}$. The ℓ–adic cohomology groups $H^i_c(X_{\overline{\mathbb{F}_q}}, \mathbb{Q}_\ell)$ and the Frobenius action on them do not change when we extend scalars from $\mathbb{F}_q$ to $\mathbb{F}_{q^t}$; only the base field (and thus the cardinality $|k|$) changes. Thus the same theorem applied with $k = \mathbb{F}_{q^t}$ gives

$$\left|\#X(\mathbb{F}_{q^t}) - \pi_r(\mathbb{F}_{q^t})\right| \;\leq\; (d-1)(d-2)\,q^{(r-1/2)t} + C_+(X)\,q^{(r-1)t} \tag{3}$$

for all $t \geq 1$, with the *same* constant $C_+(X)$, independent of t.

Step 3: replacing $\pi_r(\mathbb{F}_{q^t})$ by q^{rt}. We have

$$\pi_r(\mathbb{F}_{q^t}) = 1 + q^t + \cdots + q^{rt} = q^{rt} + \left(1 + q^t + \cdots + q^{(r-1)t}\right).$$

Therefore

$$\left|\pi_r(\mathbb{F}_{q^t}) - q^{rt}\right| = 1 + q^t + \cdots + q^{(r-1)t} \le r\, q^{(r-1)t}.$$

Combining this with (3) and the triangle inequality gives

$$\begin{aligned}\left|\#X(\mathbb{F}_{q^t}) - q^{rt}\right| &\le \left|\#X(\mathbb{F}_{q^t}) - \pi_r(\mathbb{F}_{q^t})\right| + \left|\pi_r(\mathbb{F}_{q^t}) - q^{rt}\right| \\ &\le (d-1)(d-2)\, q^{(r-1/2)t} + C_+(X)\, q^{(r-1)t} + r\, q^{(r-1)t}.\end{aligned}$$

Since $t \ge 1$ and $q \ge 2$, we have $q^{(r-1)t} \le q^{(r-1/2)t}$, so we can absorb the $q^{(r-1)t}$ terms into the $q^{(r-1/2)t}$ term by enlarging the constant. Precisely, set

$$C_X := (d-1)(d-2) + C_+(X) + r.$$

Then for all $t \ge 1$,

$$\left|\#X(\mathbb{F}_{q^t}) - q^{rt}\right| \;\le\; C_X\, q^{(r-1/2)t}.$$

Step 4: dependence of C_X on D and N. By construction, $d = \deg X \le D$, so $(d-1)(d-2) \le D^2$ is polynomial in D. The constant $r = \dim X$ is at most N.

The remaining ingredient is $C_+(X)$. In the proof of (Ghorpade and Lachaud, 2002, Prop. 11.2), $C_+(X)$ is defined in terms of the ℓ–adic Betti numbers $b_{i,\ell}(X)$ and $H_c^{2r-1}(X_{\overline{\mathbb{F}_q}}, \mathbb{Q}_\ell)$:

$$C_+(X) = \dim H_+^{2r-1}\big(X_{\overline{\mathbb{F}_q}}, \mathbb{Q}_\ell\big) + \sum_{i=0}^{2r-2} b_{i,\ell}(X) + \varepsilon_r,$$

where H_+^{2r-1} denotes the subspace corresponding to Frobenius eigenvalues of weight $< 2r-1$ and $\varepsilon_r \in \{0, 1\}$ is explicit. Hence $C_+(X)$ is bounded above by a constant multiple of the *sum* of ℓ–adic Betti numbers of X.

Now $X \subset \mathbb{P}^N$ can be defined by finitely many homogeneous polynomials of degree at most some d_0 with d_0 bounded in terms of $\deg X \le D$ and N (this is a standard fact from projective geometry, and one can make it quantitative using bounds of Bombieri, Heintz, Jelonek, etc.). If r equations of degree $\le d_0$ cut out X, Katz's theorem on sums of Betti numbers gives an explicit bound

$$\sum_i \dim H^i\big(X_{\overline{\mathbb{F}_q}}, \mathbb{Q}_\ell\big) \;\le\; B(N, r, d_0),$$

where $B(N, r, d_0)$ is an explicit function polynomial in d_0 (for fixed N and r); see Katz, *Sums of Betti numbers in arbitrary characteristic*, Theorem 3 and the following corollary. Since d_0 itself can be taken $\ll D$ (with constants depending only on N), this shows that $C_+(X)$, hence C_X, is bounded by a polynomial in D with coefficients depending on N.

□

3.3 Positive-Dimensional Case: Existence of a Small Root

We now show that in the positive-dimensional case there is always a solution whose coordinates lie in a finite extension $\mathbb{F}_{q^t}$ of *polynomially bounded* degree.

Lemma 3.3 (Degree bound for a component)**.** Let $S = \{f_1, \ldots, f_m\} \subseteq \mathbb{F}_q[x_1, \ldots, x_n]$ with $\deg(f_i) \le d$, and let X be an irreducible component of $V(S)$ of dimension r. Then there is a constant C (independent of q) such that

$$\deg X \;\le\; d^{Cn}.$$

Proof. This is a standard consequence of Bézout-type degree bounds and elimination theory. For example, one may take successive intersections of the hypersurfaces $\{f_i = 0\}$ and use the fact that the degree of an intersection is at most the product of the degrees of the hypersurfaces, together with primary decomposition in the coordinate ring. A detailed argument may be found, for instance, in Jelonek (2005). The upshot is that the degree of any irreducible component of $V(S)$ is bounded by a fixed power of d depending only on n. $\square$

We can now state and prove the small-root theorem.

Theorem 3.4 (Small root for a positive-dimensional component)**.** Let $S = \{f_1, \dots, f_m\} \subseteq \mathbb{F}_q[x_1, \dots, x_n]$ with $\deg f_i \le d$, and suppose that $V(S)$ has an irreducible component X of dimension $r > 0$ defined over $\mathbb{F}_q$. Then there exists an integer

$$1 \le t \le \mathrm{poly}(n, \log d, \log q)$$

such that S has a solution in $\mathbb{F}_{q^t}^n$.

Proof. By Lemma 3.3 there exists a constant C_0 such that $\deg X \le d^{C_0 n}$. Let $D = \deg X$. Applying Corollary 3.2 to X, we obtain a constant C_X bounded polynomially in D and n such that

$$\bigl|\#X(\mathbb{F}_{q^t}) - q^{rt}\bigr| \le C_X q^{(r-1/2)t} \qquad \text{for all } t \ge 1.$$

In particular,

$$\#X(\mathbb{F}_{q^t}) \ge q^{rt} - C_X q^{(r-1/2)t} = q^{(r-1/2)t}(q^{t/2} - C_X).$$

Thus $X(\mathbb{F}_{q^t})$ is non-empty as soon as $q^{t/2} > C_X$ (since we have $r - \tfrac{1}{2} > 0$). Let t_0 be the least integer such that $q^{t_0/2} > C_X$. Then $X(\mathbb{F}_{q^{t_0}}) \neq \emptyset$, so there exists a point $x \in X(\mathbb{F}_{q^{t_0}})$ with coordinates in $\mathbb{F}_{q^{t_0}}$. Since $X \subseteq V(S)$, this point is in fact a solution of S in $\mathbb{F}_{q^{t_0}}^n$.

It remains to bound t_0 in terms of the input parameters. By polynomial dependence of C_X on D and n, there exists a constant C_1 such that

$$C_X \le d^{C_1 n}.$$

Taking binary logarithms we get

$$\log_2 C_X \le C_1 n \log_2 d.$$

Since $q \ge 2$, we have $\log_q C_X = (\log_2 C_X)/(\log_2 q)$, and hence

$$\log_q C_X \le \frac{C_1 n \log_2 d}{\log_2 q} \;=\; \mathrm{poly}(n, \log d, \log q).$$

By the definition of t_0 we may take, for example,

$$t_0 = 2\bigl(\lfloor \log_q C_X \rfloor + 1\bigr),$$

which clearly satisfies $q^{t_0/2} > C_X$. Consequently t_0 is bounded by a polynomial in n, $\log d$ and $\log q$, as required. Setting $t = t_0$ completes the proof. $\square$

3.4 AM Protocol in the Positive-Dimensional Case

We now explain how Theorem 3.4 yields a proof that, under the positive-dimensional promise, the associated Nullstellensatz problem over $\mathbb{F}_q$ lies in the class AM (in fact, already in NP).

Theorem 3.5 (Positive-dimensional Nullstellensatz over $\mathbb{F}_q$)**.** Consider the decision problem:

> *Input:* A system $S = \{f_1, \ldots, f_m\} \subseteq \mathbb{F}_q[x_1, \ldots, x_n]$ with $\deg f_i \leq d$, together with the promise that $V(S)$ has an irreducible component X of dimension $r > 0$ defined over $\mathbb{F}_q$.
>
> *Question:* Does S have a common root in $\overline{\mathbb{F}}_q^n$?

Then this promise problem lies in AM (indeed, in NP).

Proof. By Theorem 3.4, if S is satisfiable then there exists a solution $x = (x_1, \ldots, x_n)$ with coordinates in $\mathbb{F}_{q^t}$ for some $t \leq T$, where T is bounded by a polynomial in n, $\log d$ and $\log q$. We fix a canonical representation of the extension field $\mathbb{F}_{q^t}$ as $\mathbb{F}_q[X]/(P(X))$, where P is a monic irreducible polynomial of degree t over $\mathbb{F}_q$. An element of $\mathbb{F}_{q^t}$ is then represented by a residue class of a polynomial of degree $< t$, which in turn can be encoded by t coefficients in $\mathbb{F}_q$, i.e., $t \log q$ bits.

A *certificate* for the satisfiability of S consists of the following data:

- a description of the extension field $\mathbb{F}_{q^t}$, given by an irreducible polynomial $P \in \mathbb{F}_q[X]$ of degree t;
- the coordinates $x_1, \ldots, x_n \in \mathbb{F}_{q^t}$, each written in the chosen basis of $\mathbb{F}_{q^t}$ over $\mathbb{F}_q$.

The total length of this certificate is $O(t \log q \cdot n)$ bits, which is polynomial in the input size by the bound on t.

Given such a certificate, a polynomial-time verifier (Arthur) proceeds as follows:

1. First check that P is irreducible over $\mathbb{F}_q$; this can be done in randomized or deterministic polynomial time using standard algorithms for irreducibility testing of polynomials over finite fields.
2. Using the representation $\mathbb{F}_{q^t} = \mathbb{F}_q[X]/(P)$, implement field addition and multiplication in time polynomial in t and $\log q$.
3. For each i, evaluate $f_i(x)$ in $\mathbb{F}_{q^t}$ by repeated field operations, and check whether the result is zero.

If all evaluations vanish, the verifier accepts; otherwise, it rejects.

If S is satisfiable, then by Theorem 3.4 there exists a solution in some $\mathbb{F}_{q^t}$ with $t \leq T$, so Merlin can send a correct certificate which causes Arthur to accept. Conversely, if S is unsatisfiable, then there is no tuple x in any extension field which satisfies all equations, so no purported certificate can make Arthur accept. Thus the problem lies in NP and hence in AM. □

3.5 Zero-Dimensional Case and Large Roots

We now explain why the simple strategy of using a root as a certificate can fail in the zero-dimensional case: even when $V(S)$ is finite, the solutions may live in very large field extensions.

Definition 3.2 (Zero-dimensional systems)**.** A system $S = \{f_1 = 0, \ldots, f_m = 0\} \subseteq K[x_1, \ldots, x_n]$ over a field K is said to be *zero-dimensional* if its zero set $V(S)$ in an algebraic closure $\overline{K}$ is finite. Equivalently, the coordinate ring $K[x_1, \ldots, x_n]/\langle f_1, \ldots, f_m \rangle$ is a finite-dimensional K-vector space.

The following example (adapted from Koiran's work) shows that in characteristic 3 there exist zero-dimensional systems whose solutions require field extensions of exponentially large degree.

Proposition 3.6 (Exponential extension degree in the zero-dimensional case)**.** Let K be a field of characteristic 3 and, for $n \geq 1$, consider the system S_n in variables $x_1, \ldots, x_n$ given by

$$x_1 = 2, \qquad x_{k+1}^2 = x_k \quad (1 \leq k < n).$$

Then S_n is zero-dimensional and for any solution $(x_1, \ldots, x_n) \in \overline{K}^n$ one has

$$[K(x_n) : K] \geq 2^{n-1}.$$

Proof. Each equation $x_{k+1}^2 = x_k$ is quadratic in x_{k+1}, so for fixed x_k there are at most two possibilities for x_{k+1}. Hence S_n has only finitely many solutions and is zero-dimensional.

Fix a solution and set $K_1 = K$ and $K_{k+1} = K_k(x_{k+1})$ for $1 \leq k < n$. For each k, x_{k+1} satisfies the monic quadratic

$$h_k(X) := X^2 - x_k \in K_k[X].$$

Thus $[K_{k+1} : K_k]$ is either 1 or 2. One checks inductively that h_k is irreducible over K_k for every k (equivalently, x_k is not a square in K_k), so $[K_{k+1} : K_k] = 2$ at each step. Consequently

$$[K(x_n) : K] \geq [K_n : K] = \prod_{k=1}^{n-1} [K_{k+1} : K_k] = 2^{n-1},$$

as claimed. □

Remark 3.1. Proposition 3.6 shows that, unlike in the positive-dimensional case treated in Theorem 3.4, one cannot hope for a general polynomial bound on the extension degree needed to realize a solution of a zero-dimensional system over a field of positive characteristic. From the viewpoint of complexity theory, this explains why the simple AM protocol of Theorem 3.5 does not extend verbatim to the zero-dimensional case: a single root may require an exponentially long description.

4 Conclusion and Future Directions

In this report we investigated the computational complexity of Hilbert's Nullstellensatz over both characteristic zero fields and finite fields, combining tools from effective algebraic geometry and analytic number theory. Over characteristic zero, effective Nullstellensatz bounds (Theorem 1.3) yield an explicit representation of 1 in the ideal $\langle f_1, \ldots, f_m \rangle$ with controlled degrees and coefficient sizes. This leads to a PSPACE decision procedure for the Hilbert Nullstellensatz problem via linear algebra over the integers. Under the Generalized Riemann Hypothesis, Koiran's argument then upgrades this to an AM protocol by relating satisfiability over $\mathbb{C}$ to a gap in the number of prime moduli for which the system has a solution.

Over finite fields, Deligne's proof of the Weil conjectures provides strong point-counting estimates which, together with degree bounds on components, guarantee the existence of "small" roots in the positive-dimensional case (Theorem 3.4). This yields short certificates and places

the corresponding promise problem in NP (hence AM) (Theorem 3.5). In contrast, the zero-dimensional case admits systems whose solutions require exponentially large extension degrees (Proposition 3.6), showing that naive root-based certificates are inherently too large in general.

Several directions remain open for future work:

- The AM upper bound over characteristic zero currently depends on GRH via effective Chebotarev estimates. An unconditional analogue with weaker, but still meaningful, complexity bounds would be highly desirable. or special group actions).
- For zero-dimensional varieties, it is natural to seek alternative certificate notions that might place the corresponding decision problems in lower complexity classes.

References

Pierre Deligne. La conjecture de weil : I. *Publications Mathématiques de l'IHÉS*, 43:273–307, 1974. URL https://www.numdam.org/item/PMIHES_1974__43__273_0/.

Pierre Deligne. La conjecture de Weil : II. 52:137–252, 1980. ISSN 0073-8301. doi: 10.1007/BF02684780. URL https://www.numdam.org/item/PMIHES_1980__52__137_0/.

Sudhir R. Ghorpade and Gilles Lachaud. étale cohomology, lefschetz theorems and number of points of singular varieties over finite fields. 2(3):589–631, 2002. ISSN 1609-3321. doi: 10.48550/arXiv.0808.2169. URL https://arxiv.org/abs/0808.2169.

Zbigniew Jelonek. On the effective nullstellensatz. *Inventiones Mathematicae*, 162(1):1–17, 2005. doi: 10.1007/s00222-004-0434-8.

Nicholas M. Katz. Sums of betti numbers in arbitrary characteristic. 7(1):29–44, 2001. ISSN 1071-5797. doi: 10.1006/ffta.2000.0303. URL https://doi.org/10.1006/ffta.2000.0303.

Pascal Koiran. Hilbert's nullstellensatz is in the polynomial hierarchy. *Journal of Complexity*, 12(4):273–286, 1996. ISSN 0885-064X. doi: 10.1006/jcom.1996.0019. URL https://www.sciencedirect.com/science/article/pii/S0885064X96900199.

Jeffrey C. Lagarias and Andrew M. Odlyzko. Effective versions of the chebotarev density theorem. *Algebraic Number Fields*, pages 409–464, 1977.

Serge Lang and André Weil. Number of points of varieties in finite fields. *American Journal of Mathematics*, 76(4):819–827, 1954. doi: 10.2307/2372655.