Electronic Colloquium on Computational Complexity, Revision 3 of Report No. 52 (2025)

Deterministic Depth-4 PIT and Normalization

Zeyu Guo* Siki Wang†

Abstract

In this paper, we initiate the study of deterministic PIT for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits over fields of any characteristic, where k and δ are bounded. Our main result is a deterministic polynomial-time black-box PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$ circuits, under the additional condition that one of the summands at the top Σ gate is squarefree.

Our techniques are purely algebro-geometric: they do not rely on Sylvester–Gallai-type theorems, and our PIT result holds over arbitrary fields.

The core of our proof is based on the normalization of algebraic varieties. Specifically, we carry out the analysis in the integral closure of a coordinate ring, which enjoys better algebraic properties than the original ring.

1 Introduction

Polynomial Identity Testing (PIT) is the fundamental problem of deciding whether a given multivariate algebraic circuit computes the identically zero polynomial. Equivalently, it asks whether two arithmetic circuits compute the same polynomial, by checking whether their difference is identically zero. Despite its simple formulation, PIT is one of the central problems in the study of derandomization. On the one hand, it admits an efficient randomized algorithm: by the Schwartz–Zippel Lemma [Sch80, Zip79], evaluating a degree-d polynomial at a random point over a sufficiently large field yields a correct identity test with high probability. This leads to a fast Monte Carlo algorithm that treats the circuit as a black box. On the other hand, no efficient deterministic algorithm is known, even when the circuit's structure is fully accessible. In complexity-theoretic terms, PIT lies in the class coRP, but whether it can be solved in P remains a major open question. Closing this gap is widely regarded as a central challenge in theoretical computer science. Indeed, PIT is often viewed as the algebraic analog of the classic P vs. BPP (or P vs. RP) question.

The importance of PIT lies both in its algorithmic applications and its deep connections to complexity theory. It has been used in a wide range of settings, such as primality testing [AKS04], polynomial factoring [KSS14, BSV20, KRS24, DST24], and perfect matching [Lov79, MVV87, CRS95, FGT21, ST17]. More fundamentally, PIT plays a central role in the *hardness vs. randomness* paradigm, and its derandomization is known to imply long-sought arithmetic circuit lower bounds [HS80, KI04].

Over the past few decades, there have been numerous results on deterministic PIT for various restricted models, including sparse polynomials, depth-3 circuits, low-depth circuits, and read-once

*Department of Computer Science and Engineering, The Ohio State University. Email: zguotcs@gmail.com. Supported by the NSF CAREER award (CCF-2440926).

†Computing and Mathematical Sciences Department, Caltech. Email: siki.wang@caltech.edu

ISSN 1433-8092

arithmetic branching programs, among others. See, for example, [LST24, AF22] for recent breakthroughs, including subexponential-time deterministic PIT algorithms for low-depth circuits over fields of characteristic zero or large characteristic. For surveys of earlier and recent developments, see [Sax09, Sax14, SY10, DG24]. Nonetheless, despite decades of significant progress, a general polynomial-time deterministic PIT algorithm remains elusive. As such, PIT continues to serve as a central testing ground for our understanding of algebraic structure, pseudorandomness, and computational hardness.

Depth-4 PIT. A major breakthrough by Agrawal and Vinay [AV08] revealed that PIT for depth-4 circuits essentially captures the full complexity of general PIT. In particular, for n-variate polynomials of degree $O(n)$, subexponential-size arithmetic circuits can be reduced to depth-4 circuits of subexponential size. Thus, proving exponential lower bounds for depth-4 circuits would imply exponential lower bounds for general circuits. Combined with hardness-vs-randomness connections [NW94, KI04], they further showed that a complete derandomization of depth-4 PIT, even for circuits with slightly unbounded second-layer fan-in and $O(\log n)$ bottom fan-in, would yield a quasi-polynomial-time deterministic PIT algorithm for general circuits. As a result, depth-4 PIT has become an important frontier in algebraic derandomization. Notably, no analogous constant-depth reduction is known in the Boolean setting.

Despite this equivalence, the expressive power of general depth-4 circuits makes deterministic PIT for them highly challenging. This has motivated a refined line of work focusing on the syntactic subclass $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$, consisting of depth-4 circuits with top fan-in at most k and bottom fan-in at most δ, where k and δ are bounded. The study of deterministic PIT for these circuits has led to the development of new techniques grounded in Sylvester–Gallai-type arguments and tools from algebraic geometry.

1.1 Previous Work

Sylvester–Gallai-based approach. A $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuit consists of a top Σ gate of fan-in at most k, followed by alternating layers of unbounded Π, Σ, and Π gates, where the bottom Π gates have fan-in at most δ. This model naturally generalizes bounded top fan-in depth-3 circuits ($\Sigma^{[k]}\Pi\Sigma$), for which efficient deterministic PIT algorithms are known. Over fields of characteristic zero, Sylvester–Gallai (SG)-type theorems [EK66, BM90] from combinatorial geometry have been used to establish constant-rank bounds for simple minimal $\Sigma^{[k]}\Pi\Sigma$ identities, leading to deterministic polynomial-time black-box PIT algorithms [DS07, KS09, SS13].

Gupta [Gup14] conjectured a nonlinear Sylvester–Gallai-type statement which, if true, would yield deterministic polynomial-time black-box PIT algorithms for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits with bounded k and δ over fields of characteristic zero. Motivated by this conjecture, Shpilka proved an SG-type theorem for quadratic polynomials, initiating a substantial line of work [Shp20, PS22, PS21, GOS22, OS22, GOPS23, OS24, GOS25b, GOS25a, GOS$^+$26] aimed at fully resolving the conjecture.

Peleg and Shpilka [PS21], building on earlier work [Shp20, PS22], proved a quadratic SG-type theorem that yields a deterministic polynomial-time black-box PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[2]}$ circuits. Subsequent works [OS22, GOPS23, OS24, GOS25b] extended these results and introduced new techniques, including those based on Stillman uniformity [AH20, ESS19, DLL19], culminating in the work of Garg, Oliveira, and Sengupta [GOS25a], which established constant-rank bounds for nonlinear Edelstein–Kelly configurations and gave a deterministic polynomial-time black-box PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$ circuits over fields of characteristic zero for any bounded δ. More

recently, Garg, Oliveira, Sengupta, Shalmon, and Shpilka [GOS+26] established similar results for $\Sigma^{[k]}\Pi\Sigma\Pi^{[2]}$ circuits over fields of characteristic zero for any bounded k.

However, SG-based methods remain largely confined to characteristic zero. Kayal and Saxena [KS07] showed that the constant-rank bounds implied by SG-type arguments do not hold in positive characteristic, even for the simpler $\Sigma^{[k]}\Pi\Sigma$ model.

Other work. Dutta, Dwivedi, and Saxena [DDS21] gave a deterministic quasipolynomial-time black-box PIT algorithm for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits. Due to their use of logarithmic derivatives, the result holds only when the characteristic of the base field is zero or sufficiently large. Their approach reduces the problem to black-box PIT for read-once oblivious arithmetic branching programs (ROABPs). However, obtaining a deterministic polynomial-time black-box PIT algorithm for this class remains open despite extensive effort.

We also mention the breakthrough of superpolynomial lower bounds for low-depth circuits by Limaye, Srinivasan, and Tavenas [LST24]. This result, together with the earlier work of Chou, Kumar, and Solomon [CKS18] or the work of Andrews and Forbes [AF22], yields subexponential-time black-box PIT algorithms for low-depth circuits, including $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits. These algorithms also rely on properties that only hold when the characteristic is zero or large. Subsequently, Forbes [For24] extended the superpolynomial lower bound for low-depth circuits to positive characteristics. However, whether the PIT algorithms themselves can be extended to this setting remains an open question [For24].

In summary, all of the depth-4 PIT algorithms discussed above require characteristic zero or sufficiently large characteristic. This naturally raises the following question: Is it possible to design a nontrivial, or even polynomial-time, deterministic PIT algorithm for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits over arbitrary fields? We believe the answer is yes. In this work, we make progress toward answering this question affirmatively.

As a reality check, any such result must first address the simpler case of $\Sigma^{[k]}\Pi\Sigma$ circuits. Deterministic polynomial-time algorithms for this class are indeed known in positive characteristics: they were first obtained by Kayal and Saxena in the white-box setting [KS07], and later by Saxena and Seshadhri in the black-box setting [SS12]. Notably, these results do not rely on SG-type theorems or other arguments that require assumptions on the characteristic of the base field. Our work is inspired by these approaches and seeks to extend them to the depth-4 setting.

1.2 Our Results

In this paper, we initiate the study of deterministic PIT for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits over fields of any characteristic. Our main result is a deterministic polynomial-time black-box PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$ circuits under the condition that one of the summands at the top Σ gate is squarefree.

Our techniques are purely algebro-geometric. Notably, the proof does not rely on Sylvester–Gallai-type theorems, and the result remains valid even over fields of small positive characteristic.

We begin by stating the homogeneous version of our result.

Theorem 1.1 (Main theorem, homogeneous version)**.** *Let $C_{n,d,k,\delta,\mathbb{F}}$ be the set of polynomials $F \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ over a field $\mathbb{F}$ satisfying the following conditions:*

(1) F can be expressed as a sum $F = \sum_{i=0}^{k_0-1} F_i$, where $k_0 \le k$, $F_i = \prod_{j=1}^{m_i} f_{i,j}$ for $i \in \{0, 1, \ldots, k_0 - 1\}$, and each $f_{i,j} \in \mathbb{F}[\mathbf{X}]$ is a nonzero homogeneous polynomial of degree at most δ.

(2) $\deg(F_i) = d_0$ *for some* $d_0 \le d$ *and all* $i \in \{0, 1, \ldots, k_0 - 1\}$.

(3) There exists $i \in \{0, 1, \ldots, k_0 - 1\}$ *such that* F_i *is squarefree over* $\overline{\mathbb{F}}$*, meaning that the irreducible factors of* F_i *over* $\overline{\mathbb{F}}$ *are distinct.*

Then there exists an explicit $(nd)^{O_\delta(1)}$*-sized hitting set* $\mathcal{H} \subseteq \overline{\mathbb{F}}^n$ *for* $C_{n,d,3,\delta,\mathbb{F}}$*. Equivalently, there exists a deterministic black-box PIT algorithm for* $C_{n,d,3,\delta,\mathbb{F}}$ *that makes at most* $(nd)^{O_\delta(1)}$ *evaluation queries.*

The above homogeneous result easily extends to an analogous statement for inhomogeneous polynomials. We now explain how this extension works.

Let $F = \sum_{i=0}^{k_0-1} F_i \in \mathbb{F}[X_1, \ldots, X_n]$, where $F_0, \ldots, F_{k_0-1}$ are nonzero and possibly inhomogeneous polynomials. Define the homogenization of F with respect to $F_0, \ldots, F_{k_0-1}$ as

$$\mathrm{H}(F, F_0, \ldots, F_{k_0-1}) := \sum_{i=0}^{k_0-1} F_i(X_1/X_0, \ldots, X_n/X_0) X_0^{d_0}, \quad \text{where } d_0 = \max_{0 \le i \le k_0-1}(\deg(F_i)),$$

which is a homogeneous polynomial living in $\mathbb{F}[X_0, X_1, \ldots, X_n]$.

We now state the inhomogeneous version of our result.

Corollary 1.2 (Main theorem, inhomogeneous version)**.** *Let* $C^*_{n,d,k,\delta,\mathbb{F}}$ *be the set of polynomials* $F \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ *over a field* $\mathbb{F}$ *that can be written as a sum* $F = \sum_{i=0}^{k_0-1} F_i$ *for some* $k_0 \le k$ *such that* $\mathrm{H}(F, F_0, \ldots, F_{k_0-1}) \in C_{n+1,d,k,\delta,\mathbb{F}}$*. This includes all those* $F = \sum_{i=0}^{k_0-1} F_i$ *with* $k_0 \le k$ *for which the following conditions hold:*

(1) $F_i = \prod_{j=1}^{m_i} f_{i,j}$ *for* $i \in \{0, 1, \ldots, k_0 - 1\}$*, where each* $f_{i,j} \in \mathbb{F}[\mathbf{X}]$ *is a nonzero polynomial of degree at most* δ*.*

(2) $d_0 := \max_{0 \le i \le k_0-1}(\deg(F_i))$ *is at most* d*.*

(3) F_i *is squarefree over* $\overline{\mathbb{F}}$ *for some* $i \in \{0, 1, \ldots, k_0 - 1\}$ *satisfying* $\deg(F_i) = d_0$*, or more generally, satisfying* $\deg(F_i) \ge d_0 - 1$*.*

And there exists an explicit $(nd)^{O_\delta(1)}$*-sized hitting set* $\mathcal{H} \subseteq \overline{\mathbb{F}}^n$ *for* $C^*_{n,d,3,\delta,\mathbb{F}}$*. Equivalently, there exists a deterministic black-box PIT algorithm for* $C^*_{n,d,3,\delta,\mathbb{F}}$ *that makes at most* $(nd)^{O_\delta(1)}$ *evaluation queries.*

Proof. Consider $i \in \{0, 1, \ldots, k_0 - 1\}$. Note that

$$F_i(X_1/X_0, \ldots, X_n/X_0) X_0^{d_0} = \left(\prod_{j=1}^{m_i} \left(f_{i,j}(X_1/X_0, \ldots, X_n/X_0) X_0^{\deg(f_{i,j})} \right) \right) X_0^{d_0 - \deg(F_i)}.$$

Here, each factor $f_{i,j}(X_1/X_0, \ldots, X_n/X_0) X_0^{\deg(f_{i,j})}$ is not divisible by X_0. Therefore, if F_i is squarefree over $\overline{\mathbb{F}}$ and $\deg(F_i) \ge d_0 - 1$, then $F_i(X_1/X_0, \ldots, X_n/X_0) X_0^{d_0}$ is also squarefree over $\overline{\mathbb{F}}$. So if $F = \sum_{i=0}^{k_0-1} F_i$ satisfies Item 3 in Corollary 1.2, then $\mathrm{H}(F, F_0, \ldots, F_{k_0-1})$ satisfies Item 3 in Theorem 1.1. It is straightforward to verify that Item 1 and Item 2 in Corollary 1.2 imply the corresponding items in Theorem 1.1 as well.

It remains to construct an explicit $(nd)^{O_\delta(1)}$-sized hitting set $\mathcal{H}^* \subseteq \overline{\mathbb{F}}^n$ for $C^*_{n,d,3,\delta,\mathbb{F}}$. First, apply Theorem 1.1 over $\overline{\mathbb{F}}$ to obtain an explicit hitting set $\mathcal{H} \subseteq \overline{\mathbb{F}}^{n+1}$ for $C_{n+1,d,3,\delta,\overline{\mathbb{F}}}$.

Let us first assume that $a_0 \neq 0$ for all $(a_0, \ldots, a_n) \in \mathcal{H}$. Construct the set

$$\mathcal{H}^* = \{(a_1/a_0, \ldots, a_n/a_0) : (a_0, \ldots, a_n) \in \mathcal{H}\}.$$

Consider a nonzero polynomial $F = \sum_{i=0}^{k_0-1} F_i \in C^*_{n,d,3,\delta,\mathbb{F}}$. Then $\mathrm{H}(F, F_0, \ldots, F_{k_0-1})$ is a nonzero polynomial in $C_{n+1,d,3,\delta,\mathbb{F}}$. Choose $\mathbf{a} = (a_0, \ldots, a_n) \in \mathcal{H}$ such that $\mathrm{H}(F, F_0, \ldots, F_{k_0-1})(\mathbf{a}) \neq 0$. By definition,

$$\mathrm{H}(F, F_0, \ldots, F_{k_0-1})(\mathbf{a}) = \sum_{i=0}^{k_0-1} F_i(a_1/a_0, \ldots, a_n/a_0)a_0^{d_0} = F(a_1/a_0, \ldots, a_n/a_0)a_0^{d_0}.$$

Therefore, $F(a_1/a_0, \ldots, a_n/a_0) = \mathrm{H}(F, F_0, \ldots, F_{k_0-1})(\mathbf{a})a_0^{-d_0} \neq 0$, where $(a_1/a_0, \ldots, a_n/a_0) \in \mathcal{H}^*$ by construction. So $\mathcal{H}^*$ is a hitting set for $C^*_{n,d,3,\delta,\mathbb{F}}$, whose size is at most $|\mathcal{H}| = (nd)^{O_\delta(1)}$.

Finally, we remove the assumption that $a_0 \neq 0$ for all $(a_0, \ldots, a_n) \in \mathcal{H}$. Observe that the class $C_{n+1,d,3,\delta,\overline{\mathbb{F}}}$ is closed under invertible linear transformations of the coordinates.

We may assume that $\mathbf{0} = (0, \ldots, 0) \notin \mathcal{H}$, since any non-constant homogeneous polynomial always vanishes at $\mathbf{0}$. Hence, every point in $\mathcal{H}$ has at least one nonzero coordinate.

We can deterministically and efficiently find an invertible linear transformation $\phi : \overline{\mathbb{F}}^{n+1} \to \overline{\mathbb{F}}^{n+1}$ such that every point $(a_0, \ldots, a_n) \in \phi(\mathcal{H})$ satisfies $a_0 \neq 0$. We then replace $\mathcal{H}$ by $\phi(\mathcal{H})$ and construct $\mathcal{H}^*$ as before. □

We remark that removing Item 3 from Corollary 1.2 (the squarefreeness condition) recovers the class $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$.

For general constant δ, no subexponential-time deterministic black-box PIT algorithm was previously known for these classes over fields of arbitrary small positive characteristic.

1.3 Proof Overview

To explain our ideas, we begin with a solved case: consider a nonzero polynomial $F \in \Sigma^{[2]}\Pi\Sigma\Pi^{[\delta]}$, i.e., a depth-4 circuit with top fan-in two. In this case, we can write $F = F_1 + F_2 \neq 0$, where F_1 and F_2 are products of nonzero polynomials, each of degree at most δ.

Suppose

$$F_1 = g_1 \cdots g_r \quad \text{and} \quad F_2 = h_1 \cdots h_s \tag{1}$$

are factorizations of F_1 and F_2 into irreducible polynomials, respectively. If $F_2 = cF_1$ for some $c \in \mathbb{F}^\times$, then $F = (c+1)F_1$, which is again a product of degree-($\leq \delta$) nonzero polynomials. PIT for such polynomials is straightforward. So assume this is not the case. Then the factorizations in (1) do not "match." That is, there does not exist a bijection $\sigma : [r] \to [s]$ such that g_i and $h_{\sigma(i)}$ are scalar multiples of each other for all $i \in [r]$.

It is natural to choose an affine line ℓ such that restricting to ℓ reduces the ambient dimension while preserving the non-matching structure of the factorizations. In other words, we want $F_1|_\ell$ and $F_2|_\ell$ to still have non-matching factorizations. This ensures that $(F_1 + F_2)|_\ell \neq 0$, reducing the problem to PIT for univariate polynomials, which is easy.

To ensure the factorizations remain non-matching after restriction, it suffices to preserve the degree of every irreducible factor and to ensure that any two coprime irreducible factors $g, h \in \{g_1, \ldots, g_r, h_1, \ldots, h_s\}$ remain coprime after restricting to ℓ, i.e., they do not share a common root

on the line. Geometrically, the latter condition amounts to ensuring that ℓ avoids the codimension-two variety $V(g, h)$. In [Guo24], it is shown how to construct a polynomial-sized set of lines such that most lines avoid all such varieties. This effectively (re)solves the problem.

We now move to the class $C_{n,d,3,\delta,\mathbb{F}}$. Let F be a nonzero polynomial in this class. Then we may write

$$F = F_0 + F_1 + F_2,$$

where F_0, F_1, F_2 are homogeneous of the same degree and are products of nonzero homogeneous polynomials, each of degree at most δ. Furthermore, one of the summands, say F_0, is squarefree over $\overline{\mathbb{F}}$.

In the study of Boolean circuits, one common technique is applying a restriction (i.e. partial assignment) to simplify the circuit. An analogous idea works here: we restrict F to the zero locus of an irreducible factor θ of F_0 to eliminate F_0. Algebraically, this corresponds to working in the quotient ring $\mathbb{F}[X_1, \ldots, X_n]/(\theta)$, where each F_i is replaced by $\overline{F}_i := F_i \bmod \theta$. Since θ divides F_0, we have $\overline{F}_0 = 0$. Thus, we have effectively reduced to the $k = 2$ case, though now over the quotient ring.

Why can't we directly reuse the $k = 2$ argument? The issue is that $\mathbb{F}[X_1, \ldots, X_n]/(\theta)$ is not, in general, a unique factorization domain (UFD), so the factorizations of $\overline{F}_1$ and $\overline{F}_2$ are not well-defined.

But is unique factorization truly necessary? We argue that it is not: even in the absence of unique factorization, working in rings with the weaker property of *normality* still enables us to obtain meaningful results.

Normality. Let A be an integral domain, whose field of fractions is denoted by $\mathrm{Frac}(A)$. We say A is *integrally closed* if for any monic polynomial $P(X) \in A[X]$, all roots of $P(X)$ in $\mathrm{Frac}(A)$ are in A. An irreducible affine variety is said to be *normal* if its coordinate ring is integrally closed.

It is not easy to give a purely geometric definition of normality. However, its usefulness lies in the fact that, if V is a normal variety with coordinate ring A, then for each codimension-one irreducible subvariety $\mathcal{Z} \subseteq V$, there is a well-behaved "order function" $\mathrm{ord}_{\mathcal{Z}} : \mathrm{Frac}(A) \to \mathbb{Z} \cup \{\infty\}$ indicating the order of zeros or poles of every $g \in \mathrm{Frac}(A)$ along $\mathcal{Z}$. For example, for $A = \mathbb{F}[X, Y]$ and $g = (X + Y)^2/X^3 \in \mathrm{Frac}(A)$, we have $\mathrm{ord}_{V(X+Y)}(g) = 2$ and $\mathrm{ord}_{V(X)}(g) = -3$.

If A is normal, then for $g \in A \setminus \{0\}$, one can define a "generalized factorization" of g, where the "irreducible factors" are not polynomials, but codimension-one irreducible subvarieties $\mathcal{Z}$ of V, each with multiplicity $\mathrm{ord}_{\mathcal{Z}}(g)$. (This is called the Weil divisor associated with g, written additively as $\mathrm{div}(g) := \sum_{\mathcal{Z}} \mathrm{ord}_{\mathcal{Z}}(g) \cdot \mathcal{Z}$.)

With this generalized notion of factorization, one can carry out an argument analogous to (and in fact generalizing) the one for the $k = 2$ case, assuming the variety defined by the factor θ of F_0 is normal.

Thus, normality may be viewed as a useful weakening of unique factorization. In general, however, the variety in question may even fail to be normal. To address this, we apply a standard technique from algebraic geometry known as *normalization*.

Normalization. Conceptually, the normalization of a variety V produces a normal variety $\widetilde{V}$ that best approximates V among all normal varieties. Algebraically, if V is affine, normalization corresponds to taking the integral closure $\widetilde{\mathbb{F}[V]}$ of the coordinate ring $\mathbb{F}[V]$ in its field of fractions. Our key idea is to work within the "nicer" ring $\widetilde{\mathbb{F}[V]}$ in place of $\mathbb{F}[V]$.

The strategy of enlarging a ring to recover a weak form of unique factorization has deep roots in number theory. It began with Kummer's introduction of *ideal numbers* to address the failure of unique factorization for cyclotomic integers, and was later formalized by Dedekind through the theory of *ideals*, recovering unique factorization at the level of ideal decomposition in number rings. In modern terms, this philosophy is embodied in the process of normalization—passing to the integral closure of a ring in its field of fractions—which yields an integrally closed ring that better approximates a unique factorization domain.

As an illustrative example, work over $\mathbb{R}$ and consider the plane curve C defined by $Y^2 - X^2(X+1) = 0$ (see Fig. 1). The rational function $Z = Y/X$ satisfies the monic polynomial $Z^2 - (X+1)$ on C, but is not a regular function in $\mathbb{R}[C]$, meaning that it is not well-defined on the curve. Geometrically, this reflects the fact that C has two branches at the point $(0,0)$, where the limits of Y/X approach 1 and -1, respectively. Introducing Z as a new coordinate function separates these branches. One can show that X, Y, and Z generate an integrally closed ring in $\mathrm{Frac}(\mathbb{R}[C])$, and that their defining relations are generated by $Y^2 - X^2(X+1)$, $XZ = Y$, and $Z^2 - (X+1)$. These equations define the normalization $\widetilde{C} \subseteq \mathbb{A}^3$.

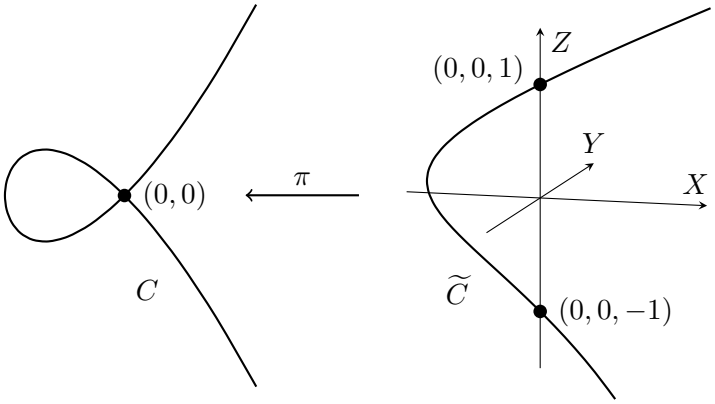


Figure 1: Normalization $\widetilde{C}$ of the curve C defined by $Y^2 - X^2(X+1) = 0$. The map $\pi : \widetilde{C} \to C$ sends $(x, y, z) \mapsto (x, y)$.

This example illustrates why normalization is useful for defining a generalized notion of factorization. On the non-normal curve C, the singular point $(0,0)$ behaves like two overlapping points, with functions like Y/X exhibiting distinct limiting behavior along each branch. After normalization, these branches are separated into two points, $p_1 = (0,0,1)$ and $p_2 = (0,0,-1)$, each admitting a well-behaved order function $\mathrm{ord}_{p_i}(\cdot)$.

A knowledgeable reader may recognize the above example as the resolution of singularities of C. Indeed, a normal variety is regular in codimension one [Sta26, Tag 0345]. Consequently, the normalization of a curve is regular; over a perfect field, it is smooth [Sta26, Tag 0B8Y] and hence provides a resolution of singularities.

In higher dimensions, by contrast, resolution of singularities is highly complex in characteristic zero [Hir64a, Hir64b], and remains an open problem in positive characteristic. Normalization, on the other hand, is significantly more tractable. A classical theorem of Emmy Noether shows that the integral closure of a coordinate ring $\mathbb{F}[V]$ is finitely generated as a module over $\mathbb{F}[V]$ (see, e.g., [Eis95]).

Constructive normalization has also been extensively studied; see [Sto68, Sei70, Sei75, Tra84,

dJ98]. In this paper, we only require normalization for curves, and we follow the framework of Trager [Tra84] to perform this task.

That said, normalization is generally regarded as a computationally expensive problem, often requiring tools such as Gröbner bases. In the context of Geometric Complexity Theory, a key challenge is the non-normality of certain orbit closures related to the determinant and permanent, a fact established by Kumar [Kum13]. In that setting, the potential utility of normalization remains unclear due to its high complexity.

In our setting, however, this complexity is not a barrier. We apply a dimensionality reduction technique that restricts the input to a carefully chosen constant-dimensional subspace, and specifically, a plane. After restriction to a plane, the number of variables and the relevant degree bounds depend only on δ. Controlling the coefficient complexity in the remaining parameters then allows us to obtain a polynomial-time hitting set construction.

Dimensionality reduction. We reduce the dimension of the ambient space by restricting to a constant-dimensional subspace (specifically, a plane) while preserving all the structural properties required for our analysis. This naturally raises the question: How should one choose such a plane?

It is not hard to show that a randomly chosen plane will work with high probability. However, using randomness would contradict the goal of this work: derandomization. Instead, we restrict to a *generic plane*. That is, we treat the parameters defining the plane as indeterminates $\mathbf{Y} = (Y_1, \ldots, Y_\ell)$, where $\ell = \Theta(n)$, and work over the function field $\mathbb{F}(\mathbf{Y})$ in place of the original base field $\mathbb{F}$. The analysis is carried out symbolically over this function field, and we eventually specialize $\mathbf{Y}$ to a tuple $\mathbf{a} \in \mathbb{F}^\ell$.

This specialization step is constructive. We identify a small collection of polynomials $Q_1, \ldots, Q_m$ in the ring $\mathbb{F}[\mathbf{Y}]$, each of *bounded degree*, such that any assignment $\mathbf{Y} \leftarrow \mathbf{a}$ satisfying $Q_i(\mathbf{a}) \neq 0$ for all i ensures that the chosen plane preserves the desired properties. Verifying such an assignment reduces to testing whether each Q_i is nonzero at $\mathbf{a}$, which is a bounded-degree PIT problem and can be solved deterministically in polynomial time [KS01].

Such generic-to-specific arguments are common in algebraic complexity. Notably, they appear in Kaltofen's work on multivariate polynomial factorization [Kal95] (see also [KSS14]), and in Sylvester–Gallai-based approaches to PIT for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\delta]}$ circuits, e.g., [Gup14]. The general philosophy—prove that a generic geometric object satisfies a given property, and then deduce that the property holds for a dense open set of specific instances—is also a standard principle in algebraic geometry.

Carrying out this analysis introduces several technical challenges, which we address in this paper. First, the function field $\mathbb{F}(\mathbf{Y})$ is not algebraically closed, so arguments must be written more carefully and sometimes more abstractly. Second, we need to compute Gröbner bases over $\mathbb{F}(\mathbf{Y})$. While degree bounds on the polynomials in a Gröbner basis are well-known [MM82, Dub90], we additionally require bounds on the *coefficient complexity*. Specifically, we need to bound the degrees of numerators and denominators of the rational functions in $\mathbb{F}(\mathbf{Y})$ that appear as coefficients. Since such bounds are not readily available in the literature, we establish them ourselves. These bounds may be of independent interest for other problems concerning algebraic pseudorandomness.

Third, although we may assume that $\mathbb{F}$ is a *perfect field* by passing to its algebraic closure, the function field $\mathbb{F}(\mathbf{Y})$ may still be non-perfect in characteristic $p > 0$. In particular, field extensions over $\mathbb{F}(\mathbf{Y})$ can be inseparable, complicating tasks such as computing radicals or primary decompositions. To address this, we adjoin p^e-th roots of the variables $\mathbf{Y}$ for sufficiently large e, thereby

passing to the extended function field $\mathbb{F}(\mathbf{Y}^{1/p^e}) := \mathbb{F}(Y_1^{1/p^e}, \ldots, Y_\ell^{1/p^e})$, over which the relevant extensions become separable.

The condition of squarefreeness. Finally, we explain where and why the squarefreeness condition is needed in our paper. This condition is specifically used in Lemma 5.5. To explain why it is needed, consider again a nonzero polynomial $F = F_0 + F_1 + F_2$, where F_0, F_1, F_2 are homogeneous of the same degree and are products of nonzero homogeneous polynomials, each of degree at most δ. When F_0 is squarefree, then, after factoring out all common divisors of F_0, F_1, and F_2 and disposing of the case $F_1 + F_2 = cF_0$, we can ensure that F_0 admits an irreducible factor θ such that $F \equiv F_1 + F_2 \not\equiv 0 \mod \theta$, allowing the approach sketched above to proceed.

In general, however, it may happen that $F \equiv F_1 + F_2 \not\equiv 0 \mod \theta^e$ but $F \equiv 0 \mod \theta$, where θ^e is the largest power of θ dividing F_0. The issue is that the geometric space associated with the ring $\mathbb{F}[X_1, \ldots, X_n]/(\theta^e)$ is not a variety, but rather a "variety with multiplicity e." Normalization recovers only the reduced structure, namely the underlying variety without multiplicities, and F may vanish identically on this variety.

In other words, while we can conveniently track and compare multiplicities along codimension-one subvarieties on the normalization of $V(\theta)$ (the variety defined by θ), the multiplicity of $V(\theta)$ itself is lost in the process. This suggests that other techniques are needed to address the general case, and we leave this important question for future work.

Organization of this paper. We begin with preliminaries in Section 2, where we introduce our notations and review relevant notions and facts from commutative algebra and algebraic geometry.

In Section 3, we develop computational tools, with a focus on Gröbner bases. Specifically, we define a measure of coefficient complexity for polynomials over the function field $\mathbb{F}(\mathbf{Y})$ in terms of the degrees of numerators and denominators of their coefficients. Using Gröbner bases, we analyze the constructivity of various fundamental computational tasks and derive bounds on coefficient complexity. These tasks include solving systems of linear equations, computing Gröbner bases, variable elimination, computing ring homomorphisms, and computing the generators of maximal ideals, among others.

In Section 4, we use the tools from Section 3 to give an essentially constructive normalization of an affine curve—equivalently, the computation of the integral closure of an affine curve's coordinate ring. We begin in Section 4.1 with basic notions from algebraic number theory, including orders, integral bases, and discriminants. In Section 4.2, we construct the normalization following the framework of Trager [Tra84], which is a function field analog of an earlier algorithm by Ford and Zassenhaus [For78] for computing the integral closure of a number ring.

In Section 5, we prove our main result (Theorem 1.1). The proof relies on the computational tools from Section 3 and the constructive normalization from Section 4.

Finally, in Section 6, we conclude the paper and discuss directions for future work.

2 Preliminaries

Denote by $[n]$ the set $\{1, 2, \ldots, n\}$. For a polynomial f and a monomial m, let $\operatorname{coeff}_f(m)$ denote the coefficient of m in f. For polynomials f and $g \neq 0$, we write $g \mid f$ to indicate that g divides f.

It is well-known that designing deterministic black-box PIT algorithms is equivalent to constructing explicit hitting sets. We now define this notion formally.

Definition 2.1 (Hitting set)**.** *A finite collection $\mathcal{H} \subseteq \mathbb{F}^n$ of points is said to be a* hitting set *for a polynomial $f \in \mathbb{F}[X_1, \ldots, X_n]$ if either $f = 0$ or $f(\mathbf{a}) \neq 0$ for some $\mathbf{a} \in \mathcal{H}$. For a family $\mathcal{C} \subseteq \mathbb{F}[X_1, \ldots, X_n]$ of polynomials, we say $\mathcal{H} \subseteq \mathbb{F}^n$ is a hitting set for $\mathcal{C}$ if it is a hitting set for every $f \in \mathcal{C}$.*

Let $\varepsilon \in [0, 1]$. We say $\mathcal{H} \subseteq \mathbb{F}^n$ is an ε-hitting set *for $f \in \mathbb{F}[X_1, \ldots, X_n]$ if either $f = 0$ or $\Pr_{\mathbf{a} \in \mathcal{H}}[f(\mathbf{a}) \neq 0] \geq 1 - \varepsilon$. We say $\mathcal{H}$ is an ε-hitting set for $\mathcal{C}$ if it is an ε-hitting set for every $f \in \mathcal{C}$.*

We focus on constructing explicit hitting sets, although the stronger notion of ε-hitting sets will also be used. Given a hitting set $\mathcal{H}$ for polynomials of degree at most d, one can boost it to an ε-hitting set for any $\varepsilon > 0$, assuming that $\mathbb{F}$ is large enough: For $|\mathcal{H}| \geq 2$, interpolate a polynomial curve C of degree at most $|\mathcal{H}| - 1$ through $\mathcal{H}$, and choose more than $d(|\mathcal{H}| - 1)/\varepsilon$ distinct points on C. The resulting set is an ε-hitting set. The case $|\mathcal{H}| = 1$ is immediate.

2.1 Commutative Algebra

All rings are assumed to be commutative with unity. The algebraic closure of a field $\mathbb{F}$ is denoted by $\overline{\mathbb{F}}$.

The set of $n \times m$ matrices over A is denoted by $A^{n \times m}$, or more generally by $A^{S \times T}$, where S and T are the index sets for rows and columns, respectively. We often use $M[i, j]$ to denote the (i, j)-th entry of a matrix M.

The ideal of a ring A generated by a set $S \subseteq A$ is denoted $\langle S \rangle$. And the ideal of A generated by $f_1, \ldots, f_k \in A$ is denoted $\langle f_1, \ldots, f_k \rangle$.

Prime, maximal, and radical ideals. A proper ideal $\mathfrak{p}$ of a ring A is a *prime ideal* if $ab \in \mathfrak{p} \implies a \in \mathfrak{p}$ or $b \in \mathfrak{p}$. Equivalently, $\mathfrak{p}$ is prime if $A/\mathfrak{p}$ is an integral domain.

A proper ideal $\mathfrak{p}$ of A is a *maximal ideal* if it is maximal among all ideals properly contained in A. Equivalently, $\mathfrak{p}$ is maximal if $A/\mathfrak{p}$ is a field. All maximal ideals are prime.

For an ideal $\mathfrak{p}$ of A, the *radical* of $\mathfrak{p}$ is the ideal $\sqrt{\mathfrak{p}} := \{a \in A : a^k \in \mathfrak{p} \text{ for some } k \geq 1\}$. An ideal $\mathfrak{p}$ is *radical* if $\mathfrak{p} = \sqrt{\mathfrak{p}}$. Prime ideals and maximal ideals are both radical.

Absolute irreducibility. A polynomial $f \in \mathbb{F}[\mathbf{X}]$ is *absolutely irreducible* over $\mathbb{F}$ if it is irreducible over the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$. This is equivalent to f being irreducible over every algebraic extension of $\mathbb{F}$. By convention, the zero polynomial is not considered irreducible.

Separability. A polynomial $f(X) \in \mathbb{F}[X]$ is said to be *separable* over $\mathbb{F}$ if its roots are distinct in the algebraic closure of $\mathbb{F}$. This is equivalent to the statement that $f(X)$ and its derivative $f'(X)$ are coprime. An algebraic element over $\mathbb{F}$ is said to be separable over $\mathbb{F}$ if its minimal polynomial over $\mathbb{F}$ is separable. The set of separable elements in an algebraic extension $\mathbb{K}/\mathbb{F}$ forms a subfield of $\mathbb{K}$ containing $\mathbb{F}$, called the *separable closure* of $\mathbb{F}$ in $\mathbb{K}$. An algebraic extension $\mathbb{K}/\mathbb{F}$ is separable if the separable closure of $\mathbb{F}$ in $\mathbb{K}$ equals $\mathbb{K}$. Otherwise, it is inseparable.

A field $\mathbb{F}$ is *perfect* if every irreducible polynomial over $\mathbb{F}$ is separable. Examples of perfect fields include fields of characteristic zero, finite fields, and algebraically closed fields.

An irreducible polynomial can be inseparable over non-perfect fields, which must have positive characteristic. We will need the following lemma to address this issue.

Lemma 2.2. *Suppose $\mathbb{F}$ is a field of characteristic $p > 0$, and $\mathbb{K}/\mathbb{F}$ is a finite extension. Let $e \geq 0$ be the largest integer such that p^e divides $[\mathbb{K} : \mathbb{F}]$. For $i \geq 0$, let $\mathbb{F}^{(i)}$ be the field $\{a^{1/p^i} : a \in \mathbb{F}\}$. Then for $e' \geq e$, every $b \in \mathbb{K}$ is separable over $\mathbb{F}^{(e')}$.*

Proof. Let $b \in \mathbb{K}$, and let $f(X)$ be its minimal polynomial over $\mathbb{F}$. Then there exists $e_0 \leq e'$ such that over $\mathbb{F}^{(e_0)}$, we may write

$$f(X) = g(X)^{p^{e_0}},$$

such that $g \in \mathbb{F}^{(e_0)}[X]$ has a monomial whose degree is not divisible by p. In particular, the derivative g' of g is nonzero. It suffices to show that b is separable over $\mathbb{F}^{(e_0)}$. Since b is a root of g, and $g' \neq 0$, it remains only to show that g is irreducible over $\mathbb{F}^{(e_0)}$.

Assume, for contradiction, that $g(X) = u(X)v(X)$ for some nonconstant $u(X), v(X) \in \mathbb{F}^{(e_0)}[X]$. Then f factors over $\mathbb{F}$ as $f(X) = U(X)V(X)$ where $U(X) = u(X)^{p^{e_0}}$ and $V(X) = v(X)^{p^{e_0}}$, contradicting the irreducibility of f over $\mathbb{F}$. □

Localization. Localization is a construction in algebra that allows us to formally invert a chosen set of elements in a ring, effectively turning them into units. For example, $\mathbb{Q}$ is a localization of $\mathbb{Z}$, where every nonzero integer has been made invertible. We now give the formal definition.

Definition 2.3 (Localization)**.** *Let M be a module over a ring A. Let S be a* multiplicatively closed subset *of A, i.e., it holds that $1 \in S$ and $ab \in S$ for $a, b \in S$. Define $S^{-1}M$ to be the set of representations of pairs $(a, s) \in M \times S$ under the equivalence relation*

$$(a, s) \equiv (b, t) \quad \iff \quad (at - bs)u = 0 \text{ for some } u \in S.$$

Write a/s or $\frac{a}{s}$ for (a, s). Call $S^{-1}M$ the localization *of M with respect to S.*

$S^{-1}A$ is a ring equipped with addition $(a/s) + (b/t) = (at + bs)/(st)$ and multiplication $(a/s) \cdot (b/t) = (ab)/(st)$. And $S^{-1}M$ is an $S^{-1}A$-module equipped with addition $(a/s) + (b/t) = (at + bs)/(st)$ and scalar multiplication $(a/s) \cdot (b/t) = (ab)/(st)$.

Intuitively, $S^{-1}A$ is the ring obtained from A by making the elements in S invertible. Note that if $0 \notin S$ and A is an integral domain, the condition that $(at - bs)u = 0$ for some $u \in S$ in Definition 2.3 is equivalent to $at - bs = 0$ since A does not have a nonzero zerodivisor.

Definition 2.4. *When $S = \{1, f, f^2, \dots\}$ for some $f \in A$, denote $S^{-1}M$ by M_f. When $S = A \setminus \mathfrak{p}$ for some prime ideal $\mathfrak{p}$ of A, denote $S^{-1}M$ by $M_\mathfrak{p}$. When A is an integral domain and $\mathfrak{p}$ is the zero ideal of A, denote $A_\mathfrak{p}$ by* $\mathrm{Frac}(A)$*, called the* field of fractions *of A.*

We also have the following fact. See, e.g., [Mag12] for a proof.

Fact 2.5. *Let A be a ring and $f \in A$. Then A_f is isomorphic to $A[X]/\langle Xf - 1\rangle$ via the map that sends a/f^i to $aX^i + \langle Xf - 1\rangle$ for $a \in A$ and $i \in \mathbb{N}$, with the inverse map sending $X + \langle Xf - 1\rangle$ to $1/f$.*

Local rings. A ring A is *local* if it has a unique maximal ideal $\mathfrak{m}$. For a ring A and a prime ideal $\mathfrak{p}$ of A, the ring $A_\mathfrak{p}$ (see Definition 2.4) is a local ring with the unique maximal ideal $\mathfrak{p}_\mathfrak{p}$.

The following statement is well-known.

Lemma 2.6 ([Mat89, Theorem 4.7])**.** *Let A be an integral domain. Then*

$$A = \bigcap_{\text{prime ideal } \mathfrak{p} \subseteq A} A_\mathfrak{p} = \bigcap_{\text{maximal ideal } \mathfrak{m} \subseteq A} A_\mathfrak{m},$$

where the intersections are taken within $\mathrm{Frac}(A)$*.*

Krull dimension. The *(Krull) dimension* of a ring A, denoted by $\dim A$, is the supremum of the lengths of chains of prime ideals $\mathfrak{p}_0 \subsetneq \cdots \subsetneq \mathfrak{p}_\ell$ of A.

If A is a finitely generated integral domain over a field $\mathbb{K}$, then $\dim A$ equals the transcendence degree of $\mathrm{Frac}(A)$ over $\mathbb{K}$. Under the same assumption, $\dim A = \dim A_{\mathfrak{m}}$ holds for every maximal ideal $\mathfrak{m}$ of A [AM69, Corollary 11.27].

We say an ideal I of a ring A has dimension k if $\dim(A/I) = k$.

Integrality and finiteness. Let $A \subseteq B$ be commutative rings. We say $b \in B$ is *integral* over A if there exists a monic polynomial $f(X) \in A[X]$ such that $f(b) = 0$. We say B is integral over A if every $b \in B$ is integral over A. The set of elements in B integral over A is called the *integral closure* of A in B, and is a ring [AM69]. The integral closure of an integral domain A in its field of fractions $\mathrm{Frac}(A)$ is simply called the integral closure of A, and denoted $\widetilde{A}$.

Lemma 2.7 ([AM69, Corollary 5.4])**.** *Suppose C is integral over B, and B is integral over A. Then C is integral over A.*

Corollary 2.8. *Suppose $A \subseteq B \subseteq C$ are commutative rings and B is integral over A. Then $c \in C$ is integral over A if and only if it is integral over B. In other words, the integral closure of A in C equals the integral closure of B in C.*

Proof. Suppose c is integral over A. Then as $A \subseteq B$, by definition, c is also integral over B.

To see the converse, let C^* be the integral closure of B in C. Suppose c is integral over B, i.e., $c \in C^*$. As C^* is integral over B and B is integral over A, by Lemma 2.7, C^* is integral over A. So c is integral over A. □

A module or algebra M over a ring A is said to be *finite* over A if M is a finitely generated A-module. Finiteness is closely related to integrality, as indicated by the following lemma.

Lemma 2.9 ([AM69, Proposition 5.1 and Corollary 5.2])**.** *A finitely generated algebra over a ring A is a finite module over A iff it is integral over A.*

We also include the following lemma, which states that integrally closedness can be checked on the localizations at prime ideals or at maximal ideals.

Lemma 2.10 ([AM69, Proposition 5.13])**.** *Let A be an integral domain. The following are equivalent:*

(1) A is integrally closed.

(2) $A_{\mathfrak{p}}$ is integrally closed for every prime ideal $\mathfrak{p}$ of A.

(3) $A_{\mathfrak{m}}$ is integrally closed for every maximal ideal $\mathfrak{m}$ of A.

Regular local rings. Let A be a Noetherian local ring with the maximal ideal $\mathfrak{m}$. Use $\dim_k \mathfrak{m}/\mathfrak{m}^2$ to denote the dimension of $\mathfrak{m}/\mathfrak{m}^2$ as a vector space over $k = A/\mathfrak{m}$. We say a local ring A is *regular* if its Krull dimension $\dim A$ equals $\dim_k \mathfrak{m}/\mathfrak{m}^2$.

Discrete valuation rings. A (nontrivial) *discrete valuation* on a field K is a mapping $v : K \to \mathbb{Z} \cup \{\infty\}$ satisfying $v(K^\times) \neq \{0\}$ and the following conditions:

(1) $v(xy) = v(x) + v(y)$,

(2) $v(x + y) \geq \min(v(x), v(y))$, and

(3) $v(x) = \infty \iff x = 0$,

for all $x, y \in K$. We say a discrete valuation v is *normalized* if $v(K^\times) = \mathbb{Z}$.

Given a discrete valuation $v : K \to \mathbb{Z} \cup \{\infty\}$, the *valuation ring* of v is defined to be $\{x \in K : v(x) \geq 0\}$. An integral domain A is a *discrete valuation ring* if there is a discrete valuation v on $\mathrm{Frac}(A)$ such that A is the valuation ring of v.

We include the following useful characterizations of discrete valuation rings.

Lemma 2.11 ([AM69, Proposition 9.2])**.** *Let A be a Noetherian[1] local domain of dimension one with the maximal ideal $\mathfrak{m}$. Then the following are equivalent:*

(1) A is a discrete valuation ring;

(2) A is integrally closed;

(3) A is a regular local ring;

(4) $\mathfrak{m}$ is a principal ideal;

(5) There exists $t \in A$ such that every nonzero ideal of A is of the form $\langle t^k \rangle$ for some $k \geq 0$.

We call such t a *uniformizer*. Given a uniformizer $t \in A$, every nonzero element $r \in A$ can be written as $r = t^k u$ where u is a unit; the (normalized) valuation v on A is determined by $v(t^k u) = k$.

Valuation at a point. Let A be a (Noetherian) integrally closed domain of dimension one, not necessarily local. Let $\mathfrak{m}$ be a maximal ideal of A. By Lemma 2.10, the localization $A_\mathfrak{m}$ is an integrally closed local domain of dimension one. So by Lemma 2.11, it is also a regular local ring and a discrete valuation ring. We often denote the corresponding (normalized) discrete valuation by $\mathrm{ord}_\mathfrak{m}(\cdot)$.

From a geometric point of view, the maximal ideal $\mathfrak{m}$ corresponds to a point p, and $\mathrm{ord}_\mathfrak{m}(f)$ indicates the order of zero or pole of f at p. If $\mathrm{ord}_\mathfrak{m}(f) \geq 0$, then $\mathrm{ord}_\mathfrak{m}(f)$ is the order of zero of f at p; otherwise, $-\mathrm{ord}_\mathfrak{m}(f)$ is the order of the pole of f at p.

Base change. Let A be a ring. For an A-module M and an A-algebra B, their tensor product $M \otimes_A B$ over A is defined to be the A-module generated by the set of elements $\{a \otimes b : a \in M, b \in B\}$ subject to the A-bilinear relations $a \otimes b + a' \otimes b = (a + a') \otimes b$, $a \otimes b + a \otimes b' = a \otimes (b + b')$, and $c(a \otimes b) = (ca) \otimes b = a \otimes (cb)$ for $a, a' \in M$, $b, b' \in B$, and $c \in A$. The A-module $M \otimes_A B$ is also a B-module. Furthermore, if M is an A-algebra, then $M \otimes_A B$ is a B-algebra.

Intuitively, $M \otimes_A B$ is obtained from M by changing the ring of scalars from A to B. For example, if M is an A-algebra $A[X_1, \ldots, X_n]/\langle r_1, \ldots, r_m \rangle$ and $A \to B$ is a ring extension, then $M \otimes_A B$ is simply $B[X_1, \ldots, X_n]/\langle r_1, \ldots, r_m \rangle$.

[1] A ring A is Noetherian if every ideal of A is finitely generated. All rings considered in this paper are Noetherian.

Noether normalization. The Noether normalization lemma is an important lemma in commutative algebra. Roughly speaking, it states that a finitely generated algebra R over a field is not too far from a polynomial ring, in the sense that R is a finite module over a subring that is isomorphic to a polynomial ring.

Lemma 2.12 (Noether normalization lemma)**.** *Let $\mathbb{K}/\mathbb{F}$ be a field extension. Let I be an ideal of $A = \mathbb{K}[X_1, \ldots, X_n]$ such that the Krull dimension of A/I is k. Then for almost all*[2] *$\mathbf{c} = (c_{1,1}, \ldots, c_{k,n}) \in \mathbb{F}^{kn}$, it holds that:*

(1) the elements $a_1, \ldots, a_k \in A/I$ defined by $a_i = (\sum_{j=1}^{n} c_{i,j} X_j) + I$ are algebraically independent over $\mathbb{K}$, and

(2) A/I is finite over $\mathbb{K}[a_1, \ldots, a_k]$.

In particular, such a vector $\mathbf{c}$ exists if $\mathbb{F}$ is infinite.

See, e.g., [GVJZ23, Lemma 4.8] and its proof.

2.2 Algebraic Geometry

Let $\mathbb{F}$ be a field. We use $\mathbb{A}^n_{\mathbb{F}}$ (or simply $\mathbb{A}^n$) to denote the affine n-space over $\mathbb{F}$, and $\mathbb{P}^n_{\mathbb{F}}$ (or simply $\mathbb{P}^n$) to denote the projective n-space over $\mathbb{F}$. Since $\mathbb{F}$ is not necessarily algebraically closed, affine and projective spaces over $\mathbb{F}$—as well as varieties defined over $\mathbb{F}$ within them—must be carefully formalized, rather than viewed simply as sets of solutions. We refer the reader to [Mum88] and [Vak24] for formal definitions of these objects.

An affine variety V over a field $\mathbb{F}$ is equipped with a ring $\mathbb{F}[V]$, called the *coordinate ring* of V. The elements of $\mathbb{F}[V]$ are called *regular functions* on V. Intuitively, these are algebraic functions that are well-defined on all of V. For example, the coordinate ring of $\mathbb{A}^n_{\mathbb{F}}$ is simply $\mathbb{F}[X_1, \ldots, X_n]$.

The (closed) points of V correspond bijectively to the maximal ideals of $\mathbb{F}[V]$. For a point $p \in V$, the set of regular functions vanishing at p is precisely the maximal ideal corresponding to p.

Any radical ideal I of $\mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ defines an affine variety in $\mathbb{A}^n_{\mathbb{F}}$, which we denote by $V(I)$. The coordinate ring of $V(I)$ is simply $\mathbb{F}[\mathbf{X}]/I$. A regular function f on $\mathbb{A}^n_{\mathbb{F}}$ can be restricted to the variety $V(I)$, yielding the regular function $f|_{V(I)} := f + I \in \mathbb{F}[\mathbf{X}]/I$ on $V(I)$.

Define $V(S) := V(\sqrt{\langle S \rangle})$ for a set S and $V(f_1, \ldots, f_k) := V(\sqrt{\langle f_1, \ldots, f_k \rangle})$ for $f_1, \ldots, f_k \in \mathbb{F}[\mathbf{X}]$. Affine varieties of the form $V(f)$ with a single polynomial f are called (affine) hypersurfaces.

A projective space $\mathbb{P}^n_{\mathbb{F}}$ has $n+1$ homogeneous coordinates $X_0, \ldots, X_n$. A projective variety in $\mathbb{P}^n_{\mathbb{F}}$ is defined by a set of homogeneous polynomials in $\mathbb{F}[X_0, \ldots, X_n]$. If the variety is defined by a single homogeneous polynomial, it is called a *(projective) hypersurface*.

The space $\mathbb{P}^n_{\mathbb{F}}$ can be covered by $n+1$ standard affine open subsets $U_0, \ldots, U_n$, where U_i is defined as the complement of the projective hyperplane defined by $X_i = 0$, and each U_i is isomorphic to $\mathbb{A}^n_{\mathbb{F}}$. The coordinate ring of U_i is $\mathbb{F}\left[\frac{X_0}{X_i}, \ldots, \frac{X_{i-1}}{X_i}, \frac{X_{i+1}}{X_i}, \ldots, \frac{X_n}{X_i}\right]$. More generally, a projective variety $V \subseteq \mathbb{P}^n_{\mathbb{F}}$ can be covered by the affine varieties $U_0 \cap V, \ldots, U_n \cap V$.

A *rational function* on $\mathbb{P}^n_{\mathbb{F}}$ is an expression of the form $F = \frac{P}{Q}$, where P and Q are homogeneous polynomials of the same degree in $X_0, \ldots, X_n$, with $Q \neq 0$. Such a function can be restricted to the affine chart U_i by substituting $X_j \mapsto \frac{X_j}{X_i}$ for all $j = 0, \ldots, n$, which yields a rational function on $U_i \cong \mathbb{A}^n_{\mathbb{F}}$.

[2] "Almost all" here means that there exists a nonzero polynomial $Q \in \overline{\mathbb{F}}[X_{1,1}, \ldots, X_{k,n}]$ such that the claim holds for all $\mathbf{c} \in \mathbb{F}^{kn}$ satisfying $Q(\mathbf{c}) \neq 0$.

Irreducibility. An (affine or projective) variety is said to be *irreducible* if it cannot be written as the union of two proper closed subvarieties. Otherwise, it is *reducible*. Every variety V can be uniquely expressed as a finite union of maximal irreducible subvarieties, called the *irreducible components* of V.

Suppose V is an irreducible affine variety. Then $\mathbb{F}[V]$ is an integral domain. In this case, the field of fractions $\mathrm{Frac}(\mathbb{F}[V])$ is called the *function field* of V, and is denoted by $\mathbb{F}(V)$.

Given an irreducible projective variety $V \subseteq \mathbb{P}^n_{\mathbb{F}}$, we can restrict a rational function $F = \frac{P}{Q}$ on $\mathbb{P}^n_{\mathbb{F}}$ to V, provided that Q does not vanish identically on V. In fact, such a restriction can be computed locally on any nonempty open subset of V, such as $V \cap U_i$ for some i [Vak24, Exercise 5.2.I].

Let $p \in V$ correspond to a maximal ideal $\mathfrak{m} \subset \mathbb{F}[V]$. A rational function $f \in \mathbb{F}(V)$ is said to be *regular at* p if f lies in the local ring $\mathbb{F}[V]_{\mathfrak{m}}$. Intuitively, this means f is well-defined at (and near) the point p.

Indeed, we can evaluate $f \in \mathbb{F}[V]_{\mathfrak{m}}$ at p via the natural quotient map

$$f \mapsto f + \mathfrak{m}_{\mathfrak{m}} \in \kappa_{\mathfrak{m}} := \mathbb{F}[V]_{\mathfrak{m}}/\mathfrak{m}_{\mathfrak{m}} \cong \mathbb{F}[V]/\mathfrak{m}.$$

This value lies in the field $\kappa_{\mathfrak{m}}$, called the *residue field* of the point p (or of the ideal $\mathfrak{m}$). It is a finite field extension of $\mathbb{F}$.

Normality. An irreducible affine variety is said to be *normal* if its coordinate ring is an integrally closed domain.[3] An irreducible projective variety $V \subseteq \mathbb{P}^n_{\mathbb{F}}$ is normal if each of the affine pieces $V \cap U_0, \ldots, V \cap U_n$ is a normal affine variety.

Dimension. Let V be an (affine or projective) variety over an algebraically closed field $\mathbb{F}$. The *dimension* of V, denoted $\dim V$, is the largest integer d such that there exists a chain of irreducible closed subvarieties $Z_0 \subsetneq Z_1 \subsetneq \cdots \subsetneq Z_d \subseteq V$, where each Z_i is a proper irreducible subvariety of Z_{i+1}. In other words, the dimension of V is the length of the longest strictly increasing chain of irreducible closed subsets in V.

If V is a variety over a field $\mathbb{F}$ that is not algebraically closed, its dimension is defined as $\dim V := \dim V_{\overline{\mathbb{F}}}$, where $V_{\overline{\mathbb{F}}}$ denotes the base change of V to the algebraic closure $\overline{\mathbb{F}}$ of $\mathbb{F}$.

In the case of an affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}}$, the dimension of V coincides with the Krull dimension of its coordinate ring $\mathbb{F}[V]$.

Varieties of dimension one are called curves.

Degree. Suppose $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ is an irreducible affine variety over an algebraically closed field $\mathbb{F}$. The *degree* of V, denoted $\deg(V)$, is defined as the number of isolated points in the intersection of V with a general affine linear subspace of codimension d, where $d = \dim V$.

Now suppose $V \subseteq \mathbb{P}^n_{\mathbb{F}}$ is an irreducible projective variety over an algebraically closed field $\mathbb{F}$. The degree of V is defined as the number of points in the intersection of V with a general linear subspace $L \subseteq \mathbb{P}^n_{\mathbb{F}}$ of codimension d, where $d = \dim V$.

For a reducible variety, we define its degree to be the sum of the degrees of all its irreducible components, including components of different dimensions.

[3] More generally, a (possibly reducible) affine variety is normal if its coordinate ring is a finite product of integrally closed domains [Sta26, Tag 030C].

If $\mathbb{F}$ is not algebraically closed, we define $\deg(V) := \deg(V_{\overline{\mathbb{F}}})$, where $V_{\overline{\mathbb{F}}}$ denotes the base change of V to the algebraic closure $\overline{\mathbb{F}}$.[4]

It can be shown that the degree of a point p (corresponding to a maximal ideal $\mathfrak{m}$) equals the degree of its residue field extension, that is, $[\kappa_{\mathfrak{m}} : \mathbb{F}]$.

Bézout's inequality. We need the following version of Bézout's inequality.

Lemma 2.13 (Bézout's inequality [HS80, Hei83])**.** *Let V and V' be closed affine subvarieties of $\mathbb{A}^n_{\mathbb{F}}$. Then $\deg(V \cap V') \leq \deg(V) \cdot \deg(V')$.*

The following statement follows from [BM93, Proposition 3.5].

Lemma 2.14. *Let $I = \langle f_1, \ldots, f_k \rangle$ be a zero-dimensional ideal of $\mathbb{F}[X_1, \ldots, X_n]$, and let $d = \max(\deg(f_1), \ldots, \deg(f_k))$. Then $\dim_{\mathbb{F}}(\mathbb{F}[\mathbf{X}]/I) \leq d^n$. In particular, $[\kappa_{\mathfrak{m}} : \mathbb{F}] \leq d^n$ for any maximal ideal $\mathfrak{m}$ containing I.*

3 Computational Tools

The main purpose of this section is to give bounds on the complexity of various algebraic procedures. The bounds will be given in terms of the degrees of the input polynomials, as well as the degrees of their coefficients $\frac{c}{c'}$, since the coefficients come from the field of rational functions $\mathbb{F}(Y_1, \ldots, Y_\ell)$. To do this, we develop various computational tools, with a focus on Gröbner bases, and establish bounds of the form $O_{s_1,\ldots,s_r}(1)$, where $s_1, \ldots, s_r$ are parameters. While such bounds obscure the specific dependence on $s_1, \ldots, s_r$, they suffice for our purposes, as these parameters will be treated as constants in our application. Additionally, although it is theoretically possible to make these bounds explicit, doing so would be tedious and distracting, as it involves repeated composition of intermediate results.

That said, one can verify that all bounds in this section are at most constant-height towers of exponentials in the degrees and the number of variables. In most cases, the bounds are at least doubly exponential due to the use of Gröbner bases. As mentioned, this is acceptable since the degrees and number of variables will eventually be fixed constants depending only on the bottom fan-in δ.

Throughout this section, let $\mathbb{K} = \mathbb{F}(\mathbf{Y}) = \mathbb{F}(Y_1, \ldots, Y_\ell)$, where $\mathbb{F}$ is an infinite field.

To control the complexity of coefficients in $\mathbb{K}$, we introduce the following families of subsets of $\mathbb{K}$ and $\mathbb{K}[\mathbf{X}]$.

Definition 3.1 (Field elements and polynomials with bounded complexity)**.** *For an integer $d \geq 0$, let $\mathbb{K}_{\leq d}$ denote the set of $a \in \mathbb{K}$ such that $a = \frac{c}{c'}$ for some $c \in \mathbb{F}[\mathbf{Y}]$ and $c' \in \mathbb{F}[\mathbf{Y}] \setminus \{0\}$ satisfying $\deg_{\mathbf{Y}}(c), \deg_{\mathbf{Y}}(c') \leq d$.*

For integers $n, d, d' \geq 0$ and variables $\mathbf{X} = (X_1, \ldots, X_n)$, let $\mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$ denote the set of polynomials $f \in \mathbb{K}[\mathbf{X}]$ of degree at most d such that the coefficients of f are all in $\mathbb{K}_{\leq d'}$.

Lemma 3.2. *The following hold:*

[4] The base change $V_{\overline{\mathbb{F}}}$ is not necessarily a variety and must be interpreted as a *scheme* [Mum88, Vak24] to correctly define the degree. For example, consider the affine variety defined by $X^p - T$ in $\mathbb{A}^1_{\mathbb{F}_p(T)}$. After base change to $\overline{\mathbb{F}_p(T)}$, it becomes a "point of multiplicity p," since $X^p - T = (X - T^{1/p})^p$.

(1) $ca \in \mathbb{K}_{\leq d}$ *for* $a \in \mathbb{K}_{\leq d}$ *and* $c \in \mathbb{F}$.

(2) $a + b, ab \in \mathbb{K}_{\leq d+d'}$ *for* $a \in \mathbb{K}_{\leq d}$ *and* $b \in \mathbb{K}_{\leq d'}$.

(3) $1/a \in \mathbb{K}_{\leq d}$ *for* $a \in \mathbb{K}_{\leq d} \setminus \{0\}$.

Proof. The claims follow straightforwardly by definition. □

Lemma 3.3. *Let* $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$. *For* $f \in \mathbb{K}_{\leq d}[\mathbf{X}]_{\leq d_1}$ *and* $g \in \mathbb{K}_{\leq d}[\mathbf{X}]_{\leq d_2}$, *it holds that* $fg \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d_1+d_2}$ *for some* $d' = O_{d_1+d_2,d,n}(1)$.

Proof. Each coefficient of fg is the sum of $O_{d_1+d_2,n}(1)$ terms of the form $\mathrm{coeff}_f(m_1) \cdot \mathrm{coeff}_g(m_2)$, where m_1 is a monomial of f and m_2 is a monomial of g. The claim follows by Lemma 3.2 (2). □

To address the issue of inseparable extensions, we introduce the following definitions.

Definition 3.4. *Let* $p = \mathrm{char}(\mathbb{F})$, *and let* $e \geq 0$. *We use the shorthand*

$$\mathbb{F}(\mathbf{Y}^{1/p^e}) := \mathbb{F}(Y_1^{1/p^e}, \ldots, Y_\ell^{1/p^e}) \quad \textit{and} \quad \mathbb{F}[\mathbf{Y}^{1/p^e}] := \mathbb{F}[Y_1^{1/p^e}, \ldots, Y_\ell^{1/p^e}].$$

For convenience, if $\mathrm{char}(\mathbb{F}) = 0$, *we interpret* $p^e = 1$, *even though this is not formally correct.*

Define $\mathbb{K}^{(e)} := \mathbb{F}(\mathbf{Y}^{1/p^e})$. *Let* $\mathbb{K}^{(e)}_{\leq d}$ *be the set of elements* $a \in \mathbb{K}^{(e)}$ *such that* $a = \frac{c}{c'}$ *for some* $c \in \mathbb{F}[\mathbf{Y}^{1/p^e}]$, $c' \in \mathbb{F}[\mathbf{Y}^{1/p^e}] \setminus \{0\}$, *and* $\deg_{\mathbf{Y}^{1/p^e}}(c), \deg_{\mathbf{Y}^{1/p^e}}(c') \leq d$. *That is, both* c *and* c' *have degree at most* d *in the variables* Y_i^{1/p^e}.

For integers $n, d, d' \geq 0$ *and variables* $\mathbf{X} = (X_1, \ldots, X_n)$, *define* $\mathbb{K}^{(e)}_{\leq d'}[\mathbf{X}]_{\leq d}$ *to be the set of polynomials* $f \in \mathbb{K}^{(e)}[\mathbf{X}]$ *of degree at most* d, *whose coefficients are all in* $\mathbb{K}^{(e)}_{\leq d'}$.

Note that the $\mathbb{F}$-isomorphism $\mathbb{K}^{(e)} \to \mathbb{K}$ sending Y_i^{1/p^e} to Y_i induces a bijection

$$\mathbb{K}^{(e)}_{\leq d'}[\mathbf{X}]_{\leq d} \to \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}.$$

In this way, statements about polynomials in $\mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$ apply to those in $\mathbb{K}^{(e)}_{\leq d'}[\mathbf{X}]_{\leq d}$ via this transformation. On the other hand, there is also a natural inclusion

$$\mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d} \hookrightarrow \mathbb{K}^{(e)}_{\leq p^e d'}[\mathbf{X}]_{\leq d}$$

given by the identity map $f \mapsto f$.

Solving a system of linear equations. The following lemma bounds the complexity of a solution of a system of linear equations.

Lemma 3.5. *Let* $A = (a_{i,j})_{i \in [N], j \in [M]} \in \mathbb{K}^{N \times M}$ *and* $\mathbf{b} = (b_1, \ldots, b_N) \in \mathbb{K}^N$ *such that all entries of* A *and those of* $\mathbf{b}$ *are in* $\mathbb{K}_{\leq d}$. *Suppose the system of linear equations* $A\mathbf{x} = \mathbf{b}$ *has a solution. Then it has a solution* $\mathbf{x} \in \mathbb{K}^M$ *whose entries are all in* $\mathbb{K}_{\leq d'}$, *where* $d' = O_{N,M,d}(1)$.

Proof. Suppose A has rank k, and without loss of generality, we may assume the top-leftmost $k \times k$ minor B of A has full rank. Let $\widetilde{A} = (a_{i,j})_{i\in[k],j\in[M]}$ and $\widetilde{\mathbf{b}} = (b_1, \ldots, b_k)$. As $A\mathbf{x} = \mathbf{b}$ has a solution and the first k rows of A span the others, we know $\mathbf{x} = (x_1, \ldots, x_M)$ is a solution of $A\mathbf{x} = \mathbf{b}$ if and only if it is a solution of $\widetilde{A}\mathbf{x} = \widetilde{\mathbf{b}}$. By Cramer's rule, the latter has a solution given by

$$x_i = \begin{cases} \frac{\det(B_i)}{\det(B)} & 1 \le i \le k \\ 0 & k < i \le M, \end{cases} \tag{2}$$

where B_i is the matrix formed by replacing the i-th column of B by $\widetilde{\mathbf{b}}$. As the entries of B and B_i are in $\mathbb{K}_{\le d}$, by Lemma 3.2, the entries x_i given by (2) are in $\mathbb{K}_{\le d'}$ for some $d' = O_{N,M,d}(1)$. □

The following lemma bounds the complexity of the coefficients of a factor of a polynomial f in terms of its degree and the complexity of its coefficients.

Lemma 3.6. *Let $f \in \mathbb{K}_{\le d'}[\mathbf{X}]_{\le d}$ be a nonzero polynomial in n variables. Let g be a factor of f. Then there exists $c \in \mathbb{K}^\times$ such that $cg \in \mathbb{K}_{\le d''}[\mathbf{X}]_{\le d}$ for some $d'' = O_{d,d',n}(1)$.*

The proof of Lemma 3.6 is deferred to Appendix A.

Gröbner bases. For $\mathbf{X} = \{X_1, \ldots, X_n\}$, denote by $\mathcal{M}_{\mathbf{X}}$ the set of monomials in $\mathbf{X}$.

Let $\preceq$ be a total order on $\mathcal{M}_{\mathbf{X}}$. We say $\preceq$ is a *monomial order* if (1) $m_1 \preceq m_2 \implies m_1 m \preceq m_2 m$, and (2) $1 \preceq m$ for all $m \in \mathcal{M}_{\mathbf{X}}$. Write $m_1 \prec m_2$ if $m_1 \preceq m_2$ and $m_1 \neq m_2$.

A monomial order is *degree-compatible* if $m_1 \preceq m_2 \implies \deg(m_1) \le \deg(m_2)$. Examples of degree-compatible monomial orders include graded lexicographic and graded reverse lexicographic orders.

Fix a monomial order $\preceq$, for $0 \neq f \in \mathbb{K}[\mathbf{X}]$, the *leading monomial* of f, denoted by $\mathrm{LM}(f)$, is the monomial appearing in f that is greatest under $\preceq$. Its coefficient is called the *leading coefficient* of f and denoted $\mathrm{LC}(f)$. The term $\mathrm{LC}(f) \cdot \mathrm{LM}(f)$ is called the *leading term* of f and denoted $\mathrm{LT}(f)$. Define $\mathrm{LM}(0) = \mathrm{LC}(0) = \mathrm{LT}(0) = 0$.

For a set $S \subseteq \mathbb{K}[\mathbf{X}]$, denote by $\mathrm{LT}(S)$ the ideal of $\mathbb{K}[\mathbf{X}]$ generated by the set $\{\mathrm{LT}(f) : f \in S\}$. It is called *ideal of leading terms* of S, also known as the *initial ideal* of S.

Definition 3.7 (Gröbner basis)**.** *A finite generating set G of an ideal I of $\mathbb{K}[\mathbf{X}]$ is said to be a Gröbner basis for I if $\mathrm{LT}(G) = \mathrm{LT}(I)$.*

Degree bounds for Gröbner bases. We have the following degree bound on Gröbner bases.

Theorem 3.8 ([Dub90, Corollary 8.3])**.** *Let $\mathbb{K} = \mathbb{F}(\mathbf{Y})$. Let I be an ideal of $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ generated by polynomials of degree at most d. Then for any monomial order $\preceq$, I has a Gröbner basis with respect to $\preceq$ consisting only of polynomials of degree at most $2\left(\frac{d^2}{2} + d\right)^{2^{n-1}}$.*

We also need the following degree bound for the ideal membership problem. For a proof, see [MM82, Appendix].

Lemma 3.9 ([Her26, MM82])**.** *Let $f_1, \ldots, f_k \in \mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ be polynomials of degree at most d. Let $f \in \langle f_1, \ldots, f_k\rangle$ be a polynomial of degree at most d'. Then $f = \sum_{i=1}^k h_i f_i$ for some $h_1, \ldots, h_k \in \mathbb{K}[\mathbf{X}]$ of degree at most $d' + (kd)^{2^n}$.*

Now, we strengthen Theorem 3.8 and Lemma 3.9 to give the following lemma. In the case where the coefficients of $f_1, ..., f_k$ come from $\mathbb{K} = \mathbb{F}(\mathbf{Y})$, we establish complexity bounds on the Gröbner basis of an ideal $\langle f_1, ..., f_k \rangle$ by also bounding the degree of the coefficients of such a Gröbner basis.

Lemma 3.10. *Suppose $\mathbb{K} = \mathbb{F}(\mathbf{Y})$. Let I be an ideal of $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ generated by $f_1, \ldots, f_k \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$. Then for any monomial order $\preceq$, I has a Gröbner basis with respect to $\preceq$ consisting only of polynomials in $\mathbb{K}_{\leq d''}[\mathbf{X}]_{\leq D}$, where $D = 2\left(\frac{d^2}{2} + d\right)^{2^{n-1}}$ and $d'' = O_{d,d',n}(1)$.*

Proof. Note that we may assume $k \leq \binom{n+d}{d} = O_{d,n}(1)$. By Theorem 3.8, I has a Gröbner basis G consisting of polynomials of degree at most D. Consider nonzero $g \in G$ and let $m_g = \mathrm{LM}(g)$. By replacing g with $\frac{1}{\mathrm{LC}(g)}g$, we may assume $\mathrm{LT}(g) = m_g$. By Lemma 3.9, we have $g = \sum_{i=1}^k h_i f_i$ for some $h_1, \ldots, h_k$ of degree at most $D' := D + (kd)^{2^n} = O_{d,n}(1)$.

For every integer $a \in \mathbb{N}$, let S_a be the set of monomials of degree at most a. For each monomial m', as $g = \sum_{i=1}^k h_i f_i$, we have

$$\mathrm{coeff}_g(m') = \sum_{i=1}^{k} \sum_{m \in S_{D'} : m | m'} \mathrm{coeff}_{h_i}(m) \mathrm{coeff}_{f_i}(m'/m). \tag{3}$$

The condition that $\mathrm{LT}(g) = m_g$ is equivalent to that $\mathrm{coeff}_g(m_g) = 1$ and $\mathrm{coeff}_g(m') = 0$ for all $m' \in S_{d+D'}$ strictly greater than m_g with respect to $\preceq$. (Here, we only need to consider monomials $m' \in S_{d+D'}$ since $\deg(f_i) \leq d$ and $\deg(h_i) \leq D'$ for $i \in [k]$.) We formulate this condition as a system of linear equations as follows. Let S be the set of $m' \in S_{d+D'}$ satisfying either $m_g \preceq m'$ or $\deg(m') > \deg(g)$. Let $T = [k] \times S_{D'}$. Define a matrix $A \in \mathbb{K}^{S \times T}$ by

$$A[m', (i, m)] = \begin{cases} \mathrm{coeff}_{f_i}(m'/m) & m | m' \\ 0 & \text{otherwise.} \end{cases}$$

And define the column vector $\mathbf{b} \in \mathbb{K}^S$ by

$$\mathbf{b}[m'] = \begin{cases} 1 & m' = m_g \\ 0 & \text{otherwise.} \end{cases}$$

By assumption, we have $A[m', (i, m)], \mathbf{b}[m'] \in \mathbb{K}_{\leq d'}$ for all $(m', (i, m)) \in S \times T$. By Lemma 3.5, the system of linear equations $A\mathbf{x} = \mathbf{b}$ has a solution $(c_{i,m})_{(i,m) \in T}$ whose entries are all in $\mathbb{K}_{\leq d_0}$ for some

$$d_0 = O_{|S|,|T|,d'}(1) = O_{d,d',n}(1),$$

where the last equality holds since $|S| \leq |S_{d+D'}| = \binom{n+d+D'}{d+D'} = O_{d,n}(1)$ and $|T| = k|S_{D'}| = k\binom{n+D'}{D'} = O_{d,n}(1)$. Let $\widetilde{h}_i := \sum_{m \in S_{D'}} c_{i,m} m$ for $i \in [k]$. As $A\mathbf{x} = \mathbf{b}$, we have $\mathrm{LT}(\sum_{i=1}^k \widetilde{h}_i f_i) = m_g$ and $\deg(\sum_{i=1}^k \widetilde{h}_i f_i) \leq \deg(g)$, where the second condition holds since S includes all monomials in $S_{d+D'}$ of degree greater than $\deg(g)$, and their coefficients are therefore forced to be zero.

By Lemma 3.2, the coefficients of $\sum_{i=1}^k \widetilde{h}_i f_i$ belong to $\mathbb{K}_{\leq d''}$ for some $d'' = O_{d_0,d,d',k,n}(1) = O_{d,d',n}(1)$. Replacing h_i by $\widetilde{h}_i$ for $i \in [k]$, and replacing g by $\sum_{i=1}^k \widetilde{h}_i f_i$, do not change the leading term of g or increase its degree. Performing this replacement for each $g \in G$ preserves the fact that G is a Gröbner basis for I, since $\mathrm{LT}(G)$ remains unchanged. After the replacement, the new Gröbner basis satisfies the requirement of the lemma. □

Similarly, we establish degree bounds on the coefficients of an instance of the ideal membership problem.

Lemma 3.11. *Let $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ and $f_1, \ldots, f_k \in \mathbb{K}_{\leq D}[\mathbf{X}]_{\leq d}$. Let $f \in \langle f_1, \ldots, f_k \rangle$ be a polynomial in $\mathbb{K}_{\leq D'}[\mathbf{X}]_{\leq d'}$. Then $f = \sum_{i=1}^{k} h_i f_i$ for some $h_1, \ldots, h_k \in \mathbb{K}_{\leq D''}[\mathbf{X}]_{\leq d''}$ where $d'' := d' + (kd)^{2^n}$ and $D'' = O_{n,d,d',D,D'}(1)$.*

Proof. We may assume $k \leq \binom{n+d}{d}$. By Lemma 3.9, there exist polynomials $h_1, \ldots, h_k$ of degree at most d'' such that $f = \sum_{i=1}^{k} h_i f_i$. Viewing the coefficients of $h_1, \ldots, h_k$ as variables, we may build a system of linear equations over $\mathbb{K}$ that expresses the relation $f = \sum_{i=1}^{k} h_i f_i$. It has at most $k\binom{n+d''}{d''} = O_{n,d,d'}(1)$ variables and at most $\binom{n+d+d''}{d+d''} = O_{n,d,d'}(1)$ linear equations, obtained by comparing all coefficients through degree $d + d''$. The coefficients of these linear equations live in $\mathbb{K}_{\leq \max\{D,D'\}}$. So by Lemma 3.5, we may choose $h_1, \ldots, h_k$ such that they live in $\mathbb{K}_{\leq D''}[\mathbf{X}]_{\leq d''}$ for some $D'' = O_{n,d,d',D,D'}(1)$. □

In the following, we introduce the *elimination order* on monomials in a polynomial ring. We use it to obtain degree bounds on the Gröbner basis of the intersections of ideals.

Elimination of variables. Let $\preceq_{\mathbf{X}}$ and $\preceq_{\mathbf{Z}}$ be monomial orders on $\mathcal{M}_{\mathbf{X}}$ and $\mathcal{M}_{\mathbf{Z}}$, respectively. Denote by $\preceq = (\preceq_{\mathbf{X}}, \preceq_{\mathbf{Z}})$ the following monomial order on the monomials in both $\mathbf{X}$ and $\mathbf{Z}$:

$$m_1 m_2 \prec m_1' m_2' \iff \begin{cases} m_1 \prec m_1' \\ \text{or} \\ m_1 = m_1' \text{ and } m_2 \prec m_2' \end{cases} \qquad \text{for } m_1, m_1' \in \mathcal{M}_{\mathbf{X}} \text{ and } m_2, m_2' \in \mathcal{M}_{\mathbf{Z}}.$$

Call $\preceq$ an *elimination order* for $\mathbf{X}$.

The following is a well-known fact about eliminating variables using an elimination order. See, e.g., [AL94, Theorem 2.3.4].

Lemma 3.12 (Elimination of variables)**.** *Let $\preceq_{\mathbf{X}}$ and $\preceq_{\mathbf{Z}}$ be monomial orders on $\mathcal{M}_{\mathbf{X}}$ and $\mathcal{M}_{\mathbf{Z}}$, respectively. Let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{\mathbf{Z}})$. Let I be an ideal of $\mathbb{K}[\mathbf{X}, \mathbf{Z}]$. Let G be a Gröbner basis for the ideal I with respect to $\preceq$. Then $G \cap \mathbb{K}[\mathbf{Z}]$ is a Gröbner basis for the ideal $I \cap \mathbb{K}[\mathbf{Z}]$ of $\mathbb{K}[\mathbf{Z}]$ with respect to $\preceq_{\mathbf{Z}}$.*

We combine Lemma 3.10 and Lemma 3.12 to give the following statement, which provides degree bounds for the Gröbner basis of ideals of the form $I \cap \mathbb{K}[\mathbf{Z}]$, as well as that of its coefficients.

Lemma 3.13. *Let I be an ideal of $\mathbb{K}[\mathbf{X}, \mathbf{Z}] = \mathbb{K}[X_1, \ldots, X_n, Z_1, \ldots, Z_m]$ generated by polynomials $f_1, \ldots, f_k \in \mathbb{K}_{\leq d'}[\mathbf{X}, \mathbf{Z}]_{\leq d}$. Then for any monomial order $\preceq_{\mathbf{Z}}$ on $\mathcal{M}_{\mathbf{Z}}$, $I \cap \mathbb{K}[\mathbf{Z}]$ has a Gröbner basis with respect to $\preceq_{\mathbf{Z}}$ consisting only of polynomials in $\mathbb{K}_{\leq D'}[\mathbf{Z}]_{\leq D}$, where $D = O_{d,n+m}(1)$ and $D' = O_{d,d',n+m}(1)$.*

Proof. Let $\preceq_{\mathbf{X}}$ be a monomial order on $\mathcal{M}_{\mathbf{X}}$ and let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{\mathbf{Z}})$. Let G be a Gröbner basis for I with respect to $\preceq$. By Lemma 3.10, we may assume that G consists only of polynomials in $\mathbb{K}_{\leq D'}[\mathbf{X}, \mathbf{Z}]_{\leq D}$, where $D = 2\left(\frac{d^2}{2} + d\right)^{2^{n+m-1}} = O_{d,n+m}(1)$ and $D' = O_{d,d',n+m}(1)$. By Lemma 3.12, $G \cap \mathbb{K}[\mathbf{Z}]$ is a Gröbner basis for $I \cap \mathbb{K}[\mathbf{Z}]$ with respect to $\preceq_{\mathbf{Z}}$. The lemma follows. □

Reduction algorithm. The reduction algorithm, also known as the division algorithm, generalizes both row reduction in Gaussian elimination and long division of univariate polynomials. Running this algorithm on a polynomial f with respect to a Gröbner basis yields a unique "remainder" of f.

Definition 3.14. *Let G be a subset of $\mathbb{K}[\mathbf{X}]$. $f \in \mathbb{K}[\mathbf{X}]$ is said to be* reducible modulo G *if f has a nonzero term that is in $\mathrm{LT}(G)$. Otherwise, we say f is* reduced modulo G.

In the following lemma, we bound the complexity of a polynomial f modulo a Gröbner basis G. We again give bounds on the degree of the remainder f_G as well as that of f_G's coefficients.

Lemma 3.15. *Let $f \in \mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ and let $G = \{g_1, \ldots, g_k\}$ be a Gröbner basis for an ideal I of $\mathbb{K}[\mathbf{X}]$ with respect to a monomial order $\preceq$. Then:*

(1) There exists a unique remainder $f_G \in \mathbb{K}[\mathbf{X}]$ such that f_G is reduced mod G and $f - f_G = \sum_{i=1}^{k} h_i g_i \in I$, where $h_1, \ldots, h_k \in \mathbb{K}[\mathbf{X}]$.

(2) Suppose $f, g_1, \ldots, g_k \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$. Further assume that for $i \in [k]$, $\deg(\mathrm{LT}(g_i)) = \deg(g_i)$, which holds if $\preceq$ is degree-compatible. Then in Item 1, $f_G \in \mathbb{K}_{\leq d''}[\mathbf{X}]_{\leq d}$ and $h_1, \ldots, h_k$ may be chosen in $\mathbb{K}_{\leq d''}[\mathbf{X}]_{\leq d}$ as well, where $d'' = O_{d,d',n}(1)$.

The reduction algorithm outputting f_G and $h_1, \ldots, h_k$ is given as follows.

Algorithm 1 The reduction algorithm

Input: f, $G = \{g_1, \ldots, g_k\}$
Output: f_G, $h_1, \ldots, h_k$

```
f_G ← f; h_1, ..., h_k ← 0
while f_G has a term T divisible by LT(g_i) for some i ∈ [k] do
    f_G ← f_G − (T / LT(g_i)) g_i
    h_i ← h_i + T / LT(g_i)
end while
return f_G, h_1, ..., h_k
```

Algorithm 1 terminates in general, even without the assumption that $\deg(\mathrm{LT}(g_i)) = \deg(g_i)$ for $i \in [k]$. This fact follows from Hilbert's basis theorem or Dickson's lemma [BW98]. We make the additional assumption that $\deg(\mathrm{LT}(g_i)) = \deg(g_i)$ for $i \in [k]$, which leads to a simpler proof that bounds the degrees of the coefficients of $f_G, h_1, \ldots, h_k$.

Proof of Lemma 3.15. Item 1 is standard (see, e.g., [AL94]).

We claim $\deg(f_G) \leq \deg(f)$ and $\deg(h_i g_i) \leq \deg(f)$ for $i \in [k]$. This follows from the following induction: At the beginning of the algorithm, these bounds hold since $f_G = f$ and $h_1, \ldots, h_k = 0$. In each iteration, the degree of $\Delta := \frac{T}{\mathrm{LT}(g_i)} g_i$ is $\deg(T) - \deg(\mathrm{LT}(g_i)) + \deg(g_i) = \deg(T) \leq \deg(f_G) \leq \deg(f)$. We subtract Δ from f_G, which does not increase the degree of f_G since the degrees of the newly added monomials are bounded by $\deg(\Delta) \leq \deg(f_G)$. Similarly, we add Δ/g_i to h_i, which preserves the property that $\deg(h_i g_i) \leq \deg(f)$ since $\deg((\Delta/g_i) \cdot g_i) = \deg(\Delta) \leq \deg(f)$. So the degree bounds hold during and at the end of the algorithm.

Now we prove Item 2. By the degree bounds just proved, we have $\deg(f_G), \deg(h_1), \ldots, \deg(h_k) \leq \deg(f) \leq d$.

We have

$$f_G = f - \sum_{i=1}^{k} h_i g_i = f - \sum_{i=1}^{k} \sum_{m \in S_i} \mathrm{coeff}_{h_i}(m) m g_i,$$

where S_i is the set of monomials m such that $\deg(m \cdot g_i) \le \deg(f)$. Therefore, for each monomial m' of degree at most $\deg(f)$, we have

$$\mathrm{coeff}_{f_G}(m') = \mathrm{coeff}_f(m') - \sum_{i=1}^{k} \sum_{m \in S_i : m \mid m'} \mathrm{coeff}_{h_i}(m) \mathrm{coeff}_{g_i}(m'/m). \tag{4}$$

By keeping only one polynomial for each distinct leading monomial in G, we may assume that $k \le \binom{n+d}{d}$. Let S be the set of monomials m' of degree at most $\deg(f)$ such that $m' \in \mathrm{LT}(G)$, and let $T = \{(i, m) : i \in [k],\ m \in S_i\}$. Note that $|S| \le \binom{n+d}{d}$ and $|T| = \sum_{i=1}^{k} |S_i| \le \binom{n+d}{d}^2$. By definition, a polynomial g satisfying $\deg(g) \le \deg(f)$ is reduced mod G if and only if $\mathrm{coeff}_g(m') = 0$ for all $m' \in S$. So the polynomial $f_G = f - \sum_{i=1}^{k} h_i g_i$ with $h_i = \sum_{m \in S_i} \mathrm{coeff}_{h_i}(m) m$ is reduced mod G if and only if the column vectors $\mathbf{x} = (\mathrm{coeff}_{h_i}(m))_{(i,m) \in T}$ and $\mathbf{b} = (\mathrm{coeff}_f(m'))_{m' \in S}$ satisfy $A\mathbf{x} = \mathbf{b}$, where the matrix $A \in \mathbb{K}^{S \times T}$ is given by

$$A[m', (i, m)] = \begin{cases} \mathrm{coeff}_{g_i}(m'/m) & m \mid m' \\ 0 & \text{otherwise.} \end{cases}$$

By assumption, we have $A[m', (i, m)], \mathbf{b}[m'] \in \mathbb{K}_{\le d'}$ for all $(m', (i, m)) \in S \times T$. By Lemma 3.5, we may choose the coefficients of $h_1, \ldots, h_k$ to be in $\mathbb{K}_{\le d_0}$ for some $d_0 = O_{|S|,|T|,d'}(1) = O_{d,d',n}(1)$.

By (4) and Lemma 3.2, for such $h_1, \ldots, h_k$, the coefficients of f_G are all in $\mathbb{K}_{\le d_1}$, where $d_1 = O_{d,d',n}(1)$. Choose $d'' := \max(d_0, d_1) = O_{d,d',n}(1)$. Then $f_G, h_1, \ldots, h_k \in \mathbb{K}_{\le d''}[\mathbf{X}]_{\le d}$ by definition. □

Definition 3.16. *f_G in Lemma 3.15 is called the* remainder *of f modulo G with respect to $\preceq$.*

The following is a variant of Lemma 3.15 for elimination orders. It states that the degree of the remainder f_G and the degrees of its coefficients are bounded.

Lemma 3.17. *Let $f \in \mathbb{K}[\mathbf{X}, \mathbf{Z}] = \mathbb{K}[X_1, \ldots, X_n, Z_1, \ldots, Z_m]$. Let $\preceq_{\mathbf{X}}$ and $\preceq_{\mathbf{Z}}$ be degree-compatible monomial orders on $\mathcal{M}_{\mathbf{X}}$ and on $\mathcal{M}_{\mathbf{Z}}$, respectively, and let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{\mathbf{Z}})$. Let $G = \{g_1, \ldots, g_k\}$ be a Gröbner basis for an ideal I of $\mathbb{K}[\mathbf{X}, \mathbf{Z}]$ with respect to $\preceq$. Suppose $f, g_1, \ldots, g_k \in \mathbb{K}_{\le d'}[\mathbf{X}, \mathbf{Z}]_{\le d}$. Then the remainder f_G of f is in $\mathbb{K}_{\le d_2}[\mathbf{X}, \mathbf{Z}]_{\le d_1}$, where $d_1 = O_{d,n+m}(1)$ and $d_2 = O_{d,d',n+m}(1)$.*

Proof. Let $N = \binom{n+d}{d}$ and $B = d(N+1)$. Run the reduction algorithm with respect to G.

Whenever a term is reduced, the $\mathbf{X}$-part of every newly introduced monomial is no larger, with respect to $\preceq_{\mathbf{X}}$, than the $\mathbf{X}$-part of the term being reduced. Since $\preceq_{\mathbf{X}}$ is degree-compatible, the degree in $\mathbf{X}$ never exceeds d. Thus at most N different $\mathbf{X}$-monomials can occur.

If a reduction leaves the $\mathbf{X}$-part unchanged, then the degree in $\mathbf{Z}$ does not increase, because $\preceq_{\mathbf{Z}}$ is degree-compatible. If the $\mathbf{X}$-part strictly decreases, the degree in $\mathbf{Z}$ increases by at most d. Along any chain of terms produced by successive reductions, the latter event occurs at most $N-1$ times. Consequently, every intermediate monomial has degree in $\mathbf{Z}$ at most dN and total degree at most B. The multipliers h_i produced by the algorithm also have degree at most B.

For the coefficient bound, we first replace G by a subset of G with the same leading-term ideal and with at most $\binom{n+m+d}{d}$ elements. This does not change the remainder. Introduce unknown

coefficients for the polynomials h_i of degree at most B and for a remainder supported on standard monomials of degree at most B. The identity

$$f = f_G + \sum_i h_i g_i$$

is a consistent linear system, obtained by comparing coefficients through degree $B + d$. Its size is $O_{d,n+m}(1)$, and its entries have complexity bounded in terms of d'. By Lemma 3.5, it has a solution with coefficient complexity $O_{d,d',n+m}(1)$. Uniqueness of the reduced remainder identifies the resulting polynomial with f_G. □

Ring homomorphisms. The following lemma gives quantitative, effective degree bounds on ring homomorphisms and the induced isomorphisms between quotient rings. While Part (1) is the standard First Isomorphism Theorem, Part (2) bounds the complexity of the Gröbner basis of the kernel of a homomorphism, and Part (3) shows that if an element in the image is represented by a polynomial with bounded complexity, then so is its preimage.

Lemma 3.18. *Let $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ and $\mathbb{K}[\mathbf{Z}] = \mathbb{K}[Z_1, \ldots, Z_m]$. Let $f_1, \ldots, f_k, g_1, \ldots, g_m \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$. Let $I = \langle f_1, \ldots, f_k \rangle \subseteq \mathbb{K}[\mathbf{X}]$. Let ϕ be the $\mathbb{K}$-linear ring homomorphism*

$$\begin{aligned} \phi : \mathbb{K}[\mathbf{Z}] &\to \mathbb{K}[\mathbf{X}]/I \\ Z_i &\mapsto g_i + I, \quad i = 1, 2, \ldots, m. \end{aligned}$$

Let A be the image of ϕ, i.e., A is the subring of $\mathbb{K}[\mathbf{X}]/I$ generated by $g_1 + I, \ldots, g_m + I$. Then:

(1) ϕ induces an isomorphism $\mathbb{K}[Z_1, \ldots, Z_m]/\ker(\phi) \cong A$.

(2) For any monomial order $\preceq_{\mathbf{Z}}$ on $\mathcal{M}_{\mathbf{Z}}$, $\ker(\phi)$ has a Gröbner basis with respect to $\preceq_{\mathbf{Z}}$ consisting only of polynomials in $\mathbb{K}_{\leq D'}[\mathbf{Z}]_{\leq D}$, where $D = O_{d,n+m}(1)$ and $D' = O_{d,d',n+m}(1)$.

(3) Let $f \in \mathbb{K}[\mathbf{X}]$ such that $f + I \in A$ and $f \in \mathbb{K}_{\leq d_2}[\mathbf{X}]_{\leq d_1}$. Then there exists $h \in \mathbb{K}[\mathbf{Z}]$ such that $\phi(h) = f + I$ and $h \in \mathbb{K}_{\leq d_4}[\mathbf{Z}]_{\leq d_3}$, where $d_3 = O_{d,d_1,n+m}(1)$ and $d_4 = O_{d,d',d_1,d_2,n+m}(1)$.

Proof. Item 1 holds by the first isomorphism theorem. Let $J = \langle f_1, \ldots, f_k, Z_1 - g_1, \ldots, Z_m - g_m \rangle \subseteq \mathbb{K}[\mathbf{X}, \mathbf{Z}]$. Then $\ker(\phi) = J \cap \mathbb{K}[\mathbf{Z}]$ [AL94, Theorem 2.4.10]. Item 2 then follows from Lemma 3.13.

It remains to prove Item 3. Let $\preceq_{\mathbf{X}}$ and $\preceq_{\mathbf{Z}}$ be degree-compatible monomial orders on $\mathcal{M}_{\mathbf{X}}$ and on $\mathcal{M}_{\mathbf{Z}}$, respectively. Let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{\mathbf{Z}})$. Let G be a Gröbner basis of J with respect to $\preceq$. By Lemma 3.10, we may assume that G consists only of polynomials in $\mathbb{K}_{\leq d_6}[\mathbf{X}, \mathbf{Z}]_{\leq d_5}$, where $d_5 = O_{d,n+m}(1)$ and $d_6 = O_{d,d',n+m}(1)$.

By assumption, $f + I$ is in A and hence in the image of ϕ. Let h be the remainder of f modulo G with respect to $\preceq$. Then $h \in \mathbb{K}[\mathbf{Z}]$ and $\phi(h) = f + I$ by [AL94, Theorem 2.4.11]. Apply Lemma 3.17 with degree parameter $\max\{d_1, d_5\}$ and coefficient-complexity parameter $\max\{d_2, d_6\}$. It follows that $h \in \mathbb{K}_{\leq d_4}[\mathbf{Z}]_{\leq d_3}$, where $d_3 = O_{d,d_1,n+m}(1)$ and $d_4 = O_{d,d',d_1,d_2,n+m}(1)$. □

Primitive element theorem. We now prove a quantitative form of the primitive element theorem, modifying a proof in [ZS75].

Theorem 3.19 (Primitive element theorem, quantitative form)**.** *Let* $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \dots, X_n]$. *Suppose* $f_1, f_2, \dots, f_k \in \mathbb{K}_{\le d'}[\mathbf{X}]_{\le d}$ *generate a zero-dimensional ideal* I. *Suppose* $\mathfrak{m}$ *is a maximal ideal of* $\mathbb{K}[\mathbf{X}]$ *containing* I *and* $\mathbb{L} := \mathbb{K}[\mathbf{X}]/\mathfrak{m}$ *is a finite separable extension of* $\mathbb{K}$. *Let* $\alpha_i := X_i + \mathfrak{m} \in \mathbb{L}$ *for* $i \in [n]$, *so that* $\mathbb{L} = \mathbb{K}(\alpha_1, \dots, \alpha_n)$. *Then there exists a nonzero polynomial* $Q \in \mathbb{L}[\mathbf{Z}] = \mathbb{L}[Z_1, \dots, Z_n]$ *such that for every nonzero* $\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{F}^n$ *satisfying* $Q(\mathbf{c}) \neq 0$, *there exist* $P_0^{\mathbf{c}}, P_1^{\mathbf{c}}, \dots, P_n^{\mathbf{c}} \in \mathbb{K}[T]$ *such that the following hold:*

(1) Let $\beta_{\mathbf{c}} = \sum_{i=1}^n c_i \alpha_i \in \mathbb{L}$. *Then* $P_0^{\mathbf{c}}(\beta_{\mathbf{c}}) \neq 0$ *and* $\alpha_i = \frac{P_i^{\mathbf{c}}(\beta_{\mathbf{c}})}{P_0^{\mathbf{c}}(\beta_{\mathbf{c}})}$ *for* $i \in [n]$. *In particular,* $\mathbb{L} = \mathbb{K}(\beta_{\mathbf{c}})$.

(2) $P_0^{\mathbf{c}}, \dots, P_n^{\mathbf{c}} \in \mathbb{K}_{\le d_0}[T]_{\le d^n}$ *for some* $d_0 = O_{d,d',n}(1)$.

In particular, the above hold for almost all $\mathbf{c} \in \mathbb{F}^n$.

Proof. Let $\widetilde{\mathbb{K}} = \mathbb{K}(\mathbf{Z}) = \mathbb{K}(Z_1, \dots, Z_n)$ and $\widetilde{\mathbb{L}} = \mathbb{L}(\mathbf{Z}) = \mathbb{L}(Z_1, \dots, Z_n) = \widetilde{\mathbb{K}}(\alpha_1, \dots, \alpha_n)$. By assumption, $\alpha_1, \dots, \alpha_n$ are separable over $\mathbb{K}$. So they are also separable over $\widetilde{\mathbb{K}}$. It follows that $\widetilde{\mathbb{L}}$ is a finite separable extension of $\widetilde{\mathbb{K}}$.

Let $\widehat{\alpha_i} = X_i + I \in \mathbb{K}[\mathbf{X}]/I$ for $i \in [n]$. Let $\widehat{\beta}(\mathbf{Z}) = Z_1\widehat{\alpha_1} + Z_2\widehat{\alpha_2} + \cdots + Z_n\widehat{\alpha_n} \in (\mathbb{K}[\mathbf{X}]/I) \otimes_{\mathbb{K}} \widetilde{\mathbb{K}}$. We have $\dim_{\mathbb{K}}(\mathbb{K}[\mathbf{X}]/I) \le d^n$ by Lemma 2.14. So $\dim_{\widetilde{\mathbb{K}}}((\mathbb{K}[\mathbf{X}]/I) \otimes_{\mathbb{K}} \widetilde{\mathbb{K}}) \le d^n$. Therefore, $\widehat{\beta}(\mathbf{Z})$ has a minimal polynomial over $\widetilde{\mathbb{K}}$, which we denote by $\widehat{F}$, and its degree is at most d^n.

We may construct $\widehat{F}$ as follows. Let $\preceq_{T,\mathbf{Z}}$ be an elimination order for T on $\mathcal{M}_{T,\mathbf{Z}}$, and let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{T,\mathbf{Z}})$ be an elimination order for $\mathbf{X}$ on $\mathcal{M}_{\mathbf{X},T,\mathbf{Z}}$. Let J be the ideal

$$J = \langle f_1, \dots, f_k, T - (Z_1X_1 + \cdots + Z_nX_n) \rangle$$

of $\mathbb{K}[\mathbf{X}, T, \mathbf{Z}]$. The ideal J is the preimage of the ideal $\left\langle T - \widehat{\beta}(\mathbf{Z}) \right\rangle$ of $(\mathbb{K}[\mathbf{X}]/I)[T, \mathbf{Z}]$ under the natural quotient map $\mathbb{K}[\mathbf{X}, T, \mathbf{Z}] \to (\mathbb{K}[\mathbf{X}]/I)[T, \mathbf{Z}]$. Let G be a Gröbner basis for J with respect to $\preceq$. Then $G \cap \mathbb{K}[T, \mathbf{Z}]$ is a Gröbner basis for $J \cap \mathbb{K}[T, \mathbf{Z}]$ with respect to $\preceq_{T,\mathbf{Z}}$ by Lemma 3.12. Choose $\widehat{g} \in G \cap \mathbb{K}[T, \mathbf{Z}]$ such that $\mathrm{LM}(\widehat{g})$ has the form mT^e, where $m \in \mathcal{M}_{\mathbf{Z}}$ and e is minimized. By the choice of $\preceq$, we have $\widehat{g} \in \mathbb{K}[T, \mathbf{Z}]$. Note that such $\widehat{g}$ exists since it can be obtained by clearing the denominators of the coefficients of $\widehat{F}(T)$. By Lemma 3.10, we may assume $\widehat{g} \in \mathbb{K}_{\le d_1}[T, \mathbf{Z}]_{\le D}$ with $D = O_{d,n}(1)$ and $d_1 = O_{d,d',n}(1)$.

Write $\widehat{g} = \sum_{i=0}^e h_i T^i$ with $h_i \in \mathbb{K}[\mathbf{Z}]$ for $i = 0, 1, \dots, e$. Then $h_e \neq 0$. As e is minimized and J is the preimage of $\left\langle T - \widehat{\beta}(\mathbf{Z}) \right\rangle$ in $\mathbb{K}[\mathbf{X}, T, \mathbf{Z}]$, $\widehat{g}/h_e = \sum_{i=0}^e (h_i/h_e) T^i$ is precisely the minimal polynomial $\widehat{F}(T)$ of $\widehat{\beta}(\mathbf{Z})$ over $\widetilde{\mathbb{K}}$.

Consider

$$\beta(\mathbf{Z}) := Z_1\alpha_1 + Z_2\alpha_2 + \cdots + Z_n\alpha_n \in \mathbb{L}[\mathbf{Z}] \subseteq \widetilde{\mathbb{L}}.$$

Let $F(T) \in \widetilde{\mathbb{K}}[T]$ be the minimal polynomial of $\beta(\mathbf{Z})$ over $\widetilde{\mathbb{K}}$. As $\mathfrak{m} \supseteq I$, F is a factor of $\widehat{F}$. In particular, $\deg(F) \le \deg(\widehat{F}) \le d^n$. As $\widehat{g}/\widehat{F} = h_e \in \widetilde{\mathbb{K}}^\times$, F is a factor of $\widehat{g}$, both viewed as polynomials over $\widetilde{\mathbb{K}}$. So $\widehat{g}$ has a factor $g \in \mathbb{K}[T, \mathbf{Z}]$ such that $g/F \in \widetilde{\mathbb{K}}^\times$. Moreover, by Lemma 3.6, after replacing g with cg for a suitable $c \in \mathbb{K}^\times$, we may assume that $g \in \mathbb{K}_{\le d_2}[T, \mathbf{Z}]_{\le D}$, where $d_2 = O_{D,d_1,n}(1) = O_{d,d',n}(1)$.

As $g/F \in \widetilde{\mathbb{K}}^\times$ and $F = F(T)$ is the minimal polynomial of $\beta(\mathbf{Z})$ over $\widetilde{\mathbb{K}}$, we have

$$g(Z_1\alpha_1 + Z_2\alpha_2 + \cdots + Z_n\alpha_n, Z_1, \dots, Z_n) = g(\beta(\mathbf{Z}), \mathbf{Z}) = 0. \tag{5}$$

Taking the partial derivatives of (5) with respect to $Z_1, \ldots, Z_n$ and applying the chain rule for multivariate polynomials, we obtain

$$\alpha_i g_0(\beta(\mathbf{Z}), \mathbf{Z}) + g_i(\beta(\mathbf{Z}), \mathbf{Z}) = 0, \quad i = 1, 2, \ldots, n, \tag{6}$$

where $g_0(T, \mathbf{Z}) := \frac{\partial g(T,\mathbf{Z})}{\partial T}$ and $g_i(T, \mathbf{Z}) := \frac{\partial g(T,\mathbf{Z})}{\partial Z_i}$ for $i \in [n]$.

As $F = F(T)$ is the minimal polynomial of $\beta(\mathbf{Z}) \in \widetilde{\mathbb{L}}$ over $\widetilde{\mathbb{K}}$ and $\widetilde{\mathbb{L}}$ is a finite separable extension of $\widetilde{\mathbb{K}}$, we have $F'(\beta(\mathbf{Z})) \neq 0$. And as $g_0(T, \mathbf{Z}) = \frac{\partial g(T,\mathbf{Z})}{\partial T}$ and $g/F \in \widetilde{\mathbb{K}}^\times$, we have

$$g_0(\beta(\mathbf{Z}), \mathbf{Z}) = (g/F) \cdot F'(\beta(\mathbf{Z})) \neq 0.$$

Define $Q(\mathbf{Z}) = g_0(\beta(\mathbf{Z}), \mathbf{Z}) \in \mathbb{L}[\mathbf{Z}]$. For every $\mathbf{c} = (c_1, \ldots, c_n) \in \mathbb{F}^n$, define $P_0^{\mathbf{c}}(T) = g_0(T, \mathbf{c}) \in \mathbb{K}[T]$ and $P_i^{\mathbf{c}}(T) = -g_i(T, \mathbf{c}) \in \mathbb{K}[T]$ for $i \in [n]$.

Consider $\mathbf{c} \in \mathbb{F}^n$ such that $Q(\mathbf{c}) \neq 0$. Note that $\beta(\mathbf{c}) = \sum_{i=1}^n c_i \alpha_i = \beta_{\mathbf{c}}$ and $P_0^{\mathbf{c}}(\beta_{\mathbf{c}}) = g_0(\beta(\mathbf{c}), \mathbf{c}) = Q(\mathbf{c}) \neq 0$. We have by (6) that

$$\alpha_i = -\frac{g_i(\beta(\mathbf{c}), \mathbf{c})}{g_0(\beta(\mathbf{c}), \mathbf{c})} = \frac{P_i^{\mathbf{c}}(\beta_{\mathbf{c}})}{P_0^{\mathbf{c}}(\beta_{\mathbf{c}})}, \quad i = 1, 2, \ldots, n,$$

By definition, $P_0^{\mathbf{c}}(T) = g_0(T, \mathbf{c})$ and $P_i^{\mathbf{c}}(T) = -g_i(T, \mathbf{c})$ for $i \in [n]$. So we have $\deg(P_i^{\mathbf{c}}) \leq \deg_T(g_i) \leq \deg_T(g) \leq e \leq d^n$ for $i = 0, 1, \ldots, n$. As already noted above, $g \in \mathbb{K}_{\leq d_2}[T, \mathbf{Z}]_{\leq D}$, so the coefficients of g are all in $\mathbb{K}_{\leq d_2}$. For $i = 0, 1, \ldots, n$ and $j = 0, 1, \ldots, \deg(P_i^{\mathbf{c}})$, the coefficient of T^j in $P_i^{\mathbf{c}}$ is a linear combination of the coefficients of the monomials of degree j in T of g_i over $\mathbb{F}$, and g_i has at most $\binom{n+D}{D}$ such monomials since $\deg(g_i) \leq \deg(g) \leq D$. Therefore, by Lemma 3.2, the coefficients of $P_0^{\mathbf{c}}, \ldots, P_n^{\mathbf{c}}$ are in $\mathbb{K}_{\leq d_0}$, where $d_0 = \binom{n+D}{D} d_2 = O_{d,d',n}(1)$. It follows that $P_0^{\mathbf{c}}, \ldots, P_n^{\mathbf{c}} \in \mathbb{K}_{\leq d_0}[T]_{\leq d^n}$. □

Extracting a maximal ideal. For I a zero-dimensional ideal of $\mathbb{K}[\mathbf{X}]$, we need to solve the problem of finding the generators of a maximal ideal $\mathfrak{m}$ containing I or, more precisely, to bound the complexity of these generators. This problem is a special case of finding the radical and the primary decomposition of a zero-dimensional ideal, with coefficient bounds. The following lemma provides a direct solution, providing the required bounds on the complexity of the generators of $\mathfrak{m}$.

Lemma 3.20. *Let $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ and $f_1, \ldots, f_k \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$. Suppose $I = \langle f_1, \ldots, f_k \rangle$ is a zero-dimensional ideal of $\mathbb{K}[\mathbf{X}]$. Let $\mathfrak{m}$ be a maximal ideal of $\mathbb{K}[\mathbf{X}]$ containing I such that $\mathbb{K}[\mathbf{X}]/\mathfrak{m}$ is a finite separable extension of $\mathbb{K}$. Then $\mathfrak{m}$ is generated by polynomials in $\mathbb{K}_{\leq d''}[\mathbf{X}]_{\leq d^n+1}$, where $d'' = O_{d,d',n}(1)$.*

Proof. We follow the approach of [GTZ88]. Let $\alpha_i = X_i + \mathfrak{m} \in \mathbb{K}[\mathbf{X}]/\mathfrak{m}$ for $i \in [n]$. By Theorem 3.19, there exist $\mathbf{c} = (c_1, \ldots, c_n) \in \mathbb{F}^n$ and $P_0^{\mathbf{c}}, P_1^{\mathbf{c}}, \ldots, P_n^{\mathbf{c}} \in \mathbb{K}_{\leq d_0}[T]_{\leq d^n}$ for some $d_0 = O_{d,d',n}(1)$ such that $\beta_{\mathbf{c}} = \sum_{i=1}^n c_i \alpha_i$ satisfies $P_0^{\mathbf{c}}(\beta_{\mathbf{c}}) \neq 0$ and $\alpha_i = \frac{P_i^{\mathbf{c}}(\beta_{\mathbf{c}})}{P_0^{\mathbf{c}}(\beta_{\mathbf{c}})}$ for $i \in [n]$.

Let $T = \sum_{i=1}^n c_i X_i \in \mathbb{K}[\mathbf{X}]$. Let $\mathfrak{m}_0 = \mathfrak{m} \cap \mathbb{K}[T]$, which is a prime ideal of $\mathbb{K}[T]$. Let $I_0 = I \cap \mathbb{K}[T]$. As $\mathbb{K}[T]$ is a PID, $\mathfrak{m}_0$ and I_0 are generated by polynomials $g, h \in \mathbb{K}[T]$ over $\mathbb{K}$, respectively, and $g | h$. The inclusion $\mathbb{K}[T] \hookrightarrow \mathbb{K}[\mathbf{X}]$ induces an inclusion $\mathbb{K}[T]/I_0 \hookrightarrow \mathbb{K}[\mathbf{X}]/I$. So we have

$$\deg(g) \leq \deg(h) = \dim_{\mathbb{K}}(\mathbb{K}[T]/I_0) \leq \dim_{\mathbb{K}}(\mathbb{K}[\mathbf{X}]/I) \leq d^n,$$

where the last inequality holds by Lemma 2.14. By Lemma 3.13 and the fact that $I_0 = \langle h \rangle$, we may assume that $h \in \mathbb{K}_{\le d_1}[T]_{\le d^n}$ for some $d_1 = O_{d,d',n}(1)$.[5] As $g|h$, by Lemma 3.6, we may assume that $g \in \mathbb{K}_{\le d_2}[T]_{\le d^n}$ for some $d_2 = O_{d,d',n}(1)$.

For $i \in [n]$, let $g_i = P_0^{\mathbf{c}}(T)X_i - P_i^{\mathbf{c}}(T) \in \mathbb{K}[\mathbf{X}]$, whose degree is at most $d^n + 1$. And the coefficients of each g_i are in $\mathbb{K}_{\le d_3}$ for some $d_3 = O_{d_0,d,n}(1) = O_{d,d',n}(1)$ since $P_0^{\mathbf{c}}, P_i^{\mathbf{c}} \in \mathbb{K}_{\le d_0}[T]_{\le d^n}$. Let $d'' = \max(d_2, d_3)$. Then $g, g_1, \ldots, g_n \in \mathbb{K}_{\le d''}[\mathbf{X}]_{\le d^n+1}$.

For $i \in [n]$, by the definition of α_i, the image of g_i in $\mathbb{K}[\mathbf{X}]/\mathfrak{m}$ is $P_0^{\mathbf{c}}(\beta_{\mathbf{c}})\alpha_i - P_i^{\mathbf{c}}(\beta_{\mathbf{c}})$, which is zero since $\alpha_i = \frac{P_i^{\mathbf{c}}(\beta_{\mathbf{c}})}{P_0^{\mathbf{c}}(\beta_{\mathbf{c}})}$. So $g_1, \ldots, g_n \in \mathfrak{m}$.

Note that $\mathbb{K}[\mathbf{X}]/\langle g, g_1, \ldots, g_n \rangle \cong \mathbb{K}[T]/\langle g \rangle = \mathbb{K}[T]/\mathfrak{m}_0$. This follows by noting that $P_0^{\mathbf{c}}(T) \in \mathbb{K}[T]$ is invertible modulo $\langle g \rangle = \mathfrak{m}_0$ (since $P_0^{\mathbf{c}}(\beta_{\mathbf{c}}) \neq 0$ and $\mathfrak{m}_0 = \mathfrak{m} \cap \mathbb{K}[T]$ is a maximal ideal, modulo which every nonzero element is invertible), and therefore, we can use the relations $g_i = P_0^{\mathbf{c}}(T)X_i - P_i^{\mathbf{c}}(T) = 0$, $i = 1, \ldots, n$, to eliminate $X_1, \ldots, X_n$. As $\mathbb{K}[T]/\mathfrak{m}_0$ is a field, $\langle g, g_1, \ldots, g_n \rangle$ is a maximal ideal of $\mathbb{K}[\mathbf{X}]$. As $g, g_1, \ldots, g_n \in \mathfrak{m}$ and $\mathfrak{m}$ is also a maximal ideal of $\mathbb{K}[\mathbf{X}]$, we have $\mathfrak{m} = \langle g, g_1, \ldots, g_n \rangle$. The lemma follows. □

Moving on, we introduce the notion of an *idealizer* in an integral domain, which will be used in the next section to compute the normalization of curves. We begin by introducing several related definitions and then prove a lemma that bounds the complexity of the preimage of an idealizer.

Idealizer. Let A be an integral domain and J be an ideal of A. The *idealizer* of J is defined as

$$\mathrm{Id}_A(J) := \{a \in \mathrm{Frac}(A) : aJ \subseteq J\}.$$

It is the largest subring of $\mathrm{Frac}(A)$ in which J is still an ideal.

A closely related notion is the *ideal quotient* of two ideals. For ideals I and J of a ring R, define the ideal quotient $(I : J) := \{a \in R : aJ \subseteq I\}$, which is an ideal of R.

Lemma 3.21. *Let $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ and $f_1, \ldots, f_k, g_1, \ldots, g_m \in \mathbb{K}_{\le d'}[\mathbf{X}]_{\le d}$. Suppose $I := \langle f_1, \ldots, f_k \rangle$ is a prime ideal of $\mathbb{K}[\mathbf{X}]$, or equivalently, $A := \mathbb{K}[\mathbf{X}]/I$ is an integral domain. Let J be the ideal $\langle g_1 + I, \ldots, g_m + I \rangle$ of A.*

Assume that $c \cdot \mathrm{Id}_A(J) \subseteq A$, where $c = f + I \in A \setminus \{0\}$ and $f \in \mathbb{K}_{\le d'}[\mathbf{X}]_{\le d}$. Then with respect to any monomial order, the preimage of $c \cdot \mathrm{Id}_A(J)$ in $\mathbb{K}[\mathbf{X}]$ under the natural quotient map $\mathbb{K}[\mathbf{X}] \to A$ has a Gröbner basis consisting only of polynomials in $\mathbb{K}_{\le D'}[\mathbf{X}]_{\le D}$, where $D = O_{d,n}(1)$ and $D' = O_{d,d',n}(1)$.

Proof. If J is the zero ideal, simply choose $\{1\}$ as the Gröbner basis. So assume $J \neq \{0\}$. By keeping linearly independent subsets of the nonzero polynomials among the f_i and among the g_i, we may assume $k, m \le \binom{n+d}{d}$ and $g_1, \ldots, g_m \neq 0$. For any ideal I_0 of A, denote by $\widehat{I_0}$ the preimage of I_0 in $\mathbb{K}[\mathbf{X}]$ under the natural quotient map, i.e., $\widehat{I_0} = \{a \in \mathbb{K}[\mathbf{X}] : a + I \in I_0\}$.

As $c \neq 0$ and A is an integral domain, we know c is a non-zero-divisor of A. As $c \cdot \mathrm{Id}_A(J) \subseteq A$, we know

$$c \cdot \mathrm{Id}_A(J) = \{b \in A : b = ca, aJ \subseteq J\} = \left\{b \in A : \frac{b}{c}J \subseteq J\right\} = \{b \in A : bJ \subseteq cJ\} = (cJ : J).$$

[5] As $c_1, \ldots, c_n \in \mathbb{F}$, we may assume that T is one of $X_1, \ldots, X_n$ by performing an $\mathbb{F}$-linear change of coordinates. This transformation may increase the degrees of the numerators and denominators of the coefficients, but only by a factor of $O_{d,n}(1)$, by Lemma 3.2.

It follows that

$$\widehat{c \cdot \mathrm{Id}_A(J)} = \widehat{(cJ : J)} = (\widehat{cJ} : \widehat{J}). \tag{7}$$

By definition, $\widehat{J} = \langle f_1, \ldots, f_k, g_1, \ldots, g_m \rangle$ and $\widehat{cJ} = \langle f_1, \ldots, f_k, fg_1, \ldots, fg_m \rangle$. It follows from [AL94, Lemmas 2.3.10 and 2.3.11] that

$$\begin{aligned}(\widehat{cJ} : \widehat{J}) &= \bigcap_{h \in \{f_1, \ldots, f_k, g_1, \ldots, g_m\}} \frac{1}{h}(\langle f_1, \ldots, f_k, fg_1, \ldots, fg_m \rangle \cap \langle h \rangle) \\ &= \bigcap_{i=1}^{m} \frac{1}{g_i}(\langle f_1, \ldots, f_k, fg_1, \ldots, fg_m \rangle \cap \langle g_i \rangle) \\ &= \bigcap_{i=1}^{m} \frac{1}{g_i} I_i,\end{aligned} \tag{8}$$

where $I_i := \langle f_1, \ldots, f_k, fg_1, \ldots, fg_m \rangle \cap \langle g_i \rangle$ for $i \in [m]$. By [AL94, Proposition 2.3.5], we can compute this intersection by introducing a new variable T and then eliminating it:

$$\begin{aligned}I_i &= (T \cdot \langle f_1, \ldots, f_k, fg_1, \ldots, fg_m \rangle + (1 - T) \cdot \langle g_i \rangle) \cap \mathbb{K}[\mathbf{X}] \\ &= \langle Tf_1, \ldots, Tf_k, Tfg_1, \ldots, Tfg_m, (1 - T)g_i \rangle \cap \mathbb{K}[\mathbf{X}].\end{aligned} \tag{9}$$

As $f, f_1, \ldots, f_k, g_1, \ldots, g_m \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$, by Lemma 3.3, we have $Tf_1, \ldots, Tf_k, Tfg_1, \ldots, Tfg_m, (1-T)g_1, \ldots, (1-T)g_m \in \mathbb{K}_{\leq d_0}[\mathbf{X}, T]_{\leq 2d+1}$ for some $d_0 = O_{d,d',n}(1)$. By (9) and Lemma 3.13, for $i \in [m]$, I_i has a Gröbner basis $G_i \subseteq \mathbb{K}_{\leq d_2}[\mathbf{X}]_{\leq d_1}$, where $d_1 = O_{d,n}(1)$ and $d_2 = O_{d,d',n}(1)$.

Consider $i \in [m]$. The polynomials in G_i are all divisible by g_i since $I_i \subseteq \langle g_i \rangle$. Let $\frac{1}{g_i}G_i = \{\frac{g}{g_i} : g \in G_i\}$. Note that $\mathrm{LT}(\frac{1}{g_i}G_i) = \mathrm{LT}(\frac{1}{g_i}I_i)$ since $\mathrm{LT}(G_i) = \mathrm{LT}(I_i)$. After rescaling the elements of $\frac{1}{g_i}G_i$ by nonzero elements of $\mathbb{K}$, Lemma 3.6 gives a Gröbner basis of the ideal $\frac{1}{g_i}I_i$ contained in $\mathbb{K}_{\leq d_3}[\mathbf{X}]_{\leq d_1}$, where $d_3 = O_{d,d',n}(1)$.

By [AL94, Exercise 2.3.8], we have

$$\bigcap_{i=1}^{m} \frac{1}{g_i} I_i = I^* \cap \mathbb{K}[\mathbf{X}] \tag{10}$$

where

$$I^* := \left\langle 1 - \sum_{i=1}^{m} T_i \right\rangle + T_1 \cdot \frac{1}{g_1} I_1 + \cdots + T_m \cdot \frac{1}{g_m} I_m \subseteq \mathbb{K}[\mathbf{X}, T_1, \ldots, T_m].$$

The ideal I^* has the set of generators $\{1 - \sum_{i=1}^{m} T_i\} \cup \{T_i \cdot g : g \in \frac{1}{g_i}G_i, i \in [m]\}$. These generators are in $\mathbb{K}_{\leq d_3}[\mathbf{X}, \mathbf{T}]_{\leq d_1+1}$ as $\frac{1}{g_i}G_i \subseteq \mathbb{K}_{\leq d_3}[\mathbf{X}]_{\leq d_1}$. By (10) and Lemma 3.13, $\bigcap_{i=1}^{m} \frac{1}{g_i}I_i$ has a Gröbner basis contained in $\mathbb{K}_{\leq D'}[\mathbf{X}]_{\leq D}$, where $D = O_{d,n}(1)$ and $D' = O_{d,d',n}(1)$. Finally, by (7) and (8), we know that $\bigcap_{i=1}^{m} \frac{1}{g_i}I_i$ is the preimage $\widehat{c \cdot \mathrm{Id}_A(J)}$ of $c \cdot \mathrm{Id}_A(J)$ under the natural quotient map $\mathbb{K}[\mathbf{X}] \to A$, concluding the proof. □

4 Normalization of Curves

In this section, we present a constructive process for normalizing affine curves, following the framework of Trager [Tra84], which itself is a function field analog of an algorithm by Ford and Zassenhaus [For78]. We start with an order contained in the coordinate ring and iteratively adjoin integral

elements obtained from idealizers. Allowing purely inseparable extensions of bounded exponent, we continue until the resulting order is geometrically normal over its base field.

We revisit these arguments primarily because the bounds on coefficient complexity that we need are not provided in Trager [Tra84] or other literature. Moreover, our treatment differs in several respects from that of Trager. In particular, Trager uses a different method for computing the idealizer, which imposes restrictions on the characteristic of the base field. In contrast, our method uses Gröbner bases and is characteristic-free.

Throughout this section, let $\mathbb{K} = \mathbb{F}(\mathbf{Y}) = \mathbb{F}(Y_1, \ldots, Y_\ell)$, where $\mathbb{F}$ is an algebraically closed field.

4.1 Orders, Integral Bases, and Discriminants

In this subsection, we introduce the core algebraic objects used to compute the integral closure. In particular, Theorem 4.5 provides a useful criterion for determining when an order is integrally closed.

Throughout this subsection, let $\mathbb{L}$ be a finite separable extension of $\mathbb{K}(X)$ of degree s, and let $\mathcal{O}_\mathbb{L}$ denote the integral closure of $\mathbb{K}[X]$ in $\mathbb{L}$.

Orders. A subring $\mathcal{O} \subseteq \mathcal{O}_\mathbb{L}$ is said to be a $(\mathbb{K}[X], \mathbb{L})$-order if it is a finite $\mathbb{K}[X]$-module and $\mathcal{O} \otimes_{\mathbb{K}[X]} \mathbb{K}(X) = \mathbb{L}$ (i.e., the elements in $\mathcal{O}$ span $\mathbb{L}$ over $\mathbb{K}(X)$).

Integral bases. It is known that every $(\mathbb{K}[X], \mathbb{L})$-order is a free $\mathbb{K}[X]$-module of rank $[\mathbb{L} : \mathbb{K}(X)] = s$. For a $(\mathbb{K}[X], \mathbb{L})$-order $\mathcal{O}$, we say $\mathbf{b} = (b_1, \ldots, b_s) \in \mathcal{O}^s$ is an *integral basis* of $\mathcal{O}$ if $b_1, \ldots, b_s$ form a basis of $\mathcal{O}$ over $\mathbb{K}[X]$.

Discriminants. Let $\mathbf{b} = (b_1, \ldots, b_s) \in \mathbb{L}^s$. As $\mathbb{L}/\mathbb{K}(X)$ is separable, by Galois theory, $\mathbb{L}$ has s distinct embeddings $\sigma_1, \ldots, \sigma_s$ into the Galois closure of $\mathbb{L}$ over $\mathbb{K}(X)$. The *discriminant* of $\mathbf{b}$ is defined to be

$$\operatorname{disc}(\mathbf{b}) = \det \begin{pmatrix} \sigma_1(b_1) & \sigma_1(b_2) & \cdots & \sigma_1(b_s) \\ \sigma_2(b_1) & \sigma_2(b_2) & \cdots & \sigma_2(b_s) \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_s(b_1) & \sigma_s(b_2) & \cdots & \sigma_s(b_s) \end{pmatrix}^2.$$

It is fixed by the Galois group of the Galois closure of $\mathbb{L}/\mathbb{K}(X)$, and hence is an element of $\mathbb{K}(X)$.

We also need the definition of the discriminant of a univariate polynomial. For simplicity, we only give the definition for univariate monic polynomials. Let $f(T) = \sum_{i=0}^{d} a_i T^i \in R[T]$ be a univariate monic (i.e., $a_d = 1$) polynomial of degree d over a ring R. The *discriminant* $\operatorname{disc}(f)$ of f is defined to be $(-1)^{d(d-1)/2}\operatorname{Res}(f, f') \in R$, where $\operatorname{Res}(f, f')$ denotes the resultant of f and

$f' = \sum_{i=0}^{d-1}(i+1)a_{i+1}T^i$, given by

$$\mathrm{Res}(f, f') = \det \begin{pmatrix} a_d & a_{d-1} & a_{d-2} & \cdots & a_0 & 0 & \cdots & 0 \\ 0 & a_d & a_{d-1} & \cdots & a_1 & a_0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & a_d & a_{d-1} & a_{d-2} & \cdots & a_0 \\ da_d & (d-1)a_{d-1} & (d-2)a_{d-2} & \cdots & a_1 & 0 & \cdots & 0 \\ 0 & da_d & (d-1)a_{d-1} & \cdots & 2a_2 & a_1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & da_d & (d-1)a_{d-1} & (d-2)a_{d-2} & \cdots & a_1 \end{pmatrix}.$$

The following lemma follows straightforwardly from the definition of $\mathrm{disc}(f)$ and Lemma 3.2.

Lemma 4.1. *Let $R = \mathbb{K}[X]$. Suppose $f \in R[T]$ is monic, and is in $\mathbb{K}_{\le d'}[X, T]_{\le d}$ when viewed as a polynomial in both X and T over $\mathbb{K}$. Then $\mathrm{disc}(f) \in \mathbb{K}_{\le D'}[X]_{\le D}$ for some $D = O_d(1)$ and $D' = O_{d,d'}(1)$.*

We now list some facts about the discriminant of a tuple, the discriminant of a polynomial, and their relations. These facts can be found in, e.g., [Tra84].

Lemma 4.2. *We have the following facts:*

(1) Let $\mathbf{b} = (b_1, \ldots, b_s) \in \mathbb{L}^s$. Then $\mathrm{disc}(\mathbf{b}) \neq 0$ if and only if $b_1, \ldots, b_s$ are linearly independent over $\mathbb{K}(X)$.

(2) Let $\mathbf{b} = (b_1, \ldots, b_s) \in \mathbb{L}^s$ such that $b_1, \ldots, b_s$ are integral over $\mathbb{K}[X]$. Then $\mathrm{disc}(\mathbf{b}) \in \mathbb{K}[X]$.

(3) Let $\mathbf{b}, \mathbf{b}' \in \mathbb{L}^s$, where $\mathrm{disc}(\mathbf{b}') \neq 0$. Suppose $\mathbf{b} = A \cdot \mathbf{b}'$ for some $A \in \mathbb{K}[X]^{s\times s}$. Then $\mathrm{disc}(\mathbf{b}) = \det(A)^2 \cdot \mathrm{disc}(\mathbf{b}')$. And $\mathbf{b}$ and $\mathbf{b}'$ generate the same $\mathbb{K}[X]$-module if and only if A is invertible as a matrix over $\mathbb{K}[X]$, i.e., $\det(A) \in \mathbb{K}[X]^\times = \mathbb{K}^\times$.

(4) Suppose $\mathbb{L} = \mathbb{K}(X)(\alpha)$ and $\alpha \in \mathbb{L}$ is integral over $\mathbb{K}[X]$. Let f be the minimal polynomial of α over $\mathbb{K}(X)$, which is a monic polynomial with all coefficients in $\mathbb{K}[X]$ [Mat89, Theorem 9.2]. Then $\mathrm{disc}(1, \alpha, \ldots, \alpha^{s-1}) = \pm\mathrm{disc}(f)$.

Recall that $\mathcal{O}_\mathbb{L}$ denotes the integral closure of $\mathbb{K}[X]$ in $\mathbb{L}$. We have the following statement, whose proof can be found in, e.g., [Eis95, Proof of Proposition 13.14].

Lemma 4.3. *Let $\mathcal{O} \subseteq \mathcal{O}_\mathbb{L}$ be a $(\mathbb{K}[X], \mathbb{L})$-order with an integral basis $\mathbf{b}$. Then*

$$\mathcal{O}_\mathbb{L} \subseteq \frac{1}{\mathrm{disc}(\mathbf{b})}\mathcal{O} \subseteq \frac{1}{\mathrm{disc}(\mathbf{b})}\mathcal{O}_\mathbb{L}.$$

Definition 4.4 (Discriminant ideal)**.** *Let $\mathcal{O} \subseteq \mathcal{O}_\mathbb{L}$ be a $(\mathbb{K}[X], \mathbb{L})$-order with an integral basis $\mathbf{b}$. Define $\mathfrak{D}_{\mathcal{O}/\mathbb{K}[X]}$ to be the ideal of $\mathbb{K}[X]$ generated by $\mathrm{disc}(\mathbf{b})$, called the* discriminant ideal *of $\mathcal{O}$. By Lemma 4.2(3), $\mathfrak{D}_{\mathcal{O}/\mathbb{K}[X]}$ is well-defined and does not depend on the choice of $\mathbf{b}$.*

The following theorem gives a criterion for $\mathcal{O}$ being integrally closed.

Theorem 4.5 ([Tra84])**.** *$\mathcal{O}$ is integrally closed if and only if the idealizer $\mathrm{Id}_\mathcal{O}(\mathfrak{m})$ of every maximal ideal $\mathfrak{m}$ of $\mathcal{O}$ containing $\mathfrak{D}_{\mathcal{O}/\mathbb{K}[X]}$ equals $\mathcal{O}$.*

The next lemma states that taking the idealizer does not introduce non-integral elements.

Lemma 4.6 ([Tra84])**.** *The idealizer* $\mathrm{Id}_{\mathcal{O}}(I)$ *is contained in the integral closure of* $\mathcal{O}$ *for any nonzero ideal* I *of* $\mathcal{O}$*.*

A k-algebra B is said to be *geometrically normal* over k if $B \otimes_k k'$ is normal for every field extension k'/k. We will use the following criterion.

Lemma 4.7. *Let* k *be a field, let* $E/k(x)$ *be a finite separable extension, and let* B *be a* $(k[x], E)$*-order. Let* $Q(x)$ *generate its discriminant ideal. Suppose* B *is integrally closed and* $B/\mathfrak{m}$ *is separable over* k *for every maximal ideal* $\mathfrak{m}$ *of* B *containing* $Q(x)$*. Then* B *is geometrically normal over* k*.*

Proof. The discriminant is invertible after inverting Q, so B_Q is finite étale over $k[x]_Q$ [Sta26, Tag 0BJF], and hence smooth over k. At each maximal ideal $\mathfrak{m}$ containing Q, the local ring $B_{\mathfrak{m}}$ is a normal one-dimensional Noetherian local domain, hence a discrete valuation ring. It is therefore regular; together with the separability of $B/\mathfrak{m}$ over k, this implies smoothness at $\mathfrak{m}$ [Sta26, Tag 00TV]. It follows that $\mathrm{Spec}(B)$ is smooth over k [Sta26, Tag 01V5], and is therefore geometrically normal [Sta26, Tag 056T]. □

4.2 Finding the Integral Closure

The main result of this section is the following theorem, which gives a quantitative version of the normalization of affine curves, or equivalently, of the computation of the integral closure of an order in a function field of transcendence degree one. Specifically, we give explicit bounds of the form $O_{d,d',n}(1)$ on the complexity of the defining data of the integral closure, including its generators and relations.

A technical difficulty is handling inseparable field extensions, since the base field $\mathbb{F}$ may have positive characteristic. To address this, we allow variables of the form $Y_i^{1/p}$, thereby extending the field $\mathbb{K}$ to $\mathbb{K}^{(e)}$, where e is a parameter. As before, the precise dependence on d, d', and n is unimportant, since these parameters will be treated as constants in the main argument of the next section.

Theorem 4.8. *Let* $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ *and* $f_1, \ldots, f_k \in \mathbb{K}_{\leq d'}[\mathbf{X}]_{\leq d}$*. Let* α *be an element in the* $\mathbb{F}$*-linear span of* $X_1, \ldots, X_n$ *such that the natural ring homomorphism*

$$\mathbb{K}[\alpha] \to A := \mathbb{K}[\mathbf{X}]/\langle f_1, \ldots, f_k \rangle$$

sending α *to* $\alpha + \langle f_1, \ldots, f_k \rangle$ *is injective and makes* A *a finite* $\mathbb{K}[\alpha]$*-module. Moreover, suppose the following hold:*

(1) For any algebraic extension $\mathbb{L}$ *of* $\mathbb{K}$*,* $A_{\mathbb{L}} := \mathbb{L}[\mathbf{X}]/\langle f_1, \ldots, f_k \rangle$ *is an integral domain of Krull dimension one.*

(2) $\mathrm{Frac}(A)$ *is a finite separable extension of* $\mathbb{K}(\alpha)$*.*

Then there exist $D, D', m, k', e \in \mathbb{N}$ *and* $g_1, \ldots, g_n, h_1, \ldots, h_{k'} \in \mathbb{K}^{(e)}[\mathbf{Z}] = \mathbb{K}^{(e)}[Z_1, \ldots, Z_m]$ *such that the following hold:*

(1) $D, D', m, k', p^e = O_{d,d',n}(1)$ *and* $g_1, \ldots, g_n, h_1, \ldots, h_{k'} \in \mathbb{K}^{(e)}_{\leq D'}[\mathbf{Z}]_{\leq D}$*.*

(2) The map

$$\phi : A_{\mathbb{K}^{(e)}} \to \mathbb{K}^{(e)}[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$$
$$X_i + \langle f_1, \ldots, f_k\rangle \mapsto g_i + \langle h_1, \ldots, h_{k'}\rangle, \quad i = 1, 2, \ldots, n.$$

defines an injective $\mathbb{K}^{(e)}$-linear ring homomorphism.

(3) $\mathbb{K}^{(e)}[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$ is isomorphic to the integral closure of $A_{\mathbb{K}^{(e)}}$, and this isomorphism composed with ϕ is the natural inclusion of $A_{\mathbb{K}^{(e)}}$ in its integral closure.

(4) $\mathbb{K}^{(e)}[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$ is geometrically normal over $\mathbb{K}^{(e)}$.

Remark 1. The subring $\mathbb{K}[\alpha]$ in Theorem 4.8 can be obtained via Noether normalization, though we defer this step to the next section. For our purposes, α is only in the $\mathbb{F}$-linear span of $X_1, ..., X_n$, which is independent of the Y variables.

Remark 2. The assumptions that $A_{\mathbb{L}}$ is an integral domain and that $\mathrm{Frac}(A)$ is separable over $\mathbb{K}(\alpha)$ are imposed for simplicity; they hold in our setting. It may be possible to eliminate these assumptions, but doing so could require additional tools, such as primary decomposition.

In the following, we adopt the notations and assumptions of Theorem 4.8.

By assumption, $\mathrm{Frac}(A)$ is a finite separable extension of $\mathbb{K}(\alpha)$. We view α as a variable as it is transcendental over $\mathbb{K}$. By the primitive element theorem [ZS75, §II.9, Theorem 19], we can fix $X = \sum_{i=1}^{n} c_i X_i$ with $c_i \in \mathbb{F}$ such that $\overline{X} := X + \langle f_1, \ldots, f_k\rangle$ generates the field $\mathrm{Frac}(A)$ over $\mathbb{K}(\alpha)$. Let $s := [\mathrm{Frac}(A) : \mathbb{K}(\alpha)]$ and $\mathbf{b}_0 = (1, \overline{X}, \ldots, \overline{X}^{s-1})$.

To prove Theorem 4.8, we first prove several lemmas. First, the following lemma bounds the complexity of the discriminant of $\mathbf{b}_0$.

Lemma 4.9. $\mathrm{disc}(\mathbf{b}_0) \in \mathbb{K}_{\leq d_2}[\alpha]_{\leq d_1}$ *for some* $d_1 = O_{d,n}(1)$ *and* $d_2 = O_{d,d',n}(1)$.

Proof. By assumption, A is a finite $\mathbb{K}[\alpha]$-module. So by Lemma 2.9, it is integral over $\mathbb{K}[\alpha]$. Therefore, by Lemma 4.2 (4), the minimal polynomial F of $\overline{X}$ over $\mathbb{K}(\alpha)$ is a monic polynomial over $\mathbb{K}[\alpha]$, and $\mathrm{disc}(\mathbf{b}_0) = \pm\mathrm{disc}(F)$.

We first find F as follows: Introduce new variables T and Z. Let $\preceq_{T,Z}$ be the elimination order for T on $\mathcal{M}_{T,Z}$. Let $\preceq = (\preceq_{\mathbf{X}}, \preceq_{T,Z})$ be an elimination order for $\mathbf{X}$ on $\mathcal{M}_{\mathbf{X},T,Z}$. Let G be a Gröbner basis for the ideal $I = \langle f_1, \ldots, f_k, T - X, Z - \alpha\rangle$ of $\mathbb{K}[\mathbf{X}, T, Z]$ with respect to $\preceq$. By Lemma 3.10, we may assume $G \subseteq \mathbb{K}_{\leq d_4}[\mathbf{X}, T, Z]_{\leq d_3}$ for some $d_3 = O_{d,n}(1)$ and $d_4 = O_{d,d',n}(1)$. By Lemma 3.12, $G \cap \mathbb{K}[T, Z]$ is a Gröbner basis for $I \cap \mathbb{K}[T, Z]$ with respect to $\preceq_{T,Z}$. Choose $g \in G \cap \mathbb{K}[T, Z]$ whose leading monomial is T^s, where $s = \deg_T(F)$. Such an element exists because F is monic in T and belongs to $I \cap \mathbb{K}[T, Z]$. Since $I \cap \mathbb{K}[T, Z]$ is generated by $F(T, Z)$, this g is a nonzero scalar multiple of $F(T, Z)$. Let $u \in \mathbb{K}^{\times}$ be the leading coefficient of g. Then $u^{-1}g(T, \alpha) \in (\mathbb{K}[\alpha])[T]$ is exactly the minimal polynomial F of $\overline{X}$ over $\mathbb{K}[\alpha]$. By Lemma 3.2, $F \in \mathbb{K}_{\leq d_5}[T, \alpha]_{\leq d_3}$ for some $d_5 = O_{d,d',n}(1)$. Finally, by Lemma 4.1, we have $\mathrm{disc}(\mathbf{b}_0) = \pm\mathrm{disc}(F) \in \mathbb{K}_{\leq d_2}[\alpha]_{\leq d_1}$, where $d_1 = O_{d,n}(1)$ and $d_2 = O_{d,d',n}(1)$. □

Keep the notations and assumptions in Theorem 4.8. For an algebraic extension $\mathbb{L}$ of $\mathbb{K}$, we denote by $I_{\mathbb{L}}$ the ideal $\langle f_1, \ldots, f_k\rangle$ of $\mathbb{L}[\mathbf{X}]$.

Suppose $A_{\mathbb{L}}$ and $\mathcal{O}$ are both $(\mathbb{L}[\alpha], \mathrm{Frac}(A_{\mathbb{L}}))$-orders and $A_{\mathbb{L}} \subseteq \mathcal{O}$. By Lemma 4.3 and the fact that $A_{\mathbb{L}} = \mathbb{L}[\mathbf{X}]/I_{\mathbb{L}}$, every element in $\mathcal{O}$ is of the form $\frac{a+I_{\mathbb{L}}}{\mathrm{disc}(\mathbf{b}_0)}$ for some $a \in \mathbb{L}[\mathbf{X}]$. The following

lemma shows how to translate a finite set of $\mathbb{L}[\alpha]$-algebra generators of $\mathcal{O}$ in this form into an alternative description in terms of generators and relations over $\mathbb{L}$, which is more convenient to work with. Here, we restrict our attention to the case where $\mathbb{L} = \mathbb{K}^{(e)}$.

Lemma 4.10. *Let $\mathbb{L} = \mathbb{K}^{(e)}$ for some integer $e \geq 0$. Let $a_1, \dots, a_t \in \mathbb{K}^{(e)}_{\leq d_2}[\mathbf{X}]_{\leq d_1}$, where $t = O_{d_1,n}(1)$. Assume that $\frac{a_1+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)}, \dots, \frac{a_t+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)} \in (A_{\mathbb{L}})_{\operatorname{disc}(\mathbf{b}_0)}$ generate an algebra $\mathcal{O}$ over $\mathbb{L}[\alpha]$ that is a $(\mathbb{L}[\alpha], \operatorname{Frac}(A_{\mathbb{L}}))$-order. Then there exist $r_1, \dots, r_{t+1} \in \mathbb{L}[\mathbf{X}, U]$ and $h_1, \dots, h_{k'} \in \mathbb{L}[\mathbf{Z}] = \mathbb{L}[Z_1, \dots, Z_{t+1}]$, where $k' = O_{d,d_1,n}(1)$, such that the following hold:*

(1) $r_1, \dots, r_{t+1} \in \mathbb{K}^{(e)}_{\leq d_2}[\mathbf{X}, U]_{\leq d_1+1}$ and $h_1, \dots, h_{k'} \in \mathbb{K}^{(e)}_{\leq d_4}[\mathbf{Z}]_{\leq d_3}$, where $d_3 = O_{d,d_1,n}(1)$ and $d_4 = O_{d,p^e d',d_1,d_2,n}(1)$. Moreover, $r_{t+1} = \alpha$.

(2) The map

$$\psi : \mathbb{L}[\mathbf{Z}]/\langle h_1, \dots, h_{k'}\rangle \to \mathbb{L}[\mathbf{X}, U]/\langle f_1, \dots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle \cong A_{\mathbb{L}}[U]/\langle \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$$
$$Z_i + \langle h_1, \dots, h_{k'}\rangle \mapsto r_i + \langle f_1, \dots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle, \quad i = 1, 2, \dots, t+1.$$

defines an injective $\mathbb{L}$-linear ring homomorphism.

(3) Let

$$\theta : A_{\mathbb{L}}[U]/\langle \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle \to (A_{\mathbb{L}})_{\operatorname{disc}(\mathbf{b}_0)}$$

be the $A_{\mathbb{L}}$-linear ring isomorphism sending $U + \langle \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$ to $1/\operatorname{disc}(\mathbf{b}_0)$ given by Fact 2.5. Then the image of $\theta \circ \psi$ is $\mathcal{O}$.

Proof. By assumption, $\frac{a_1+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)}, \dots, \frac{a_t+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)}$ and α generate the algebra $\mathcal{O}$ over $\mathbb{L}$. Let $r_i = a_i \cdot U$ for $i \in [t]$ and $r_{t+1} = \alpha$. Then $r_1, \dots, r_{t+1} \in \mathbb{K}^{(e)}_{\leq d_2}[\mathbf{X}, U]_{\leq d_1+1}$. Also note that we may assume $t \leq \binom{n+d_1}{d_1}$ by replacing $a_1, \dots, a_t$ with a linearly independent subset over $\mathbb{L}$.

Identifying $\mathbb{L}[\mathbf{X}, U]/\langle f_1, \dots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$ with $(A_{\mathbb{L}})_{\operatorname{disc}(\mathbf{b}_0)}$ via θ and using the definition of $r_1, \dots, r_{t+1}$, we see that the image of the map $\mathbb{L}[\mathbf{Z}] \to \mathbb{L}[\mathbf{X}, U]/\langle f_1, \dots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$ sending Z_i to $r_i + \langle f_1, \dots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$ for $i \in [t+1]$ is precisely $\mathcal{O}$. The kernel of this map is generated by a collection of polynomials $h_1, \dots, h_{k'} \in \mathbb{L}[\mathbf{Z}]$.

By Lemma 3.18 and Lemma 4.9, we may choose $h_1, \dots, h_{k'}$ to be in $\mathbb{K}^{(e)}_{\leq d_4}[\mathbf{Z}]_{\leq d_3}$ for sufficiently large $d_3 = O_{d,d_1,n}(1)$ and $d_4 = O_{d,p^e d',d_1,d_2,n}(1)$. We may also assume $k' \leq \binom{t+1+d_3}{d_3} = O_{d,d_1,n}(1)$. □

The next lemma shows that, unless an order is already geometrically normal, we can obtain a strictly larger order after a purely inseparable extension of bounded exponent.

Lemma 4.11. *Let $\mathbb{L} = \mathbb{K}^{(e)}$ for some integer $e \geq 0$. Let $a_1, \dots, a_t \in \mathbb{K}^{(e)}_{\leq d_2}[\mathbf{X}]_{\leq d_1}$. Assume $\frac{a_1+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)}, \dots, \frac{a_t+I_{\mathbb{L}}}{\operatorname{disc}(\mathbf{b}_0)}$ generate an algebra $\mathcal{O}$ over $\mathbb{L}[\alpha]$ that is a $(\mathbb{L}[\alpha], \operatorname{Frac}(A_{\mathbb{L}}))$-order. Assume moreover that $\mathbb{L}[\alpha, \overline{X}] \subseteq \mathcal{O}$. Then there exist d_3, d_4, and $e' \geq e$ such that $p^{e'}, d_3, d_4 = O_{d,d',d_1,d_2,n,p^e}(1)$, and one of the following is true:*

(1) $\mathcal{O}$ is geometrically normal over $\mathbb{L}$.

(2) There exists $a_{t+1} \in \mathbb{K}^{(e')}_{\leq d_4}[\mathbf{X}]_{\leq d_3}$ such that $\frac{a_1+I_{\mathbb{L}'}}{\operatorname{disc}(\mathbf{b}_0)}, \dots, \frac{a_{t+1}+I_{\mathbb{L}'}}{\operatorname{disc}(\mathbf{b}_0)}$ generate an algebra $\mathcal{O}'$ over $\mathbb{L}'[\alpha]$ that is a $(\mathbb{L}'[\alpha], \operatorname{Frac}(A_{\mathbb{L}'}))$-order strictly larger than $\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'$, where $\mathbb{L}' = \mathbb{K}^{(e')}$.

Proof. Assume that $\mathcal{O}$ is not geometrically normal over $\mathbb{L}$. We will prove that Item 2 holds.

We may assume $t \leq \binom{n+d_1}{d_1}$ by choosing a subset of linearly independent elements from $\{a_1, \ldots, a_t\}$. By Lemma 4.10, there exist $k' = O_{d,d_1,n}(1)$, $r_1, \ldots, r_{t+1} \in \mathbb{K}^{(e)}_{\leq d_6}[\mathbf{X}, U]_{\leq d_5}$, and $h_1, \ldots, h_{k'} \in \mathbb{K}^{(e)}_{\leq d_6}[\mathbf{Z}]_{\leq d_5}$ for some $d_5 = O_{d,d_1,n}(1)$ and $d_6 = O_{d,d',d_1,d_2,p^e,n}(1)$ such that $r_{t+1} = \alpha$, the ring homomorphism

$$\begin{aligned}\psi : \mathbb{L}[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle &\to \mathbb{L}[\mathbf{X}, U]/\langle f_1, \ldots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle \\ Z_i + \langle h_1, \ldots, h_{k'}\rangle &\mapsto r_i + \langle f_1, \ldots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle, \quad i = 1, 2, \ldots, t+1.\end{aligned}$$

is injective, and the image of ψ equals $\mathcal{O}$ if we identify $\mathbb{L}[\mathbf{X}, U]/\langle f_1, \ldots, f_k, \operatorname{disc}(\mathbf{b}_0)\cdot U - 1\rangle$ with $(A_{\mathbb{L}})_{\operatorname{disc}(\mathbf{b}_0)}$. We now use the above data to construct a_{t+1}.

If $\mathfrak{D}_{\mathcal{O}/\mathbb{L}[\alpha]}$ were the unit ideal, then $\mathcal{O}$ would be finite étale over $\mathbb{L}[\alpha]$ [Sta26, Tag 0BJF], hence smooth over $\mathbb{L}$ and therefore geometrically normal over $\mathbb{L}$ [Sta26, Tag 056T], contradicting our assumption. So $1 \notin \mathfrak{D}_{\mathcal{O}/\mathbb{L}[\alpha]}$. By Lemma 4.2 (3), $\mathfrak{D}_{\mathcal{O}/\mathbb{L}[\alpha]}$ is generated by some factor $Q(\alpha) \in \mathbb{L}[\alpha] = \mathbb{K}^{(e)}[\alpha]$ of $\operatorname{disc}(\mathbf{b}_0)$. And by Lemma 4.9 and Lemma 3.6, $Q(\alpha) \in \mathbb{K}^{(e)}_{\leq d_8}[\alpha]_{\leq d_7}$ for some $d_7 = O_{d,n}(1)$ and $d_8 = O_{d,d',n,p^e}(1)$.

Since $\mathcal{O}$ is finite free over $\mathbb{L}[\alpha]$ and $Q(\alpha)$ is nonconstant, $\mathcal{O}/Q(\alpha)\mathcal{O}$ is a nonzero finite-dimensional $\mathbb{L}$-algebra. So $Q(\alpha)\mathcal{O}$ is a zero-dimensional ideal of $\mathcal{O}$. Identify $\mathbb{L}[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$ with $\mathcal{O}$. Then $Q(Z_{t+1}) + \langle h_1, \ldots, h_{k'}\rangle$ is identified with $Q(\alpha)$ since $r_{t+1} = \alpha$.

In characteristic $p > 0$, for each maximal ideal $\mathfrak{m}$ of $\mathcal{O}$ containing $Q(\alpha)$, let $e_{\mathfrak{m}}$ be the largest nonnegative integer such that $p^{e_{\mathfrak{m}}}$ divides $[\mathcal{O}/\mathfrak{m} : \mathbb{L}]$, and let $e_0 = \max_{\mathfrak{m}} e_{\mathfrak{m}}$. By Lemma 2.14,

$$p^{e_0} \leq \max\{\deg(Q), \deg(h_1), \ldots, \deg(h_{k'})\}^{t+1} = O_{d,d_1,n,p^e}(1).$$

In characteristic zero, let $e_0 = 0$. Let $e' = e + e_0$, $\mathbb{L}' = \mathbb{K}^{(e')}$, and $B = \mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'$. Then $p^{e'} = O_{d,d_1,n,p^e}(1)$.

Since $A_{\mathbb{L}'}$ is a domain, scalar extension identifies B with a subring of $(A_{\mathbb{L}'})_{\operatorname{disc}(\mathbf{b}_0)}$. In particular, B is again an $(\mathbb{L}'[\alpha], \operatorname{Frac}(A_{\mathbb{L}'}))$-order. Scalar extension preserves an integral basis and its discriminant, so $Q(\alpha)$ also generates the discriminant ideal of B. Every maximal ideal $\mathfrak{n}$ of B containing $Q(\alpha)$ contracts to a maximal ideal $\mathfrak{m}$ of $\mathcal{O}$ containing $Q(\alpha)$. Its residue field is generated over $\mathbb{L}'$ by the image of $\mathcal{O}/\mathfrak{m}$, and is separable over $\mathbb{L}'$ by Lemma 2.2; here we use that $\mathbb{F}$ is perfect, so $\mathbb{L}' = \mathbb{L}^{1/p^{e_0}}$ in positive characteristic. In characteristic zero, separability is automatic.

If B were integrally closed, then Lemma 4.7 would make it geometrically normal over $\mathbb{L}'$. This would also make $\mathcal{O}$ geometrically normal over $\mathbb{L}$, since $\mathbb{L}$ and $\mathbb{L}'$ have the same perfect closure [Sta26, Tag 037Y], contradicting our assumption. Hence B is not integrally closed. By Theorem 4.5 and Lemma 4.6, there is a maximal ideal $\mathfrak{m}_{\mathbb{L}'}$ of B containing $Q(\alpha)$ whose idealizer is integral over B and strictly larger than it. Its residue field is separable over $\mathbb{L}'$ by the preceding argument.

We now find an element in $\operatorname{Id}_{\mathcal{O}\otimes_{\mathbb{L}}\mathbb{L}'}(\mathfrak{m}_{\mathbb{L}'}) \setminus (\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}')$. Let $c = Q(Z_{t+1}) + \langle h_1, \ldots, h_{k'}\rangle \in \mathbb{L}'[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$. View $\mathfrak{m}_{\mathbb{L}'}$ as a maximal ideal of $\mathbb{L}'[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$ containing c. Let $\widetilde{\mathfrak{m}}_{\mathbb{L}'}$ be the preimage of $\mathfrak{m}_{\mathbb{L}'}$ in $\mathbb{L}'[\mathbf{Z}]$ under the natural quotient map, which is a maximal ideal of $\mathbb{L}'[\mathbf{Z}]$ and contains the zero-dimensional ideal $\langle h_1, \ldots, h_{k'}, Q(Z_{t+1})\rangle$. As $(\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}')/\mathfrak{m}_{\mathbb{L}'}$ is separable, by Lemma 3.20, $\widetilde{\mathfrak{m}}_{\mathbb{L}'}$ is generated by polynomials in $\mathbb{K}^{(e')}_{\leq d_{10}}[\mathbf{Z}]_{\leq d_9}$, where $d_9 = O_{d,d_1,n}(1)$ and $d_{10} = O_{d,d',d_1,d_2,n,p^e}(1)$.

As $Q(\alpha)$ generates $\mathfrak{D}_{\mathcal{O}/\mathbb{L}[\alpha]}$, identifying $\mathbb{L}'[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$ with $\mathcal{O}\otimes_{\mathbb{L}}\mathbb{L}'$ and applying Lemma 4.3 and Lemma 4.6 shows that $c \cdot \operatorname{Id}_{\mathbb{L}'[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle}(\mathfrak{m}_{\mathbb{L}'}) \subseteq \mathbb{L}'[\mathbf{Z}]/\langle h_1, \ldots, h_{k'}\rangle$. Let $J \subseteq \mathbb{L}'[\mathbf{Z}]$ be the

preimage of $c \cdot \mathrm{Id}_{\mathbb{L}'[\mathbf{Z}]/\langle h_1,\dots,h_{k'}\rangle}(\mathfrak{m}_{\mathbb{L}'})$ under the natural quotient map $\mathbb{L}'[\mathbf{Z}] \to \mathbb{L}'[\mathbf{Z}]/\langle h_1,\dots,h_{k'}\rangle$. By Lemma 3.21, J has a Gröbner basis G_1 contained in $\mathbb{K}^{(e')}_{\le d_{12}}[\mathbf{Z}]_{\le d_{11}}$, where $d_{11} = O_{d,d_1,n}(1)$ and $d_{12} = O_{d,d',d_1,d_2,n,p^e}(1)$.

Let $R(\alpha) = \mathrm{disc}(\mathbf{b}_0)/Q(\alpha) \in \mathbb{L}[\alpha]$. Choose $\gamma \in G_1$ whose image in $\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'$ does not belong to $Q(\alpha)(\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}')$. Then

$$\eta := \frac{\psi_{\mathbb{L}'}(\gamma)}{Q(\alpha)} \in \mathrm{Id}_{\mathcal{O}\otimes_{\mathbb{L}}\mathbb{L}'}(\mathfrak{m}_{\mathbb{L}'}) \setminus (\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'),$$

where $\psi_{\mathbb{L}'}$ denotes the map $\psi \otimes_{\mathbb{L}} \mathbb{L}'$, and we identify its image with $\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'$.

By Lemma 4.6, η is integral over $\mathcal{O} \otimes_{\mathbb{L}} \mathbb{L}'$, and hence over $\mathbb{L}'[\alpha]$. Applying Lemma 4.3 to the initial order $\mathbb{L}'[\alpha, \overline{X}]$ gives

$$\mathrm{disc}(\mathbf{b}_0)\eta \in \mathbb{L}'[\alpha, \overline{X}] \subseteq A_{\mathbb{L}'}.$$

Its image in $(A_{\mathbb{L}'})_{\mathrm{disc}(\mathbf{b}_0)}$ is represented by $R(\alpha)\gamma(r_1,\dots,r_{t+1})$. Let G_2 be a Gröbner basis of $\langle f_1,\dots,f_k,\mathrm{disc}(\mathbf{b}_0)U-1\rangle$ with respect to an elimination order for U that is degree-compatible in $\mathbf{X}$. Define a_{t+1} to be the remainder of $R(\alpha)\gamma(r_1,\dots,r_{t+1})$ modulo G_2. Since $R(\alpha)\gamma(r_1,\dots,r_{t+1})$ represents $\mathrm{disc}(\mathbf{b}_0)\eta$, which belongs to $A_{\mathbb{L}'}$, a_{t+1} is congruent modulo $\langle f_1,\dots,f_k,\mathrm{disc}(\mathbf{b}_0)U-1\rangle$ to some $b \in \mathbb{L}'[\mathbf{X}]$. The remainder of b modulo G_2 must also be a_{t+1}. As the monomial order eliminates U, reducing b cannot introduce U, so $a_{t+1} \in \mathbb{L}'[\mathbf{X}]$.

The bounds from Lemma 3.6, Lemma 3.10, and Lemma 3.17 show that $a_{t+1} \in \mathbb{K}^{(e')}_{\le d_4}[\mathbf{X}]_{\le d_3}$ for suitable $d_3, d_4 = O_{d,d',d_1,d_2,n,p^e}(1)$. Moreover, $\frac{a_{t+1}+I_{\mathbb{L}'}}{\mathrm{disc}(\mathbf{b}_0)} = \eta$. Adjoining this element therefore gives the required strictly larger order. □

Finally, we prove Theorem 4.8 by repeatedly applying Lemma 4.11. The key idea is that the "difference" between the initial order and the maximal order (i.e., the integral closure) can be bounded via the discriminant, which in turn yields an upper bound on the number of steps (see Lemma 4.2 (3) and Lemma 4.3). Since the increase in the parameters at each step is bounded by Lemma 4.11, the final parameters can be bounded as well.

Proof of Theorem 4.8. We start from the order $\mathcal{O}_0 = \mathbb{K}[\alpha, \overline{X}]$. At each step, apply Lemma 4.11. Unless $\mathcal{O}_{i-1}$ is geometrically normal over $\mathbb{K}^{(e_{i-1})}$, this gives an order $\mathcal{O}_i$ that strictly contains $\mathcal{O}_{i-1} \otimes_{\mathbb{K}^{(e_{i-1})}} \mathbb{K}^{(e_i)}$. We stop when the order is geometrically normal over its current base field. Let τ denote the number of steps; the argument below bounds τ and proves termination.

For $i = 0, 1, \dots, \tau$, the generators of $\mathcal{O}_i$ are $(\alpha_1 + I_{\mathbb{K}^{(e_i)}})/\mathrm{disc}(\mathbf{b}_0), \dots, (\alpha_{i+1} + I_{\mathbb{K}^{(e_i)}})/\mathrm{disc}(\mathbf{b}_0)$, where $\alpha_1, \dots, \alpha_{i+1} \in \mathbb{K}^{(e_i)}_{\le d'_i}[\mathbf{X}]_{\le d_i}$. By Lemma 4.11, the parameters d_i, d'_i, p^{e_i} satisfy

$$d_0, d'_0, p^{e_0} = O_{d,d',n}(1)$$

and

$$d_i, d'_i, p^{e_i} \le F_{d,d',n}(d_{i-1}, d'_{i-1}, p^{e_{i-1}}) \quad \text{for } i > 0,$$

where $F_{d,d',n}$ is some non-decreasing function depending only on d, d', and n.[6] Define $F(0) = \max(d_0, d'_0, p^{e_0})$ and $F(i) = F_{d,d',n}(F(i-1), F(i-1), F(i-1))$. Then $d_i, d'_i, p^{e_i} \le F(i)$.

[6] We say a function F is non-decreasing if it is non-decreasing in each variable when the other variables are fixed. Note that for any function $F_0(\cdot,\cdot,\cdot)$, we can find a non-decreasing function $F(\cdot,\cdot,\cdot)$ that bounds F_0 from above by letting $F(a,b,c) = \sup_{a'\le a, b'\le b, c'\le c} F_0(a',b',c')$.

Scalar extension preserves the discriminant of an integral basis and hence its degree. By Lemma 4.2 (3), replacing $\mathcal{O}_{i-1}$ by $\mathcal{O}_i$ decreases the degree of the generator of the discriminant ideal by at least two. So the process must terminate with a geometrically normal order after at most $\tau \leq \deg_\alpha(\mathrm{disc}(\mathbf{b}_0))/2 = O_{d,n}(1)$ steps: otherwise Lemma 4.11 would give another strict enlargement and eventually force a negative discriminant degree. Therefore, the final parameters d_τ, d'_τ, and p^{e_τ} are bounded by $F(\tau) = O_{d,d',n}(1)$. It also follows by construction that $m, k' = O_{d,d',n}(1)$.

Note that $\mathcal{O}_\tau$ equals the integral closure of $A_{\mathbb{K}^{(e_\tau)}}$. This follows from Corollary 2.8 and the fact that $A_{\mathbb{K}^{(e_\tau)}}$ is integral over $\mathbb{K}^{(e_\tau)}[\alpha]$.

To find the polynomials g_i, use the map ψ associated with $\mathcal{O}_\tau$ given by Lemma 4.10, and then use Lemma 3.18 (3) to compute the inverse of $X_i + \langle f_1, \ldots, f_k, \mathrm{disc}(\mathbf{b}_0) \cdot U - 1\rangle$ under ψ for $i \in [n]$. □

5 PIT via Normalization

In this section, let $\mathbb{F}$ be a field and let $\mathbb{K} = \mathbb{F}(\mathbf{Y}) = \mathbb{F}(Y_{0,1}, \ldots, Y_{0,n}, Y_{1,1}, \ldots, Y_{1,n}, Y_{2,1}, \ldots, Y_{2,n})$.

5.1 Restricting to a Generic Affine Plane

We start by discussing the algebra resulting from restricting to a generic affine plane.

Definition 5.1 (Restriction to a generic affine plane)**.** *For a polynomial $f \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ over a field $\mathbb{F}$, define* $\mathrm{res}(f)$ *to be the polynomial*

$$f\left(Y_{0,1} + Y_{1,1}Z_1 + Y_{2,1}Z_2, \ldots, Y_{0,n} + Y_{1,n}Z_1 + Y_{2,n}Z_2\right) \in \mathbb{F}[\mathbf{Y}][Z_1, Z_2] \subseteq \mathbb{K}[Z_1, Z_2]. \tag{11}$$

Lemma 5.2. *Let $n \geq 2$. Suppose $\mathbb{F}$ is algebraically closed and $f \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ is irreducible over $\mathbb{F}$. Then* $\mathrm{res}(f) \in \mathbb{K}[\mathbf{Z}]$ *is absolutely irreducible over $\mathbb{K} = \mathbb{F}(\mathbf{Y})$.*

For a proof of Lemma 5.2, see [Kal95, Lemma 7].[7]

The next lemma states that for any algebraic extension $\mathbb{L}$ of $\mathbb{K}$, the ring $\mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f)\rangle$ is an integral domain whose field of fractions is a finite separable extension of the function field $\mathbb{L}(\overline{X}_{c_1,c_2})$, where $\overline{X}_{c_1,c_2}$ is defined below, for $(c_1, c_2) \in (\mathbb{F}^\times)^2$. Together with Noether normalization, this lemma will later be used to show that the assumptions of Theorem 4.8 can always be satisfied.

Lemma 5.3. *Let $n \geq 2$. Suppose $\mathbb{F}$ is algebraically closed and $f \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ is an irreducible polynomial over $\mathbb{F}$ of degree $d > 0$. Let $\mathbb{L}$ be an algebraic extension of $\mathbb{K} = \mathbb{F}(\mathbf{Y})$. Then $\mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f)\rangle$ is an integral domain. Moreover, for every $(c_1, c_2) \in (\mathbb{F}^\times)^2$, the field of fractions $\mathbb{E}$ of $\mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f)\rangle$ is a finite separable extension of $\mathbb{L}(\overline{X}_{c_1,c_2})$, where $\overline{X}_{c_1,c_2} := c_1Z_1 + c_2Z_2 + \langle \mathrm{res}(f)\rangle \in \mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f)\rangle \subseteq \mathbb{E}$.*

Proof. By Lemma 5.2, $\mathrm{res}(f)$ is absolutely irreducible over $\mathbb{K}$. So it is irreducible over $\mathbb{L}$. Therefore, $\mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f)\rangle$ is an integral domain.

[7] The statement [Kal95, Lemma 7] is slightly different, where the expression $Y_{0,1} + Y_{1,1}Z_1 + Y_{2,1}Z_2$ in (11) is replaced by $Y_{0,1} + Z_1$, resulting in another polynomial φ_2. However, the proof can be adapted to prove Lemma 5.2. Alternatively, we can recover (11) from φ_2 by performing the invertible $\mathbb{K}$-linear variable substitution $Z_1 \mapsto Y_{1,1}Z_1 + Y_{2,1}Z_2$, followed by applying the field automorphism of $\mathbb{K}$ that maps $Y_{1,i} \mapsto Y_{1,i}/Y_{1,1}$ and $Y_{2,i} \mapsto Y_{2,i} - Y_{2,1}Y_{1,i}/Y_{1,1}$ for $i \in [n] \setminus \{1\}$. Both transformations preserve absolute irreducibility.

Consider $(c_1, c_2) \in (\mathbb{F}^\times)^2$ and let $X_{c_1,c_2} := c_1 Z_1 + c_2 Z_2$. Then $Z_2 = c_2^{-1}(X_{c_1,c_2} - c_1 Z_1)$. So $\mathbb{L}[Z_1, Z_2] = \mathbb{L}[Z_1, X_{c_1,c_2}]$. Performing the substitution $Z_2 = c_2^{-1}(X_{c_1,c_2} - c_1 Z_1)$ in (11), we have

$$\mathrm{res}(f) = f(T_1, \ldots, T_n),$$

where

$$T_i := Y_{0,i} + Y_{1,i} Z_1 + c_2^{-1} Y_{2,i}(X_{c_1,c_2} - c_1 Z_1), \quad i = 1, 2, \ldots, n \tag{12}$$

View $\mathrm{res}(f)$ as a univariate polynomial in Z_1 over $\mathbb{L}[X_{c_1,c_2}]$ (which is fine as Z_1 and X_{c_1,c_2} are algebraically independent). Its derivative, which we denote by $D_{\mathrm{res}(f)}$, can be determined by the chain rule:

$$D_{\mathrm{res}(f)} = \sum_{i=1}^{n} \frac{\partial f}{\partial X_i}(T_1, \ldots, T_n) \frac{\partial T_i}{\partial Z_1} = \sum_{i=1}^{n} \frac{\partial f}{\partial X_i}(T_1, \ldots, T_n) \cdot (Y_{1,i} - c_2^{-1} c_1 Y_{2,i}). \tag{13}$$

As $\mathbb{F}$ is algebraically closed (and hence a perfect field) and f is irreducible over $\mathbb{F}$, we know $\frac{\partial f}{\partial X_i} \neq 0$ for some $i \in [n]$. Choose $i_0 \in [n]$ and a monomial $m = \prod_{i=1}^{n} X_i^{e_i}$ that appears in $\frac{\partial f}{\partial X_{i_0}}$ such that $\deg_{\mathbf{X}}(m)$ is maximized over all choices of i_0 and m. Then by (12) and (13), $D_{\mathrm{res}(f)}$, viewed as a polynomial in $\mathbf{Y}$, Z_1, and X_{c_1,c_2} over $\mathbb{F}$, has a monomial $X_{c_1,c_2}^{\deg_{\mathbf{X}}(m)} \left(\prod_{i=1}^{n} Y_{2,i}^{e_i}\right) Y_{1,i_0}$ that comes from the term $\frac{\partial f}{\partial X_{i_0}}(T_1, \ldots, T_n) \cdot Y_{1,i_0}$ in (13) and is not canceled by other monomials. So $D_{\mathrm{res}(f)} \neq 0$.

As the degree of $D_{\mathrm{res}(f)}$ is smaller than that of $\mathrm{res}(f)$, both viewed as univariate polynomials in Z_1 over $\mathbb{L}[X_{c_1,c_2}]$, we have $D_{\mathrm{res}(f)} \notin \langle \mathrm{res}(f) \rangle$.

Let $f_d = f_d(X_1, \ldots, X_n)$ be the homogeneous degree-d component of f, where $d = \deg(f)$. View $\mathrm{res}(f)$ as a univariate polynomial in Z_1 over $\mathbb{L}[X_{c_1,c_2}]$. Then we can write $\mathrm{res}(f) = \sum_{i=0}^{d} a_i Z_1^i$, where $a_i \in \mathbb{L}[X_{c_1,c_2}]$. By (12), we have

$$a_d = f_d(Y_{1,1} - c_2^{-1} c_1 Y_{2,1}, \ldots, Y_{1,n} - c_2^{-1} c_1 Y_{2,n}) \neq 0$$

We also have $a_d \notin \langle \mathrm{res}(f) \rangle$ as $d > 0$ and a_d is independent of Z_1.

Note that $\mathbb{L}[X_{c_1,c_2}] \cap \langle \mathrm{res}(f) \rangle = 0$ since $d > 0$. So we have an inclusion

$$\mathbb{L}[X_{c_1,c_2}] \cong \mathbb{L}[X_{c_1,c_2}] / (\mathbb{L}[X_{c_1,c_2}] \cap \langle \mathrm{res}(f) \rangle) \hookrightarrow \mathbb{L}[Z_1, Z_2] / \langle \mathrm{res}(f) \rangle$$

Taking the fields of fractions, we see that $\mathbb{L}(X_{c_1,c_2})$ may be identified with $\mathbb{L}(\overline{X}_{c_1,c_2}) \subseteq \mathbb{E}$. And the minimal polynomial of Z_1 over $\mathbb{L}(X_{c_1,c_2})$ is $F(T) := \sum_{i=0}^{d} (a_i/a_d) T^i \in \mathbb{L}(X_{c_1,c_2})[T]$. We also have $F'(Z_1 + \langle \mathrm{res}(f) \rangle) = (D_{\mathrm{res}(f)} + \langle \mathrm{res}(f) \rangle)/(a_d + \langle \mathrm{res}(f) \rangle) \neq 0$ since $a_d, D_{\mathrm{res}(f)} \notin \langle \mathrm{res}(f) \rangle$. It follows that $Z_1 + \langle \mathrm{res}(f) \rangle \in \mathbb{E}$ is separable over $\mathbb{L}(\overline{X}_{c_1,c_2})$. A symmetric argument shows that $Z_2 + \langle \mathrm{res}(f) \rangle \in \mathbb{E}$ is separable over $\mathbb{L}(\overline{X}_{c_1,c_2})$. As $\mathbb{E}$ is generated by $Z_1 + \langle \mathrm{res}(f) \rangle$ and $Z_2 + \langle \mathrm{res}(f) \rangle$, we conclude that $\mathbb{E}$ is a finite separable extension of $\mathbb{L}(\overline{X}_{c_1,c_2})$. □

5.2 Proof of the Main Theorem

For convenience, we state our main theorem again.

Theorem 5.4 (Main theorem, homogeneous version)**.** *Let $C_{n,d,k,\delta,\mathbb{F}}$ be the set of polynomials $F \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ over a field $\mathbb{F}$ satisfying the following conditions:*

(1) F can be expressed as a sum $F = \sum_{i=0}^{k_0-1} F_i$, where $k_0 \leq k$, $F_i = \prod_{j=1}^{m_i} f_{i,j}$ for $i \in \{0, 1, \ldots, k_0 - 1\}$, and each $f_{i,j} \in \mathbb{F}[\mathbf{X}]$ is a nonzero homogeneous polynomial of degree at most δ.

(2) $\deg(F_i) = d_0$ *for some* $d_0 \le d$ *and all* $i \in \{0, 1, \ldots, k_0 - 1\}$.

(3) F_i *is squarefree over* $\overline{\mathbb{F}}$ *for some* $i \in \{0, 1, \ldots, k_0 - 1\}$*, meaning that the irreducible factors of* F_i *over* $\overline{\mathbb{F}}$ *are distinct.*

Then there exists an explicit $(nd)^{O_\delta(1)}$*-sized hitting set* $\mathcal{H} \subseteq \overline{\mathbb{F}}^n$ *for* $C_{n,d,3,\delta,\mathbb{F}}$.

Assumptions. We make some further assumptions to simplify the discussion, and briefly justify them:

(1) F is non-constant. Otherwise, any nonempty set will be a hitting set for F by definition.

(2) $\mathbb{F}$ is algebraically closed. This can be guaranteed by extending $\mathbb{F}$ to $\overline{\mathbb{F}}$.

(3) $k_0 = 3$. This is because Theorem 5.4 is known to hold when $k_0 = 1$ or $k_0 = 2$ even without the condition of squarefreeness. So $F = F_0 + F_1 + F_2$.

(4) $n \ge 3$. This is because there exists an explicit hitting set of size $d^{O(1)}$ for constant-variate polynomials of degree at most d. This follows from either Kronecker substitutions or explicit hitting set constructions for sparse polynomials [KS01].

(5) F_0 is squarefree. Item 3 of Theorem 5.4 states that some F_i is squarefree. So this assumption can be guaranteed by permuting the summands F_i.

(6) The GCD of F_0, F_1, and F_2 is 1. This is because we can take out the GCD of F_0, F_1, and F_2, if it is nontrivial, from each F_i. See [Gup14] for a more detailed discussion.

(7) $f_{i,j}$ is an irreducible non-constant polynomial for $i = 0, 1, 2$ and $j \in [m_i]$. This can be guaranteed by replacing $f_{i,j}$ with its irreducible factors over $\mathbb{F}$, and absorbing the constant $f_{i,j}$ into other factors.

(8) There exists $j_0 \in [m_0]$ such that f_{0,j_0} does not divide any polynomial in $\{F_1, F_2, F_1 + F_2\}$. By permuting the factors of F_0, we may assume, without loss of generality, that $j_0 = 1$.

Assumption (8) is justified by the following lemma.

Lemma 5.5. *There exists an explicit* $(nd)^{O_\delta(1)}$*-sized hitting set for all* $F \in C_{n,d,3,\delta,\mathbb{F}}$ *satisfying Assumptions* (1)–(7) *and the additional assumption that every irreducible factor of* F_0 *over* $\mathbb{F}$ *divides some polynomial in* $\{F_1, F_2, F_1 + F_2\}$.

Proof. If $F_1 + F_2 = cF_0$ for some $c \in \mathbb{F}$, then $F = F_0 + F_1 + F_2 = (c + 1)F_0$, which is a product of polynomials of degree at most δ. This is the case where $k_0 = 1$, which is already solved as mentioned. So we assume this is not the case. In particular, $F_1 + F_2 \neq 0$. Then, $\deg(F_1 + F_2) = \deg(F_0) = d_0$. As $F_1 + F_2$ is not of the form cF_0 with $c \in \mathbb{F}$ but has degree $d_0 = \deg(F_0)$, we see that F_0 does not divide $F_1 + F_2$. As $F_0 = \prod_{j=1}^{m_0} f_{0,j}$ is squarefree, we conclude that for some $j \in [m_0]$, the factor $f_{0,j}$ does not divide $F_1 + F_2$. Fix such j.

By assumption, $f_{0,j}$ divides F_1 or F_2, but not both since it does not divide $F_1 + F_2$. By symmetry, we may assume $f_{0,j}$ divides F_1 but not F_2. By Assumption (7), $V(f_{0,j})$ is an irreducible hypersurface of $\mathbb{A}^n$. As $f_{0,j}$ divides F_1, we have $V(f_{0,j}, F_1) = V(f_{0,j})$ and therefore its codimension in $\mathbb{A}^n$ equals one. On the other hand, as $f_{0,j}$ does not divide $F_2 = \prod_{i=1}^{m_2} f_{2,i}$, the codimension

of $V(f_{0,j}, f_{2,i})$ in $\mathbb{A}^n$ equals two for all $i \in [m_2]$. We also have $\deg(V(f_{0,j})) = \deg(f_{0,j}) \le \delta$ and $\deg(V(f_{0,j}, f_{2,i})) \le \delta^2$ for $i \in [m_2]$ by Bézout's inequality.

Guo [Guo24, Theorem 1.6] showed how to explicitly construct a set $\mathcal{H}$ of affine planes in $\mathbb{A}^n$ of size at most $n^{O(\delta^2)}d^{O(1)}$ such that for at least one $P \in \mathcal{H}$, we have $\dim(V(f_{0,j}) \cap P) = 1$ and $\dim(V(f_{0,j}, f_{2,i}) \cap P) \le 0$ for all $i \in [m_2]$. Fix such P. Then

$$\dim(V(f_{0,j}, F_1) \cap P) = \dim(V(f_{0,j}) \cap P) = 1$$

and

$$\dim(V(f_{0,j}, F_2) \cap P) = \dim\left(\bigcup_{i=1}^{m_2}(V(f_{0,j}, f_{2,i}) \cap P)\right) = \max_{i\in[m_2]}(\dim(V(f_{0,j}, f_{2,i}) \cap P)) \le 0.$$

Note that this implies that $F_1 + F_2$ is not identically zero when restricting to $V(f_{0,j}) \cap P$, since otherwise we would have

$$V(f_{0,j}, F_1) \cap P = V(f_{0,j}) \cap P \cap V(F_1) = V(f_{0,j}) \cap P \cap V(-F_2) = V(f_{0,j}, F_2) \cap P,$$

contradicting the fact that $\dim(V(f_{0,j}, F_1) \cap P) \ne \dim(V(f_{0,j}, F_2) \cap P)$.

As $f_{0,j}$ divides F_0, we have $F = F_0 + F_1 + F_2 \equiv F_1 + F_2 \not\equiv 0$ when restricting to $V(f_{0,j}) \cap P$. So F is not identically zero on P. As F restricted to P is a bivariate polynomial of degree $\deg(F) = d_0 \le d$, we know how to construct an explicit hitting set of size $d^{O(1)}$ on P for F. Finally, while we do not know which affine plane $P \in \mathcal{H}$ works, we could construct an explicit hitting set for each affine plane in $\mathcal{H}$ and then take their union as the final hitting set, whose size is at most $|\mathcal{H}| \cdot d^{O(1)} \le n^{O(\delta^2)}d^{O(1)}$. □

From now on, we assume Assumptions (1)–(8). Next, we introduce the projective analogue of Definition 5.1.

Definition 5.6 (Restriction to a generic projective plane)**.** *Let $f \in \mathbb{F}[\mathbf{X}] = \mathbb{F}[X_1, \ldots, X_n]$ be a homogeneous polynomial. Define* proj.res(f) *to be the homogeneous polynomial*

$$f\left(Y_{0,1}\widehat{Z}_0 + Y_{1,1}\widehat{Z}_1 + Y_{2,1}\widehat{Z}_2, \ldots, Y_{0,n}\widehat{Z}_0 + Y_{1,n}\widehat{Z}_1 + Y_{2,n}\widehat{Z}_2\right) \in \mathbb{F}[\mathbf{Y}][\widehat{Z}_0, \widehat{Z}_1, \widehat{Z}_2] \subseteq \mathbb{K}[\widehat{Z}_0, \widehat{Z}_1, \widehat{Z}_2]. \tag{14}$$

Consider the projective space $\mathbb{P}^2_{\mathbb{K}}$ over $\mathbb{K} = \mathbb{F}(Y_{0,1}, Y_{1,1}, Y_{2,1}, \ldots, Y_{0,n}, Y_{1,n}, Y_{2,n})$ with homogeneous coordinates $\widehat{Z}_0$, $\widehat{Z}_1$, and $\widehat{Z}_2$. For $i = 0, 1, 2$, let $U_i \cong \mathbb{A}^2_{\mathbb{K}}$ be the affine open chart of $\mathbb{P}^2_{\mathbb{K}}$ defined by $\widehat{Z}_i \ne 0$.

The homogeneous polynomial proj.res$(f_{0,1})$ defines a projective curve $C \subseteq \mathbb{P}^2_{\mathbb{K}}$. Identify U_0 with $\mathbb{A}^2_{\mathbb{K}}$ and let $Z_1 = \widehat{Z}_1/\widehat{Z}_0$ and $Z_2 = \widehat{Z}_2/\widehat{Z}_0$ be the coordinates of U_0. By Definition 5.1 and Definition 5.6, $C \cap U_0 \subseteq U_0$ is defined precisely by the polynomial res$(f_{0,1})$. By Assumption (7), $f_{0,1}$ is irreducible over $\mathbb{F}$. So by Lemma 5.3, res$(f_{0,1})$ is absolutely irreducible. Thus, the affine curve $C \cap U_0$ and the projective curve C are both absolutely irreducible.

Let $g =$ proj.res(F_1)/proj.res(F_2), which is a homogeneous rational function of degree $\deg(F_1) - \deg(F_2) = 0$ on the projective space $\mathbb{P}^2_{\mathbb{K}}$. By Assumption (8), $f_{0,1}$ divides neither F_1 nor F_2.

Claim 5.7. proj.res$(f_{0,1})$ *divides neither* proj.res(F_1) *nor* proj.res(F_2).

Subproof of Claim 5.7. Write $f = f_{0,1}$ and $R = \mathbb{F}[\mathbf{Y}]$, so that $\mathbb{K} = \mathrm{Frac}(R)$. Since f is irreducible, $f(\mathbf{Y}_0)$ remains irreducible after adjoining the variables $\mathbf{Y}_1, \mathbf{Y}_2, Z_1, Z_2$. The substitution $Y_{0,j} \mapsto Y_{0,j} + Y_{1,j}Z_1 + Y_{2,j}Z_2$, fixing all other variables, is an $\mathbb{F}$-algebra automorphism of $R[Z_1, Z_2]$ sending $f(\mathbf{Y}_0)$ to $\mathrm{res}(f)$. Thus $\mathrm{res}(f)$ is irreducible in $R[Z_1, Z_2]$. Since it also has positive degree in (Z_1, Z_2), its coefficients have no common nonunit factor in R; that is, it is primitive over R.

Suppose that $\mathrm{proj.res}(f) \mid \mathrm{proj.res}(F_i)$ for some $i \in \{1, 2\}$. Setting $\widehat{Z}_0 = 1$ gives $\mathrm{res}(f) \mid \mathrm{res}(F_i)$ in $\mathbb{K}[Z_1, Z_2]$. By Gauss's lemma, divisibility by a primitive polynomial over a UFD can be checked over its fraction field. Hence this divisibility already holds in $R[Z_1, Z_2]$. Setting $Z_1 = Z_2 = 0$ and $\mathbf{Y}_1 = \mathbf{Y}_2 = 0$ then yields $f(\mathbf{Y}_0) \mid F_i(\mathbf{Y}_0)$, contradicting Assumption (8). □

We divide the proof of the main theorem into two cases.

The easier case. In the easier case, $\mathrm{proj.res}(F_1)$ and $\mathrm{proj.res}(F_2)$ have the same order of vanishing at every point of the normalized curve $\widetilde{C}$, so that the quotient $\mathrm{proj.res}(F_1)/\mathrm{proj.res}(F_2)$ is well-defined everywhere on $\widetilde{C}$. We show that in this case, $g = \mathrm{proj.res}(F_1)/\mathrm{proj.res}(F_2)$ is a constant on $\widetilde{C}$, and in fact, F_1/F_2 is a constant modulo $f_{0,1}$, allowing us to merge F_1 and F_2 in the analysis. The following lemma and its proof make this precise.

Lemma 5.8. *There exists an explicit set $S \subseteq \mathbb{F}^n$ of size $n^{O(\delta^2)}d^{O(1)}$ independent of F such that, if the restriction[8] of $g = \mathrm{proj.res}(F_1)/\mathrm{proj.res}(F_2)$ to the normalization $\widetilde{C \cap U_i}$ of $C \cap U_i$ is regular for $i = 0, 1, 2$, then S is a hitting set for F.*

Proof. Suppose the restriction of $g = \mathrm{proj.res}(F_1)/\mathrm{proj.res}(F_2)$ to $\widetilde{C \cap U_i}$ is regular for $i = 0, 1, 2$. These normalizations $\widetilde{C \cap U_i}$ with $i = 0, 1, 2$ glue together to form the normalization $\widetilde{C}$ of C; see [Eis95, Proposition 4.13] and the discussion thereafter. Then $g|_C$ is a regular function on $\widetilde{C}$.

Let $A_i = \mathbb{K}[C \cap U_i]$. By Lemma 5.2, even after changing the base field $\mathbb{K}$ to $\overline{\mathbb{K}}$, the coordinate ring $A_i \otimes_{\mathbb{K}} \overline{\mathbb{K}} = \overline{\mathbb{K}}[C \cap U_i]$ is an integral domain. As $\overline{\mathbb{K}}$ is free over $\mathbb{K}$, we have an injection

$$\widetilde{A_i} \otimes_{\mathbb{K}} \overline{\mathbb{K}} \hookrightarrow \mathrm{Frac}(A_i) \otimes_{\mathbb{K}} \overline{\mathbb{K}}.$$

The ring on the right is a localization of $A_i \otimes_{\mathbb{K}} \overline{\mathbb{K}}$ and hence is a domain. Thus $\widetilde{A_i} \otimes_{\mathbb{K}} \overline{\mathbb{K}}$ is a domain as well, proving that $\widetilde{C}$ is *geometrically integral* [Sta26, Tag 05DW, Tag 0366]. Moreover, $\widetilde{C}$ is projective and hence *proper* over $\mathbb{K}$.[9] Any regular function on a geometrically integral and proper variety lives in the base field [Sta26, Tag 0BUG]. So $g|_{\widetilde{C}} \in \mathbb{K}$. It follows that $g|_C \in \mathbb{K} = \mathbb{F}(\mathbf{Y})$.

View F_0, F_1, F_2 as polynomials on $\mathbb{A}^n_{\mathbb{F}}$. Let H be the irreducible hypersurface in $\mathbb{A}^n_{\mathbb{F}}$ defined by $f_{0,1}$. As $f_{0,1}$ divides neither F_1 nor F_2, we know F_1/F_2 restricts to a nonzero rational function $(F_1/F_2)|_H \in \mathbb{F}(H)$ on H.

We claim that $(F_1/F_2)|_H \in \mathbb{F}$. Assume that this is not the case. Then for a general point $\mathbf{x} = (\mathbf{a}, \mathbf{b}, \mathbf{c}) \in \mathbb{F}^{3n}$, we have

1. $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{F}^n$ are linearly independent.

[8] The restriction of g to $\widetilde{C \cap U_i}$ means first restricting g to $C \cap U_i$, and then viewing it as a (rational) function on $\widetilde{C \cap U_i}$. Note that $C \cap U_i$ and $\widetilde{C \cap U_i}$ share the same function field, $\mathrm{Frac}(\mathbb{K}[C \cap U_i])$.

[9] Properness is a version of compactness in algebraic geometry. See [Vak24, §11.4]. Here we use the fact that regular functions on a geometrically integral and proper variety are constants, which is an analog of the fact in complex geometry that holomorphic functions on a connected complex compact manifold are constants.

2. $g|_C \in \mathbb{K} = \mathbb{F}(\mathbf{Y})$ is regular at $\mathbf{x}$.

3. The affine line passing through $\mathbf{a}$ and $\mathbf{b}$ intersects H at some point $\mathbf{u}$.

4. The affine line passing through $\mathbf{a}$ and $\mathbf{c}$ intersects H at some point $\mathbf{v}$.

5. F_1/F_2 is regular at $\mathbf{u}$ and $\mathbf{v}$ but $(F_1/F_2)(\mathbf{u}) \neq (F_1/F_2)(\mathbf{v})$.

Here, the last item holds because, for a fixed general vector $\mathbf{a}$, both $\mathbf{u}$ and $\mathbf{v}$ can vary independently over a dense open subset of H, depending on $\mathbf{b}$ and $\mathbf{c}$, respectively.

Fix such $\mathbf{x} = (\mathbf{a}, \mathbf{b}, \mathbf{c})$. We may write $\mathbf{u} = \alpha\mathbf{a} + (1-\alpha)\mathbf{b}$ and $\mathbf{v} = \beta\mathbf{a} + (1-\beta)\mathbf{c}$ for some $\alpha, \beta \in \mathbb{F}$. By definition, as $g|_C(\mathbf{x}) \in \mathbb{K}$ does not depend on $(\widehat{Z_0}, \widehat{Z_1}, \widehat{Z_2})$, we have

$$(F_1/F_2)(\mathbf{u}) = \frac{\text{proj.res}(F_1)}{\text{proj.res}(F_2)}\bigg|_{\mathbf{Y}=\mathbf{x},(\widehat{Z_0},\widehat{Z_1},\widehat{Z_2})=(\alpha,1-\alpha,0)} = g|_C(\mathbf{x}),$$

and similarly,

$$(F_1/F_2)(\mathbf{v}) = \frac{\text{proj.res}(F_1)}{\text{proj.res}(F_2)}\bigg|_{\mathbf{Y}=\mathbf{x},(\widehat{Z_0},\widehat{Z_1},\widehat{Z_2})=(\beta,0,1-\beta)} = g|_C(\mathbf{x}).$$

So $(F_1/F_2)(\mathbf{u}) = (F_1/F_2)(\mathbf{v})$, contradicting the fact that $(F_1/F_2)(\mathbf{u}) \neq (F_1/F_2)(\mathbf{v})$.

Therefore, $(F_1/F_2)|_H \in \mathbb{F}$. Denote $(F_1/F_2)|_H$ by γ. The facts that $(F_1/F_2)|_H = \gamma$, H is defined by $f_{0,1}$, and $f_{0,1}$ divides F_0 imply that

$$F = F_0 + F_1 + F_2 \equiv F_1 + F_2 \equiv (1+\gamma)F_2 = (1+\gamma)\prod_{i=1}^{m_2} f_{2,i} \pmod{f_{0,1}}. \tag{15}$$

If $1+\gamma = 0$, then by (15), $F_1 + F_2$ would be divisible by $f_{0,1}$, contradicting Assumption (8). So $1+\gamma \neq 0$. Therefore, by (15),

$$V(f_{0,1}, F_1 + F_2) = \bigcup_{i=1}^{m_2} V(f_{0,1}, f_{2,i}),$$

where the degree of each $V(f_{0,1}, f_{2,i})$ in $\mathbb{A}^n_{\mathbb{F}}$ is bounded by δ^2 by Bézout's inequality.

As $f_{0,1}$ does not divide F_2, the codimension of each $V(f_{0,1}, f_{2,i})$ is exactly two. Apply the same plane-selection argument as in Lemma 5.5. It gives an explicit family $\mathcal{H}$ containing a plane P such that $\dim(V(f_{0,1}) \cap P) = 1$ and $\dim(V(f_{0,1}, f_{2,i}) \cap P) \leq 0$ for all $i \in [m_2]$. Hence the one-dimensional variety $V(f_{0,1}) \cap P$ is not contained in the union of these finite sets. This implies that $\prod_{i=1}^{m_2} f_{2,i}$ is nonzero when restricted to $V(f_{0,1}) \cap P$. Combining this with (15) and the fact that $1+\gamma \neq 0$, we have $F \equiv (1+\gamma)\prod_{i=1}^{m_2} f_{2,i} \not\equiv 0$ when restricted to $V(f_{0,1}) \cap P$. In particular, F is nonzero when restricted to P.

The rest of the proof is the same as the last part of the proof of Lemma 5.5: As F restricted to P is a bivariate polynomial of degree $\deg(F) = d_0 \leq d$, we know how to construct an explicit hitting set of size $d^{O(1)}$ on P for F. While we do not know which affine plane $P \in \mathcal{H}$ works, we could construct an explicit hitting set for each affine plane in $\mathcal{H}$ and then take their union as the final hitting set, whose size is at most $|\mathcal{H}| \cdot d^{O(1)} \leq n^{O(\delta^2)} d^{O(1)}$. □

The harder case. Next, we address the harder case, where proj.res(F_1) and proj.res(F_2) vanish at some point $\mathfrak{m} \in \widetilde{C}$ with orders $k_1, k_2 \geq 0$, respectively, and $k_1 \neq k_2$. The basic idea is to specialize the projective plane (i.e., by assigning $\mathbf{a} = (a_{0,1}, \ldots, a_{2,n}) \in \mathbb{F}^{3n}$ to $(Y_{0,1}, \ldots, Y_{2,n})$) so that the two orders of vanishing remain distinct after specialization. This, in turn, ensures that the polynomial F remains nonzero when restricted to the specialized plane.

To achieve this, we exploit the structure of F_1 and F_2. Recall that F_1 and F_2 are products of polynomials $f_{i,j}$ of constant degree. Since $\widetilde{C}$ is normal and the function $\mathrm{ord}_{\mathfrak{m}}(\cdot)$, which maps a function to its order of vanishing at $\mathfrak{m}$, is additive on products (as it is a discrete valuation), we can "decompose" F_1 and F_2 into their constant-degree factors $f_{i,j}$ in our analysis. This decomposition plays a key role in efficiently derandomizing the choice of $\mathbf{a}$.

We now provide the details:

Lemma 5.9. *For $i = 0, 1, 2$, there exists an explicit set $\mathcal{H}_i \subseteq \mathbb{F}^n$ of size at most $(nd)^{O_\delta(1)}$ independent of F such that, if the restriction of $g = \text{proj.res}(F_1)/\text{proj.res}(F_2)$ to the normalization $\widetilde{C \cap U_i}$ of $C \cap U_i$ is not regular, then $\mathcal{H}_i$ is a hitting set for F.*

Proof. By symmetry, it suffices to explicitly construct $\mathcal{H}_0$ and prove that this set satisfies the lemma. Let $C_0 = C \cap U_0$. Suppose the restriction of g to $\widetilde{C_0}$ is not regular. Identifying U_0 with the affine plane $\mathbb{A}^2_{\mathbb{K}}$ using the coordinates $Z_1 = \widehat{Z}_1/\widehat{Z}_0$ and $Z_2 = \widehat{Z}_2/\widehat{Z}_0$, one can see that the affine curve C_0 is defined by $\mathrm{res}(f_{0,1})$ in $\mathbb{A}^2_{\mathbb{K}}$, and $g = \text{proj.res}(F_1)/\text{proj.res}(F_2)$ restricted to $\mathbb{A}^2_{\mathbb{K}}$ becomes the rational function $\mathrm{res}(F_1)/\mathrm{res}(F_2)$. So $(\mathrm{res}(F_1)/\mathrm{res}(F_2))|_{\widetilde{C_0}}$ is not regular.

The coordinate ring of C_0 is the integral domain $A := \mathbb{K}[Z_1, Z_2]/\langle \mathrm{res}(f_{0,1}) \rangle$, and the coordinate ring of $\widetilde{C_0}$ is its integral closure $\widetilde{A}$. Define

$$g_0 := \frac{\mathrm{res}(F_1) + \langle \mathrm{res}(f_{0,1}) \rangle}{\mathrm{res}(F_2) + \langle \mathrm{res}(f_{0,1}) \rangle} \in \mathrm{Frac}(A). \tag{16}$$

As mentioned above, $(\mathrm{res}(F_1)/\mathrm{res}(F_2))|_{\widetilde{C_0}}$ is not a regular function. By definition, this means that it does not belong to the coordinate ring of $\widetilde{C_0}$, that is, $g_0 \notin \widetilde{A}$.

Step 1: Representing the normalization. We start by representing $\widetilde{A}$, or more precisely, the integral closure $\widetilde{A_{\mathbb{K}^{(e)}}}$ of $A_{\mathbb{K}^{(e)}} = A \otimes_{\mathbb{K}} \mathbb{K}^{(e)}$ for some $e \geq 0$, in terms of generators and relations. Under this representation, the subring $A_{\mathbb{K}^{(e)}}$ can also be represented explicitly. This is achieved via Theorem 4.8.

By Noether normalization (Lemma 2.12), there exist $c_1, c_2 \in \mathbb{F}^\times$ such that for $\alpha = c_1 Z_1 + c_2 Z_2$, the natural ring homomorphism $\mathbb{K}[\alpha] \to A$ sending α to $\alpha + \langle \mathrm{res}(f_{0,1}) \rangle$ is injective and makes A a finite $\mathbb{K}[\alpha]$-module. Moreover, by Lemma 5.2, for any algebraic extension $\mathbb{L}$ of $\mathbb{K}$, $A_{\mathbb{L}} := \mathbb{L}[Z_1, Z_2]/\langle \mathrm{res}(f_{0,1}) \rangle$ is an integral domain of Krull dimension one. And by Lemma 5.3, $\mathrm{Frac}(A_{\mathbb{L}})$ is a finite separable extension of $\mathbb{L}(\alpha)$. By Definition 5.1 and the fact that $\deg(f_{0,1}) \leq \delta$, we have $\mathrm{res}(f_{0,1}) \in \mathbb{K}_{\leq\delta}[Z_1, Z_2]_{\leq\delta}$. (Similarly, $\mathrm{res}(f_{i,j}) \in \mathbb{K}_{\leq\delta}[Z_1, Z_2]_{\leq\delta}$ for all $i = 0, 1, 2$ and $j \in [m_i]$.) So, by Theorem 4.8, there exist integers D, D', m, k', e satisfying $D, D', m, k', p^e = O_\delta(1)$, and $g_1, g_2, h_1, \ldots, h_{k'} \in \mathbb{K}^{(e)}[\mathbf{T}] = \mathbb{K}^{(e)}[T_1, \ldots, T_m]$ such that the following hold:

(1) $g_1, g_2, h_1, \ldots, h_{k'} \in \mathbb{K}^{(e)}_{\leq D'}[\mathbf{T}]_{\leq D}$.

(2) The map

$$\begin{aligned} \phi : A_{\mathbb{K}^{(e)}} &\to \mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle \\ Z_i + \langle \mathrm{res}(f_{0,1}) \rangle &\mapsto g_i + \langle h_1, \ldots, h_{k'} \rangle, \quad i = 1, 2. \end{aligned}$$

defines an injective $\mathbb{K}^{(e)}$-linear ring homomorphism.

(3) $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \dots, h_{k'}\rangle$ is isomorphic to the integral closure $\widetilde{A_{\mathbb{K}^{(e)}}}$ of $A_{\mathbb{K}^{(e)}}$, and this isomorphism composed with ϕ is the natural inclusion of $A_{\mathbb{K}^{(e)}}$ in its integral closure.

(4) $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \dots, h_{k'}\rangle$ is geometrically normal over $\mathbb{K}^{(e)}$.

We claim that g_0 is not in the integral closure $\widetilde{A_{\mathbb{K}^{(e)}}}$ of $A_{\mathbb{K}^{(e)}}$ either. To see this, assume to the contrary that $g_0 \in \widetilde{A_{\mathbb{K}^{(e)}}}$. Then it is integral over $A_{\mathbb{K}^{(e)}}$. And $A_{\mathbb{K}^{(e)}} = A \otimes_{\mathbb{K}} \mathbb{K}^{(e)}$ is integral over A. So by Lemma 2.7, g_0 is integral over A, contradicting the fact that $g_0 \notin \widetilde{A}$. Therefore, $g_0 \notin \widetilde{A_{\mathbb{K}^{(e)}}}$.

Step 2: Choosing a valuation detecting non-regularity. By Lemma 2.6, we have

$$\widetilde{A_{\mathbb{K}^{(e)}}} = \bigcap_{\text{maximal ideal } \mathfrak{m} \subseteq \widetilde{A_{\mathbb{K}^{(e)}}}} (\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}.^{10} \tag{17}$$

By (17), for some maximal ideal $\mathfrak{m}$ of $\widetilde{A_{\mathbb{K}^{(e)}}}$, we have $g_0 \notin (\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}$. Fix such $\mathfrak{m}$. As $\widetilde{A_{\mathbb{K}^{(e)}}}$ is integrally closed, the one-dimensional local ring $(\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}$ is a discrete valuation ring and is equipped with a normalized valuation $\mathrm{ord}_{\mathfrak{m}}(\cdot)$. An element $f \in \mathrm{Frac}(A_{\mathbb{K}^{(e)}})$ is in $(\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}$ if and only if $\mathrm{ord}_{\mathfrak{m}}(f) \geq 0$.

For $i = 1, 2$ and $j \in [m_i]$, let

$$k_{i,j} = \mathrm{ord}_{\mathfrak{m}}(\mathrm{res}(f_{i,j}) + \langle \mathrm{res}(f_{0,1})\rangle) \quad \text{and} \quad n_i = \mathrm{ord}_{\mathfrak{m}}(\mathrm{res}(F_i) + \langle \mathrm{res}(f_{0,1})\rangle). \tag{18}$$

Then for $i = 1, 2$, since $F_i = \prod_{j=1}^{m_i} f_{i,j}$, we have

$$n_i = \mathrm{ord}_{\mathfrak{m}}(\mathrm{res}(F_i) + \langle \mathrm{res}(f_{0,1})\rangle) = \sum_{j\in[m_i]} \mathrm{ord}_{\mathfrak{m}}(\mathrm{res}(f_{i,j}) + \langle \mathrm{res}(f_{0,1})\rangle) = \sum_{j\in[m_i]} k_{i,j}. \tag{19}$$

By (16) and the fact that $g_0 \notin (\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}$, we have

$$0 \leq n_1 < n_2. \tag{20}$$

Step 3: Representing the valuation. Next, we realize the valuation via polynomials of bounded complexity.

We know that $\widetilde{A_{\mathbb{K}^{(e)}}}$ may be identified with $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \dots, h_{k'}\rangle$ such that ϕ becomes the inclusion $A_{\mathbb{K}^{(e)}} \hookrightarrow \widetilde{A_{\mathbb{K}^{(e)}}}$. We now consider the latter ring $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \dots, h_{k'}\rangle$ for computational purposes.

View $\mathfrak{m}$ as a maximal ideal of $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \dots, h_{k'}\rangle$. Let $\widehat{\mathfrak{m}} \subseteq \mathbb{K}^{(e)}[\mathbf{T}]$ be the preimage of $\mathfrak{m}$ under the natural quotient map, which is a maximal ideal of $\mathbb{K}^{(e)}[\mathbf{T}]$.

Next, we find a Gröbner basis of $\widehat{\mathfrak{m}}$ whose members have low complexity. For $i \in \{0, 1, 2\}$ and $j \in [m_i]$, since $\mathrm{res}(f_{i,j}) \in \mathbb{K}_{\leq\delta}[Z_1, Z_2]_{\leq\delta}$ and $g_1, g_2 \in \mathbb{K}^{(e)}_{\leq D'}[\mathbf{T}]_{\leq D}$, where $D, D', p^e = O_\delta(1)$, we have $\mathrm{res}(f_{i,j})(g_1, g_2) \in \mathbb{K}^{(e)}_{\leq d_2}[\mathbf{T}]_{\leq d_1}$ for some $d_1, d_2 = O_\delta(1)$. By (19) and (20), there exists $j_0 \in [m_2]$ such

[10] Alternatively, (17) follows from *Algebraic Hartogs's Lemma* [Vak24, 13.5.19], where the intersection is taken over all prime ideals of codimension one. This extends to higher-dimensional normal varieties.

that $k_{2,j_0} > 0$. So $\mathrm{res}(f_{2,j_0})(g_1, g_2) \in \widehat{\mathfrak{m}} + \langle h_1, \ldots, h_{k'} \rangle$. Therefore, $\widehat{\mathfrak{m}}$ contains the zero-dimensional ideal $\langle h_1, \ldots, h_{k'}, \mathrm{res}(f_{2,j_0})(g_1, g_2) \rangle$.

Let $\mathbb{L} = \mathbb{K}^{(e)}$ and $C = \mathbb{L}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle$. By Lemma 2.14, $[C/\mathfrak{m} : \mathbb{L}] = O_\delta(1)$. In characteristic $p > 0$, let r be the largest integer such that p^r divides this degree, and let $\mathbb{L}' = \mathbb{K}^{(e+r)}$; in characteristic zero, let $r = 0$ and $\mathbb{L}' = \mathbb{L}$. Let $\mathfrak{m}'$ be the unique maximal ideal of $C_{\mathbb{L}'} := C \otimes_{\mathbb{L}} \mathbb{L}'$ above $\mathfrak{m}$. By Lemma 2.2, its residue field is separable over $\mathbb{L}'$.

By the geometric normality obtained in Step 1, $C_{\mathbb{L}'}$ remains normal. It is integral over $A_{\mathbb{L}'}$ and has the same fraction field, so it is the normalization of $A_{\mathbb{L}'}$. The inclusion of discrete valuation rings $C_{\mathfrak{m}} \subseteq (C_{\mathbb{L}'})_{\mathfrak{m}'}$ implies that the new normalized valuation restricts to a positive integer multiple of the old one: old units remain units, and an old uniformizer has positive valuation in the new ring. Replace $e, \mathfrak{m}, \widehat{\mathfrak{m}}$ by their counterparts over $\mathbb{L}'$, and redefine $k_{i,j}$ and n_i using this new valuation. Then (19) and (20) still hold. Since $p^r = O_\delta(1)$, all complexity bounds remain $O_\delta(1)$ after this scalar extension. By Lemma 3.20, $\widehat{\mathfrak{m}}$ admits a Gröbner basis $G \subseteq \mathbb{K}^{(e)}_{\leq d_4}[\mathbf{T}]_{\leq d_3}$ for some $d_3, d_4 \in O_\delta(1)$.

As $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle$ is isomorphic to (and hence identified with) $\widetilde{A_{\mathbb{K}^{(e)}}}$, the local ring

$$B := (\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle)_{\mathfrak{m}}$$

is a discrete valuation ring, naturally identified with $(\widetilde{A_{\mathbb{K}^{(e)}}})_{\mathfrak{m}}$. By a slight abuse of notation, we continue to denote its normalized valuation by $\mathrm{ord}_{\mathfrak{m}}$. As $\mathrm{ord}_{\mathfrak{m}}$ is normalized and G generates the ideal $\widehat{\mathfrak{m}}$, there exists $u \in G$ such that $\overline{u} := u + \langle h_1, \ldots, h_{k'} \rangle$ satisfies $\mathrm{ord}_{\mathfrak{m}}(\overline{u}) = 1$.

Step 4: Deriving polynomial identities. We now encode the valuation information in (18) as a collection of low-complexity polynomial identities.

Fix $i \in \{1, 2\}$ and $j \in [m_i]$, and write $I = \langle h_1, \ldots, h_{k'} \rangle$, $C = \mathbb{K}^{(e)}[\mathbf{T}]/I$, and $P = \mathrm{res}(f_{i,j})(g_1, g_2)$. Let $k = k_{i,j}$. In the discrete valuation ring $B = C_{\mathfrak{m}}$, we have $P + I = \overline{u}^k v$ for some unit $v \in B^\times$. Write $v = (\rho + I)/(\sigma + I)$ with $\rho, \sigma \in \mathbb{K}^{(e)}[\mathbf{T}] \setminus \widehat{\mathfrak{m}}$. Since $C \hookrightarrow B$ is injective, $\sigma P - u^k \rho \in I$. Choose $\tau \in \mathbb{K}^{(e)}[\mathbf{T}]$ such that $\tau\rho \equiv 1 \pmod{\widehat{\mathfrak{m}}}$. Taking $t_{i,j} = \tau\sigma$ gives

$$t_{i,j} P - u^k \in u^k \widehat{\mathfrak{m}} + I. \tag{21}$$

Moreover, $k = O_\delta(1)$: the local quotient $B/(P + I)$ has length k, so $k \leq \dim_{\mathbb{K}^{(e)}}(C/(P + I)) = O_\delta(1)$ by Lemma 2.14. Since (21) implies $u^k \in \langle P, \{u^k g : g \in G\}, h_1, \ldots, h_{k'} \rangle$, Lemma 3.11 lets us choose $t_{i,j}$ and witness polynomials of bounded complexity such that

$$t_{i,j}\mathrm{res}(f_{i,j})(g_1, g_2) = u^{k_{i,j}} \left(1 + \sum_{g \in G} a_{i,j,g} g\right) + \sum_{\ell=1}^{k'} b_{i,j,\ell} h_\ell, \tag{22}$$

where $t_{i,j} \in \mathbb{K}^{(e)}_{\leq d_6}[\mathbf{T}]_{\leq d_5}$, $a_{i,j,g}, b_{i,j,\ell} \in \mathbb{K}^{(e)}_{\leq d_8}[\mathbf{T}]_{\leq d_7}$, and $d_5, d_6, d_7, d_8 = O_\delta(1)$.

Cancelling $\overline{u}^k$ in B and reducing (22) modulo $\mathfrak{m}B$ gives $t_{i,j} v \equiv 1 \pmod{\mathfrak{m}B}$. Thus this bounded-complexity choice of $t_{i,j}$ is still invertible modulo $\widehat{\mathfrak{m}}$. Applying Lemma 3.11 to $1 \in \langle t_{i,j} \rangle + \widehat{\mathfrak{m}}$ gives

$$s_{i,j} t_{i,j} = 1 + \sum_{g \in G} c_{i,j,g} g, \tag{23}$$

with $s_{i,j}, c_{i,j,g} \in \mathbb{K}^{(e)}_{\leq d_{10}}[\mathbf{T}]_{\leq d_9}$ and $d_9, d_{10} = O_\delta(1)$.

Similarly, as $u \in \widehat{\mathfrak{m}}$ and $\langle h_1, \ldots, h_{k'} \rangle \subseteq \widehat{\mathfrak{m}}$, by Lemma 3.11, we can write

$$u = \sum_{g \in G} d_g g \quad \text{and} \quad h_\ell = \sum_{g \in G} e_{\ell,g} g \ \text{ for } \ \ell \in [k'] \tag{24}$$

with $d_g, e_{\ell,g} \in \mathbb{K}^{(e)}_{\leq d_{12}}[\mathbf{T}]_{\leq d_{11}}$ for all $g \in G$ and $\ell \in [k']$, where $d_{11}, d_{12} \in O_\delta(1)$.

Step 5: Constructing specialization conditions. We adopt the following notation: Suppose Q is a polynomial over $\mathbb{K}^{(e)}$, and $\mathbf{a} \in \mathbb{F}^{3n}$ is a point such that none of the denominators of the coefficients of Q, which are polynomials in $\mathbf{Y}^{1/p^e}$, vanishes after assigning $\mathbf{a}$ to $\mathbf{Y}^{1/p^e}$. Then the resulting polynomial after the assignment $\mathbf{Y}^{1/p^e} \leftarrow \mathbf{a}$ is a well-defined polynomial over $\mathbb{F}$. Denote this polynomial by $Q|_{\mathbf{a}}$, called the specialization of Q at $\mathbf{a}$.

For convenience, for a set $S \subseteq \mathbb{F}[\mathbf{Y}^{1/p^e}]$ and $\mathbf{a} \in \mathbb{F}^{3n}$, we say $\mathbf{a}$ is a *common non-root* of S if every $Q \in S$ is non-vanishing at $\mathbf{a}$.

We will construct a set $S^* \subseteq \mathbb{K}^{(e)}_{\leq D_0} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(d)$, where $D_0 = O_\delta(1)$, such that $\operatorname{res}(F)|_{\mathbf{a}} \neq 0$ for any common non-root $\mathbf{a}$ of S^*.

First, we construct S_0 by collecting the denominators of all coefficients of $g_1, g_2, u, h_1, \ldots, h_{k'}$, all elements of G, and all polynomials occurring in (22), (23), and (24), including the witness polynomials $t_{i,j}$, $s_{i,j}$, $a_{i,j,g}$, $b_{i,j,\ell}$, $c_{i,j,g}$, d_g, $e_{\ell,g}$ for all relevant indices. Then $|S_0| = O_\delta(d)$, and $S_0 \subseteq \mathbb{K}^{(e)}_{\leq D_0} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ for sufficiently large $D_0 = O_\delta(1)$.

We also need the following three claims, whose proofs are given in the appendix.

Claim 5.10. *There exists a set $S_1 \subseteq \mathbb{K}^{(e)}_{\leq D_1} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(1)$, where $D_1 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_1$, it holds that $1 \notin \widehat{\mathfrak{m}}|_{\mathbf{a}}$.*

Claim 5.11. *There exists a set $S_2 \subseteq \mathbb{K}^{(e)}_{\leq D_2} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(1)$, where $D_2 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_2$, the element $u|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ is a non-zerodivisor of $\mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$.*

Claim 5.12. *There exists a set $S_3 \subseteq \mathbb{K}^{(e)}_{\leq D_3} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(d)$, where $D_3 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_3$, the polynomial $\operatorname{res}(f_{i,j})|_{\mathbf{a}} \in \mathbb{F}[Z_1, Z_2]$ is well-defined for all $i \in \{0,1,2\}$ and $j \in [m_i]$, and the kernel of the $\mathbb{F}$-linear ring homomorphism $\mathbb{F}[Z_1, Z_2] \to \mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ sending Z_i to $g_i|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ contains $\langle \operatorname{res}(f_{0,1})|_{\mathbf{a}} \rangle$.*

Let $S^* = S_0 \cup S_1 \cup S_2 \cup S_3$. Then $|S^*| = O_\delta(d)$, and $S^* \subseteq \mathbb{K}^{(e)}_{\leq D_0} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ for sufficiently large $D_0 = O_\delta(1)$.

Step 6: Specialization. Consider a common non-root $\mathbf{a}$ of S^*. We will show that $\operatorname{res}(F)|_{\mathbf{a}} \neq 0$. The remainder of this step is devoted to proving this.

Define the ideal

$$\widehat{\mathfrak{m}}|_{\mathbf{a}} := \langle \{ g|_{\mathbf{a}} : g \in G \} \rangle$$

of $\mathbb{F}[\mathbf{T}]$. By (22), (23), and (24), and the fact that $S_0 \subseteq S^*$, we have

$$\left(t_{i,j}(\operatorname{res}(f_{i,j})(g_1, g_2))\right)|_{\mathbf{a}} = \left(u^{k_{i,j}} \left(1 + \sum_{g \in G} a_{i,j,g} g \right) + \sum_{\ell=1}^{k'} b_{i,j,\ell} h_\ell \right) \Bigg|_{\mathbf{a}}, \tag{25}$$

$$(s_{i,j}t_{i,j})|_{\mathbf{a}} = \left(1 + \sum_{g\in G} c_{i,j,g}g\right)\Bigg|_{\mathbf{a}}, \tag{26}$$

and

$$u|_{\mathbf{a}} \in \widehat{\mathfrak{m}}|_{\mathbf{a}} \quad \text{and} \quad \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \subseteq \widehat{\mathfrak{m}}|_{\mathbf{a}}. \tag{27}$$

Taking the product of the equations (25) over all $j \in [m_i]$ and applying (19), we obtain, for $i = 1, 2$,

$$\left(\left(\prod_{j\in[m_i]} t_{i,j}\right)\operatorname{res}(F_i)(g_1, g_2)\right)\Bigg|_{\mathbf{a}} - u^{n_i}|_{\mathbf{a}} \in u^{n_i}|_{\mathbf{a}} \cdot \widehat{\mathfrak{m}}|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle. \tag{28}$$

Let

$$\mathfrak{m}|_{\mathbf{a}} := \widehat{\mathfrak{m}}|_{\mathbf{a}} / \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle,$$

which is an ideal of $R_{\mathbf{a}} := \mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle$. By Claim 5.10, we have $1 \notin \mathfrak{m}|_{\mathbf{a}}$, i.e., $\mathfrak{m}|_{\mathbf{a}} \subsetneq R_{\mathbf{a}}$.

Let $\overline{u}|_{\mathbf{a}} := u|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \in R_{\mathbf{a}}$. By Claim 5.11, $\overline{u}|_{\mathbf{a}}$ is a non-zerodivisor of $R_{\mathbf{a}}$, and so are its powers. Therefore, for any integer $r \geq 0$, the $(R_{\mathbf{a}}/\mathfrak{m}|_{\mathbf{a}})$-module homomorphism

$$\begin{aligned} \times(\overline{u}|_{\mathbf{a}})^r : R_{\mathbf{a}}/\mathfrak{m}|_{\mathbf{a}} &\to \langle(\overline{u}|_{\mathbf{a}})^r\rangle / ((\overline{u}|_{\mathbf{a}})^r \mathfrak{m}|_{\mathbf{a}}) \\ x + \mathfrak{m}|_{\mathbf{a}} &\mapsto (\overline{u}|_{\mathbf{a}})^r x + (\overline{u}|_{\mathbf{a}})^r \mathfrak{m}|_{\mathbf{a}}. \end{aligned}$$

is injective. It is also obviously surjective. So $\times(\overline{u}|_{\mathbf{a}})^r$ is a $(R_{\mathbf{a}}/\mathfrak{m}|_{\mathbf{a}})$-module isomorphism.

As $1 \notin \mathfrak{m}|_{\mathbf{a}}$, using the above isomorphism, we see that $(\overline{u}|_{\mathbf{a}})^r$ is in $\langle(\overline{u}|_{\mathbf{a}})^r\rangle$ but not in $(\overline{u}|_{\mathbf{a}})^r \cdot \mathfrak{m}|_{\mathbf{a}}$ for $r \geq 0$.

Let $U \subseteq R_{\mathbf{a}}$ be the set of all finite products of elements in $\{t_{2,j}|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle : j \in [m_2]\}$. By (26), elements in U are already invertible modulo $\mathfrak{m}|_{\mathbf{a}}$. So localizing $R_{\mathbf{a}}/\mathfrak{m}|_{\mathbf{a}}$ and $\langle(\overline{u}|_{\mathbf{a}})^r\rangle / ((\overline{u}|_{\mathbf{a}})^r \mathfrak{m}|_{\mathbf{a}})$ with respect to the multiplicatively closed set U does not change these two modules. It follows that $(\overline{u}|_{\mathbf{a}})^r$ is not in the localization $U^{-1}((\overline{u}|_{\mathbf{a}})^r \mathfrak{m}|_{\mathbf{a}})$ for $r \geq 0$.

In particular, by (28), we have

$$\left(\left(\prod_{j\in[m_1]} t_{1,j}\right)\operatorname{res}(F_1)(g_1, g_2)\right)\Bigg|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \notin U^{-1}((\overline{u}|_{\mathbf{a}})^{n_1} \cdot \mathfrak{m}|_{\mathbf{a}}).$$

It follows that

$$\operatorname{res}(F_1)(g_1, g_2)|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \notin U^{-1}((\overline{u}|_{\mathbf{a}})^{n_1} \cdot \mathfrak{m}|_{\mathbf{a}}). \tag{29}$$

As $n_2 > n_1$ and $\overline{u}|_{\mathbf{a}} \in \mathfrak{m}|_{\mathbf{a}}$ (which holds by (27)), it also follows from (28) that

$$\left(\left(\prod_{j\in[m_2]} t_{2,j}\right)\operatorname{res}(F_2)(g_1, g_2)\right)\Bigg|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \in \langle(\overline{u}|_{\mathbf{a}})^{n_2}\rangle \subseteq (\overline{u}|_{\mathbf{a}})^{n_1} \cdot \mathfrak{m}|_{\mathbf{a}}.$$

Localizing with respect to U makes $\left(\prod_{j\in[m_2]} t_{2,j}\right)|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle$ invertible. Therefore,

$$\operatorname{res}(F_2)(g_1, g_2)|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \dots, h_{k'}|_{\mathbf{a}}\rangle \in U^{-1}((\overline{u}|_{\mathbf{a}})^{n_1} \cdot \mathfrak{m}|_{\mathbf{a}}). \tag{30}$$

It follows from (29) and (30) that

$$\operatorname{res}(F_1 + F_2)(g_1, g_2)|_{\mathbf{a}} \notin \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle . \tag{31}$$

By Claim 5.12, $\operatorname{res}(F_0)|_{\mathbf{a}}$, $\operatorname{res}(F_1)|_{\mathbf{a}}$, and $\operatorname{res}(F_2)|_{\mathbf{a}}$ are well-defined. Again by Claim 5.12, we have $\operatorname{res}(f_{0,1})(g_1, g_2)|_{\mathbf{a}} \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$. This together with (31) implies that

$$\operatorname{res}(F_1 + F_2)|_{\mathbf{a}} \notin \langle \operatorname{res}(f_{0,1})|_{\mathbf{a}} \rangle .$$

Therefore,

$$\operatorname{res}(F)|_{\mathbf{a}} = \operatorname{res}(F_0 + F_1 + F_2)|_{\mathbf{a}} \equiv \operatorname{res}(F_1 + F_2)|_{\mathbf{a}} \not\equiv 0 \pmod{\langle \operatorname{res}(f_{0,1})|_{\mathbf{a}} \rangle}.$$

In particular, $\operatorname{res}(F)|_{\mathbf{a}} \neq 0$.

Step 7: Concluding the hitting set construction. Write $q = p^e$, with $q = 1$ in characteristic zero. Since specialization assigns $\mathbf{a}$ to $\mathbf{Y}^{1/q}$, it assigns $\mathbf{a}^q$ to $\mathbf{Y}$, where powers are taken coordinatewise. Thus the preceding step shows that, for any common non-root $\mathbf{a}$ of S^*, the polynomial $F \circ L_{\mathbf{a},q}$ is nonzero, where $L_{\mathbf{a},q}(z_1, z_2) := \left(a_{0,j}^q + a_{1,j}^q z_1 + a_{2,j}^q z_2\right)_{j \in [n]}$.

Choose constants $M_\delta, C_\delta, D_0 = O_\delta(1)$, independent of F, such that $q \leq M_\delta$, $|S^*| \leq C_\delta d$, and every element of S^* has degree at most D_0 in $\mathbf{Y}^{1/q}$. Using [KS01], construct an ε-hitting set $\mathcal{H} \subseteq \mathbb{F}^{3n}$ for all nonzero $3n$-variate polynomials of degree at most D_0, where $\varepsilon < 1/(C_\delta d)$. Its size is $d^{O(1)} n^{O_\delta(1)}$. By the union bound, $\mathcal{H}$ contains a common non-root of S^*.

Let $\mathcal{Q}$ consist of all powers $p^r \leq M_\delta$ with $r \geq 0$; in characteristic zero, let $\mathcal{Q} = \{1\}$. Fix a set $T \subseteq \mathbb{F}$ of $d+1$ distinct elements and define

$$\mathcal{H}_0 = \bigcup_{q \in \mathcal{Q}} \bigcup_{\mathbf{a} \in \mathcal{H}} L_{\mathbf{a},q}(T^2).$$

For the value of q used in the preceding argument, some $\mathbf{a} \in \mathcal{H}$ makes $F \circ L_{\mathbf{a},q}$ a nonzero bivariate polynomial of degree at most d, so $T \times T$ hits it. Therefore $\mathcal{H}_0$ hits F. The construction is independent of F, and

$$|\mathcal{H}_0| \leq |\mathcal{Q}|\,|\mathcal{H}|\,(d+1)^2 = d^{O(1)} n^{O_\delta(1)} \leq (nd)^{O_\delta(1)}$$

as desired. □

Now we finish the proof of Theorem 5.4.

Proof of Theorem 5.4. As discussed at the beginning of Section 5.2, we may assume Assumptions (1)–(8). Under these assumptions, the statement follows from Lemma 5.8 and Lemma 5.9. The final hitting set is the union of the hitting set in Lemma 5.8 and the three hitting sets $\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2$ in Lemma 5.9. The resulting size remains $(nd)^{O_\delta(1)}$. □

6 Conclusions and Future Directions

In this paper, we gave a deterministic polynomial-time black-box PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$ circuits over arbitrary fields, under a squarefreeness assumption. Beyond the specific result, our main contribution is a new approach to PIT based on normalization of varieties, providing a field-independent alternative to the Sylvester–Gallai framework.

A central open problem is to remove the squarefreeness assumption and obtain PIT algorithms for the full class of $\Sigma^{[3]}\Pi\Sigma\Pi^{[\delta]}$ circuits over arbitrary fields. Our current proof uses normalization to recover factorization-like behavior after restricting to a hypersurface. In the general case, however, relevant multiplicity information may be lost upon passing to the reduced structure. Understanding how to retain or recover such higher-order information may be an important step toward further progress.

Another direction is to extend the normalization-based framework to larger top fan-in. It would also be interesting to investigate whether the techniques developed here can be combined with other structural ideas, including those arising from the Sylvester–Gallai-based literature.

AI disclosure. During revisions, we used AI to assist with proof checking and corrections, including the geometric-normality and specialization arguments. The authors take responsibility for all content.

References

[AF22] Robert Andrews and Michael A. Forbes. Ideals, determinants, and straightening: proving and using lower bounds for polynomial ideals. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*, pages 389–402. ACM, 2022. doi:10.1145/3519935.3520025.

[AH20] Tigran Ananyan and Melvin Hochster. Small subalgebras of polynomial rings and Stillman's conjecture. *Journal of the American Mathematical Society*, 33(1):291–309, 2020. doi:10.1090/jams/932.

[AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004. doi:10.4007/annals.2004.160.781.

[AL94] William Wells Adams and Philippe Loustaunau. *An Introduction to Gröbner Bases*, volume 3 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1994. doi:10.1090/gsm/003.

[AM69] Michael F. Atiyah and Ian G. Macdonald. *Introduction to Commutative Algebra*. Addison-Wesley, Reading, MA, 1969.

[AV08] Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pages 67–75. IEEE Computer Society, 2008. doi:10.1109/focs.2008.32.

[BM90] Peter Borwein and William O. J. Moser. A survey of Sylvester's problem and its generalizations. *Aequationes Mathematicae*, 40(1):111–135, 1990. doi:10.1007/bf02112289.

[BM93] Dave Bayer and David Mumford. What can be computed in algebraic geometry? arXiv preprint arXiv:alg-geom/9304003, 1993. URL: https://arxiv.org/abs/alg-geom/9304003, arXiv:alg-geom/9304003, doi:10.48550/arXiv.alg-geom/9304003.

[BSV20] Vishwas Bhargava, Shubhangi Saraf, and Ilya Volkovich. Deterministic factorization of sparse polynomials with bounded individual degree. *Journal of the ACM*, 67(2):8:1–8:28, 2020. doi:10.1145/3365667.

[BW98] Bruno Buchberger and Franz Winkler, editors. *Gröbner Bases and Applications*, volume 251 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1998. doi:10.1017/CBO9780511565847.

[CKS18] Chi-Ning Chou, Mrinal Kumar, and Noam Solomon. Hardness vs randomness for bounded depth arithmetic circuits. In Rocco A. Servedio, editor, *33rd Computational Complexity Conference (CCC 2018)*, volume 102 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:17, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2018.13.

[CRS95] Suresh Chari, Pankaj Rohatgi, and Aravind Srinivasan. Randomness-optimal unique element isolation with applications to perfect matching and related problems. *SIAM Journal on Computing*, 24(5):1036–1050, 1995. doi:10.1137/s0097539793250330.

[DDS21] Pranjal Dutta, Prateek Dwivedi, and Nitin Saxena. Deterministic identity testing paradigms for bounded top-fanin depth-4 circuits. In Valentine Kabanets, editor, *36th Computational Complexity Conference (CCC 2021)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:27, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2021.11.

[DG24] Pranjal Dutta and Sumanta Ghosh. Advances in polynomial identity testing. *ACM SIGACT News*, 55(2):54–88, 2024. Guest article in Complexity Theory Column 121. doi:10.1145/3674159.3674165.

[dJ98] Theo de Jong. An algorithm for computing the integral closure. *Journal of Symbolic Computation*, 26(3):273–277, 1998. doi:10.1006/jsco.1998.0211.

[DLL19] Jan Draisma, Michał Lasoń, and Anton Leykin. Stillman's conjecture via generic initial ideals. *Communications in Algebra*, 47(6):2384–2395, 2019. doi:10.1080/00927872.2019.1574806.

[DS07] Zeev Dvir and Amir Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM Journal on Computing*, 36(5):1404–1434, 2007. doi:10.1137/05063605x.

[DST24] Pranjal Dutta, Amit Sinhababu, and Thomas Thierauf. Derandomizing multivariate polynomial factoring for low degree factors. In Amit Kumar and Noga Ron-Zewi, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024)*, volume 317 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 75:1–75:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.APPROX/RANDOM.2024.75.

[Dub90] Thomas W. Dubé. The structure of polynomial ideals and Gröbner bases. *SIAM Journal on Computing*, 19(4):750–773, 1990. doi:10.1137/0219053.

[Eis95] David Eisenbud. *Commutative Algebra: With a View Toward Algebraic Geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. doi:10.1007/978-1-4612-5350-1.

[EK66] Michael Edelstein and Leroy M. Kelly. Bisecants of finite collections of sets in linear spaces. *Canadian Journal of Mathematics*, 18:375–380, 1966. doi:10.4153/cjm-1966-039-2.

[ESS19] Daniel Erman, Steven V. Sam, and Andrew Snowden. Big polynomial rings and Stillman's conjecture. *Inventiones mathematicae*, 218(2):413–439, 2019. doi:10.1007/s00222-019-00889-y.

[FGT21] Stephen Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in quasi-NC. *SIAM Journal on Computing*, 50(3):STOC16–218–STOC16–235, 2021. doi:10.1137/16M1097870.

[For78] David James Ford. *On the computation of the maximal order in a Dedekind domain*. PhD thesis, The Ohio State University, 1978.

[For24] Michael A. Forbes. Low-depth algebraic circuit lower bounds over any field. In Rahul Santhanam, editor, *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2024.31.

[GOPS23] Abhibhav Garg, Rafael Oliveira, Shir Peleg, and Akash Kumar Sengupta. Radical Sylvester–Gallai theorem for tuples of quadratics. In Amnon Ta-Shma, editor, *38th Computational Complexity Conference (CCC 2023)*, volume 264 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:30, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2023.20.

[GOS22] Abhibhav Garg, Rafael Oliveira, and Akash Kumar Sengupta. Robust radical Sylvester–Gallai theorem for quadratics. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2022.42.

[GOS25a] Abhibhav Garg, Rafael Oliveira, and Akash Kumar Sengupta. Rank bounds and PIT for depth-4 circuits with top fan-in 3 and constant bottom fan-in via a non-linear Edelstein–Kelly theorem. In *Proceedings of the 66th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2025)*, pages 2143–2169. IEEE Computer Society, 2025. doi:10.1109/focs63196.2025.00114.

[GOS25b] Abhibhav Garg, Rafael Oliveira, and Akash Kumar Sengupta. Uniform bounds on product Sylvester–Gallai configurations. In Oswin Aichholzer and Haitao Wang, editors,

41st International Symposium on Computational Geometry (SoCG 2025), volume 332 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 52:1–52:15, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2025.52.

[GOS+26] Abhibhav Garg, Rafael Oliveira, Akash Kumar Sengupta, Nir Shalmon, and Amir Shpilka. Rank bounds and polynomial-time PIT for $\Sigma^k\Pi\Sigma\Pi^2$ circuits. In Dana Moshkovitz, editor, *41st Computational Complexity Conference (CCC 2026)*, volume 383 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:18, Dagstuhl, Germany, 2026. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2026.17.

[GTZ88] Patrizia Gianni, Barry Trager, and Gail Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *Journal of Symbolic Computation*, 6(2–3):149–167, 1988. doi:10.1016/s0747-7171(88)80040-3.

[Guo24] Zeyu Guo. Variety evasive subspace families. *Computational Complexity*, 33(2):10, 2024. doi:10.1007/s00037-024-00256-1.

[Gup14] Ankit Gupta. Algebraic geometric techniques for depth-4 PIT & Sylvester–Gallai conjectures for varieties. Report TR14-130, Electronic Colloquium on Computational Complexity, 2014. URL: https://eccc.weizmann.ac.il/report/2014/130.

[GVJZ23] Zeyu Guo, Ben Lee Volk, Akhil Jalan, and David Zuckerman. Extractors for images of varieties. In *Proceedings of the 55th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2023)*, pages 46–59. ACM, 2023. doi:10.1145/3564246.3585109.

[Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983. doi:10.1016/0304-3975(83)90002-6.

[Her26] Grete Hermann. Die Frage der endlich vielen Schritte in der Theorie der Polynomideale. *Mathematische Annalen*, 95(1):736–788, 1926. doi:10.1007/bf01206635.

[Hir64a] Heisuke Hironaka. Resolution of singularities of an algebraic variety over a field of characteristic zero: I. *Annals of Mathematics*, 79(1):109–203, 1964. doi:10.2307/1970486.

[Hir64b] Heisuke Hironaka. Resolution of singularities of an algebraic variety over a field of characteristic zero: II. *Annals of Mathematics*, 79(2):205–326, 1964. doi:10.2307/1970547.

[HS80] Joos Heintz and Claus-Peter Schnorr. Testing polynomials which are easy to compute. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC 1980)*, pages 262–272. ACM, 1980. doi:10.1145/800141.804674.

[Kal95] Erich Kaltofen. Effective Noether irreducibility forms and applications. *Journal of Computer and System Sciences*, 50(2):274–295, 1995. doi:10.1006/jcss.1995.1023.

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1–2):1–46, 2004. doi:10.1007/s00037-004-0182-6.

[KRS24] Mrinal Kumar, Varun Ramanathan, and Ramprasad Saptharishi. Deterministic algorithms for low degree factors of constant depth circuits. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2024)*, pages 3901–3918. SIAM, 2024. doi:10.1137/1.9781611977912.137.

[KS01] Adam R. Klivans and Daniel Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 216–223. ACM, 2001. doi:10.1145/380752.380801.

[KS07] Neeraj Kayal and Nitin Saxena. Polynomial identity testing for depth 3 circuits. *Computational Complexity*, 16(2):115–138, 2007. doi:10.1007/s00037-007-0226-9.

[KS09] Neeraj Kayal and Shubhangi Saraf. Blackbox polynomial identity testing for depth 3 circuits. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 198–207. IEEE Computer Society, 2009. doi:10.1109/focs.2009.67.

[KSS14] Swastik Kopparty, Shubhangi Saraf, and Amir Shpilka. Equivalence of polynomial identity testing and deterministic multivariate polynomial factorization. In *Proceedings of the 29th IEEE Conference on Computational Complexity (CCC 2014)*, pages 169–180. IEEE Computer Society, 2014. doi:10.1109/ccc.2014.25.

[Kum13] Shrawan Kumar. Geometry of orbits of permanents and determinants. *Commentarii Mathematici Helvetici*, 88(3):759–788, 2013. doi:10.4171/cmh/302.

[Lan93] Serge Lang. *Algebra*. Addison-Wesley, Reading, MA, third edition, 1993.

[Lov79] László Lovász. On determinants, matchings, and random algorithms. In Lothar Budach, editor, *Fundamentals of Computation Theory (FCT 1979)*, pages 565–574, Berlin, 1979. Akademie-Verlag.

[LST24] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. *Communications of the ACM*, 67(2):101–108, 2024. doi:10.1145/3611094.

[Mag12] Arturo Magidin. Localisation is isomorphic to a quotient of polynomial ring. Mathematics Stack Exchange, 2012. Answer posted June 1, 2012; revised December 7, 2018. URL: https://math.stackexchange.com/a/152289.

[Mat89] Hideyuki Matsumura. *Commutative Ring Theory*, volume 8 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989. doi:10.1017/cbo9781139171762.

[MM82] Ernst W. Mayr and Albert R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46(3):305–329, 1982. doi:10.1016/0001-8708(82)90048-2.

[Mum88] David Mumford. *The Red Book of Varieties and Schemes*, volume 1358 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1988. doi:10.1007/978-3-662-21581-4.

[MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987. doi:10.1007/bf02579206.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994. doi:10.1016/s0022-0000(05)80043-1.

[OS22] Rafael Oliveira and Akash Kumar Sengupta. Radical Sylvester–Gallai theorem for cubics. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2022)*, pages 212–220. IEEE Computer Society, 2022. doi:10.1109/focs54457.2022.00027.

[OS24] Rafael Oliveira and Akash Kumar Sengupta. Strong algebras and radical Sylvester–Gallai configurations. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 95–105. ACM, 2024. doi:10.1145/3618260.3649617.

[PS21] Shir Peleg and Amir Shpilka. Polynomial time deterministic identity testing algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[2]}$ circuits via Edelstein–Kelly type theorem for quadratic polynomials. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pages 259–271. ACM, 2021. doi:10.1145/3406325.3451013.

[PS22] Shir Peleg and Amir Shpilka. A generalized Sylvester–Gallai-type theorem for quadratic polynomials. *Forum of Mathematics, Sigma*, 10:e112, 2022. doi:10.1017/fms.2022.100.

[Sax09] Nitin Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS*, 99:49–79, 2009.

[Sax14] Nitin Saxena. Progress on polynomial identity testing—II. In Manindra Agrawal and Vikraman Arvind, editors, *Perspectives in Computational Complexity: The Somenath Biswas Anniversary Volume*, volume 26 of *Progress in Computer Science and Applied Logic*, pages 131–146. Birkhäuser, Cham, 2014. doi:10.1007/978-3-319-05446-9_7.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980. doi:10.1145/322217.322225.

[Sei70] Abraham Seidenberg. Construction of the integral closure of a finite integral domain. *Rendiconti del Seminario Matematico e Fisico di Milano*, 40(1):101–120, 1970. doi:10.1007/bf02923228.

[Sei75] Abraham Seidenberg. Construction of the integral closure of a finite integral domain. II. *Proceedings of the American Mathematical Society*, 52(1):368–372, 1975. doi:10.1090/s0002-9939-1975-0424783-5.

[Shp20] Amir Shpilka. Sylvester–Gallai type theorems for quadratic polynomials. *Discrete Analysis*, 2020(13):1–34, 2020. doi:10.19086/da.14492.

[SS12] Nitin Saxena and Comandur Seshadhri. Blackbox identity testing for bounded top-fanin depth-3 circuits: The field doesn't matter. *SIAM Journal on Computing*, 41(5):1285–1298, 2012. doi:10.1137/10848232.

[SS13] Nitin Saxena and Comandur Seshadhri. From Sylvester–Gallai configurations to rank bounds: Improved blackbox identity test for depth-3 circuits. *Journal of the ACM*, 60(5):33:1–33:33, 2013. doi:10.1145/2528403.

[ST17] Ola Svensson and Jakub Tarnawski. The matching problem in general graphs is in quasi-NC. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2017)*, pages 696–707. IEEE Computer Society, 2017. doi:10.1109/focs.2017.70.

[Sta26] The Stacks project authors. The stacks project. https://stacks.math.columbia.edu, 2026.

[Sto68] Gabriel Stolzenberg. Constructive normalization of an algebraic variety. *Bulletin of the American Mathematical Society*, 74(3):595–599, 1968. doi:10.1090/s0002-9904-1968-12023-3.

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3–4):207–388, 2010. doi:10.1561/0400000039.

[Tra84] Barry Marshall Trager. *Integration of algebraic functions*. PhD thesis, Massachusetts Institute of Technology, 1984.

[Vak24] Ravi Vakil. The rising sea: Foundations of algebraic geometry. https://math.stanford.edu/~vakil/216blog/FOAGjul2724public.pdf, 2024. Draft, July 27, 2024 version.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation (EUROSAM '79)*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226, Berlin, 1979. Springer-Verlag. doi:10.1007/3-540-09519-5_73.

[ZS75] Oscar Zariski and Pierre Samuel. *Commutative Algebra I*, volume 28 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1975.

A Proof of Lemma 3.6

We now prove Lemma 3.6, which we restate below for convenience.

Lemma 3.6. *Let $f \in \mathbb{K}_{\le d'}[\mathbf{X}]_{\le d}$ be a nonzero polynomial in n variables. Let g be a factor of f. Then there exists $c \in \mathbb{K}^\times$ such that $cg \in \mathbb{K}_{\le d''}[\mathbf{X}]_{\le d}$ for some $d'' = O_{d,d',n}(1)$.*

First, we introduce some notation. Let $R = \mathbb{F}[\mathbf{Y}] \subseteq \mathbb{K} = \mathbb{F}(\mathbf{Y})$. For an irreducible polynomial $a \in R$ and $b \in \mathbb{K}^\times$, define $\mathrm{ord}_a(b)$ to be the unique integer r such that b can be written as $b = a^r \frac{s}{t}$, where $s \in R$ and $t \in R \setminus \{0\}$ are not divisible by a. For a univariate polynomial $f \in \mathbb{K}[X] \setminus \{0\}$ and an irreducible polynomial $a \in R$ in $\mathbf{Y}$, define $\mathrm{ord}_a(f) := \min_c(\mathrm{ord}_a(c))$, where c ranges over the nonzero coefficients of f. Finally, denote by $\mathrm{cont}(f)$ the *content* of f, defined as

$$\mathrm{cont}(f) := \prod_a a^{\mathrm{ord}_a(f)}$$

with the product taken over the irreducible polynomials $a \in R$ such that $\mathrm{ord}_a(f) \neq 0$. If two such polynomials differ by a unit factor, only one representative is included in the product. Note that $\mathrm{cont}(f)$ is well-defined up to multiplication by a unit of R. Also note that by definition, $f \in R[X]$ as long as $\mathrm{cont}(f) \in R$.

We need the following version of Gauss's lemma, which holds more generally when R is a UFD [Lan93, Theorem 2.1].

Lemma A.1 (Gauss's lemma)**.** $\mathrm{cont}(gh) = \mathrm{cont}(g)\mathrm{cont}(h)$ *for nonzero univariate polynomials* $g, h \in \mathbb{K}[X]$.

Proof of Lemma 3.6. Consider the Kronecker map $\phi : \mathbb{K}[X_1, \dots, X_n] \to \mathbb{K}[X]$, which is the $\mathbb{K}$-linear ring homomorphism sending X_i to $X^{(d+1)^{i-1}}$ for i. We have $\phi(f) = \phi(g) \cdot \phi(f/g)$. The terms of f (resp. g) correspond one-to-one to the terms of $\phi(f)$ (resp. $\phi(g)$), preserving coefficients. And the degree of $\phi(f)$ is at most $(d+1)^n = O_{d,n}(1)$. This reduces the problem to the univariate case.

Now we prove the lemma for the univariate case. Let $f \in \mathbb{K}_{\le d'}[X]_{\le d}$ be a nonzero polynomial. Recall that $R = \mathbb{F}[\mathbf{Y}]$. Let s be the least common multiple of all the denominators of the nonzero coefficients of f, so that $sf \in R[X]$. Let $t = \mathrm{cont}(sf) \in R \setminus \{0\}$. Then $\mathrm{cont}(\frac{s}{t}f) = 1$. Note that the coefficients of $\frac{s}{t}f$ are polynomials in $\mathbb{F}[\mathbf{Y}]$ of degree at most $d'' := (d+1)d'$. Let $c = \frac{1}{\mathrm{cont}(g)} \in \mathbb{K}^\times$. Then $\mathrm{cont}(cg) = 1$. Let $h = \frac{s}{ct} \cdot \frac{f}{g}$, so that $\frac{s}{t}f = (cg)h$. As $\mathrm{cont}(\frac{s}{t}f) = 1$ and $\mathrm{cont}(cg) = 1$, we have $\mathrm{cont}(h) = 1$ by Lemma A.1. In particular, $\frac{s}{t}f$, cg, and h are all in $R[X]$.

Write $cg = \sum_{i=0}^{\deg(cg)} a_i X^i$ and $h = \sum_{j=0}^{\deg(h)} b_j X^j$ with $a_i, b_j \in R$. Let $D = \max_{0 \le i \le \deg(cg)} \deg_{\mathbf{Y}}(a_i)$ and $D' = \max_{0 \le j \le \deg(h)} \deg_{\mathbf{Y}}(b_j)$. Choose the maximal $i_0 \in \{0, 1, \dots, \deg(cg)\}$ such that a_{i_0} has a nonzero degree-D homogeneous component, and let H be this homogeneous component. Similarly, choose the maximal $j_0 \in \{0, 1, \dots, \deg(h)\}$ such that b_{j_0} has a nonzero degree-D' homogeneous component, and let H' be this homogeneous component. By the maximality of D, D', i_0, and j_0, the degree-$(D + D')$ homogeneous component of the coefficient of $X^{i_0+j_0}$ in $(cg) \cdot h = \frac{s}{t}f$ is exactly $H \cdot H' \neq 0$. However, we know that the coefficients of $\frac{s}{t}f$ all have degrees at most d''. So

$$D + D' \le d''.$$

Therefore, the coefficients of cg are all polynomials in $\mathbb{F}[\mathbf{Y}]$ of degree at most $D \le d''$. In particular, $cg \in \mathbb{K}_{\le d''}[X]_{\le d}$, where $d'' = (d+1)d' \in O_{d,d'}(1)$. □

B Proofs of Other Claims

We now prove Claim 5.10, Claim 5.11, and Claim 5.12.

Claim 5.10. *There exists a set $S_1 \subseteq \mathbb{K}^{(e)}_{\leq D_1} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(1)$, where $D_1 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_1$, it holds that $1 \notin \widehat{\mathfrak{m}}|_{\mathbf{a}}$.*

Proof. Let $D = O_\delta(1)$ bound the degrees of the polynomials in G. By Lemma 3.9, we can choose $d_0 = O_\delta(1)$ such that, for every common non-root $\mathbf{a}$ of S_0, if $1 \in \widehat{\mathfrak{m}}|_{\mathbf{a}}$, then $1 = \sum_{g \in G} h_g(g|_{\mathbf{a}})$ holds for some $h_g \in \mathbb{F}[\mathbf{T}]$ of degree at most d_0. This choice of d_0 works for all $\mathbf{a}$, since specialization does not increase the degrees or the number of generators.

Construct a matrix M over $\mathbb{K}^{(e)}$ as follows. Its rows are indexed by pairs (μ, g), where $g \in G$ and μ is a monomial in $\mathbf{T}$ of degree at most d_0. Its columns are indexed by all monomials in $\mathbf{T}$ of degree at most $d_0 + D$. The row indexed by (μ, g) is the coefficient vector of μg. Let M^+ be obtained from M by adjoining the coefficient vector of the constant polynomial 1.

For any common non-root $\mathbf{a}$ of S_0, define $M|_{\mathbf{a}}$ and $M^+|_{\mathbf{a}}$ by specializing their entries, with the same row and column indices. Thus the rows of $M|_{\mathbf{a}}$ are precisely the coefficient vectors of $\mu(g|_{\mathbf{a}})$, including any zero or repeated rows. By the choice of d_0, we have

$$1 \in \widehat{\mathfrak{m}}|_{\mathbf{a}} \quad \Longleftrightarrow \quad \operatorname{rank}(M^+|_{\mathbf{a}}) = \operatorname{rank}(M|_{\mathbf{a}}).$$

Let $r = \operatorname{rank}(M)$. Since $1 \notin \widehat{\mathfrak{m}} = \langle G \rangle$, the adjoined row does not belong to the row span of M. Hence $\operatorname{rank}(M^+) = r + 1$. Choose an $(r+1) \times (r+1)$ submatrix N of M^+ with nonzero determinant, and write $\det(N) = \frac{Q}{P}$, where $P, Q \in \mathbb{F}[\mathbf{Y}^{1/p^e}] \setminus \{0\}$. The dimensions of M^+ and the complexities of its entries are $O_\delta(1)$, so we may choose P and Q of degree $O_\delta(1)$. Set $S_1 = \{P, Q\}$.

Now let $\mathbf{a}$ be a common non-root of $S_0 \cup S_1$. Then

$$\det(N|_{\mathbf{a}}) = \frac{Q(\mathbf{a})}{P(\mathbf{a})} \neq 0,$$

and therefore $\operatorname{rank}(M^+|_{\mathbf{a}}) \geq r + 1$. On the other hand, all $(r+1) \times (r+1)$ minors of M vanish identically and remain zero after specialization. Consequently, $\operatorname{rank}(M|_{\mathbf{a}}) \leq r$. Thus the two specialized matrices have different ranks, which implies $1 \notin \widehat{\mathfrak{m}}|_{\mathbf{a}}$. The required size and degree bounds for S_1 follow from its construction. □

Claim 5.11. *There exists a set $S_2 \subseteq \mathbb{K}^{(e)}_{\leq D_2} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(1)$, where $D_2 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_2$, the element $u|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ is a non-zerodivisor of $\mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$.*

Proof. The key observation is that it suffices to consider zerodivisor witnesses of bounded degree. Indeed, $u|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ is a zerodivisor of $\mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle$ if and only if there exists $v \in \mathbb{F}[\mathbf{T}]$ such that

$$u|_{\mathbf{a}} \cdot v \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle \quad \text{but} \quad v \notin \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle .$$

The set of v satisfying the first condition is the ideal quotient $(\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle : \langle u|_{\mathbf{a}} \rangle)$, whose Gröbner basis can be computed using [AL94, Lemmas 2.3.10 and 2.3.11]. By the degree bounds for the Gröbner bases involved, if such a v exists, then one exists with $\deg_{\mathbf{T}}(v) \leq d_0$ for some $d_0 = O_\delta(1)$.

Similarly, $u + \langle h_1, \ldots, h_{k'} \rangle$ is a zerodivisor of $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle$ if and only if there exists $v \in \mathbb{K}^{(e)}[\mathbf{T}]$ such that

$$u \cdot v \in \langle h_1, \ldots, h_{k'} \rangle \quad \text{but} \quad v \notin \langle h_1, \ldots, h_{k'} \rangle,$$

and if such a v exists, then one exists with $\deg_{\mathbf{T}}(v) \leq d_0$.

Let W be the $\mathbb{K}^{(e)}$-vector space of polynomials in $\mathbb{K}^{(e)}[\mathbf{T}]$ of degree at most d_0. Define

$$W_1 := \{v \in W : u \cdot v \in \langle h_1, \ldots, h_{k'} \rangle\}, \qquad W_2 := \{v \in W : v \in \langle h_1, \ldots, h_{k'} \rangle\}.$$

Similarly, for a specialization $\mathbf{a}$, let $W|_{\mathbf{a}}$ be the $\mathbb{F}$-vector space of polynomials in $\mathbb{F}[\mathbf{T}]$ of degree at most d_0, and define

$$W_1|_{\mathbf{a}} := \{v \in W|_{\mathbf{a}} : u|_{\mathbf{a}} \cdot v \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle\}, \quad W_2|_{\mathbf{a}} := \{v \in W|_{\mathbf{a}} : v \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}} \rangle\}.$$

Clearly $W_2 \subseteq W_1$. Since $\mathbb{K}^{(e)}[\mathbf{T}]/\langle h_1, \ldots, h_{k'} \rangle$ is an integral domain and $u + \langle h_1, \ldots, h_{k'} \rangle \neq 0$, the element $u + \langle h_1, \ldots, h_{k'} \rangle$ is a non-zerodivisor. Hence $W_1 \subseteq W_2$, and therefore

$$W_1 = W_2. \tag{32}$$

We now express W_1 and W_2 by finite systems of linear equations. Choose $d_1 = O_\delta(1)$, using Lemma 3.9, uniformly for the identities $uv = \sum_i w_i h_i$ and $v = \sum_i w_i h_i$, both over $\mathbb{K}^{(e)}$ and after every specialization at a common non-root of S_0. This is possible because the numbers of variables and generators and all relevant degrees are bounded in terms of δ. So for $v \in W_1$, there exist $w_1, \ldots, w_{k'} \in \mathbb{K}^{(e)}[\mathbf{T}]_{\leq d_1}$ with

$$uv = \sum_{i=1}^{k'} w_i h_i.$$

Viewing the coefficients of $v, w_1, \ldots, w_{k'}$ as unknowns, this identity becomes a homogeneous linear system over $\mathbb{K}^{(e)}$. Let M_1 be its coefficient matrix. Computing a basis of the solution space of this system, projecting each solution to the coordinates corresponding to v, and selecting a linearly independent subset if necessary yields a basis of W_1. Let B_1 be the matrix whose rows are these basis vectors.

Similarly, viewing the coefficients of $v, w_1, \ldots, w_{k'}$ as unknowns, the identity

$$v = \sum_{i=1}^{k'} w_i h_i$$

gives another homogeneous linear system over $\mathbb{K}^{(e)}$. Let M_2 denote its coefficient matrix. Computing a basis of the solution space of this system, projecting each solution to the coordinates corresponding to v, and selecting a linearly independent subset if necessary yields a basis of W_2. Let B_2 be the matrix whose rows are these basis vectors.

By (32), the row spans of B_1 and B_2 are equal. We now define S_2. Add to S_2 all denominators of the entries of M_1, M_2, B_1, B_2 and all intermediate matrices used to compute and project the kernel bases. Also add the numerators and denominators of the nonzero pivot minors used in these Gaussian eliminations and in selecting the bases B_1, B_2. Finally, add the denominators of the coefficients expressing each row of B_1 as a linear combination of the rows of B_2, and vice versa. Since all the linear systems above have size $O_\delta(1)$ and all their entries have complexity $O_\delta(1)$, we have $S_2 \subseteq \mathbb{K}^{(e)}_{\leq D_2} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$, where $|S_2| = O_\delta(1)$ and $D_2 = O_\delta(1)$.

Now let $\mathbf{a}$ be a common non-root of $S_0 \cup S_2$. The conditions imposed by S_2 ensure that all the matrices above specialize well and that the same Gaussian eliminations remain valid after specialization. Therefore $B_1|_{\mathbf{a}}$ and $B_2|_{\mathbf{a}}$ are bases of $W_1|_{\mathbf{a}}$ and $W_2|_{\mathbf{a}}$, respectively. Moreover, the linear relations expressing the rows of B_1 in terms of the rows of B_2, and conversely, also specialize well. Hence the row spans of $B_1|_{\mathbf{a}}$ and $B_2|_{\mathbf{a}}$ are equal, i.e., $W_1|_{\mathbf{a}} = W_2|_{\mathbf{a}}$.

Suppose, for contradiction, that $u|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ is a zerodivisor in $\mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$. By the bounded-degree reduction above, there exists $v \in W|_{\mathbf{a}}$ such that $u|_{\mathbf{a}} \cdot v \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ but $v \notin \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$, or equivalently, $v \in W_1|_{\mathbf{a}}$ but $v \notin W_2|_{\mathbf{a}}$. But this contradicts $W_1|_{\mathbf{a}} = W_2|_{\mathbf{a}}$. □

Claim 5.12. *There exists a set $S_3 \subseteq \mathbb{K}^{(e)}_{\leq D_3} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ of size $O_\delta(d)$, where $D_3 = O_\delta(1)$, such that for any common non-root $\mathbf{a} \in \mathbb{F}^{3n}$ of $S_0 \cup S_3$, the polynomial $\mathrm{res}(f_{i,j})|_{\mathbf{a}} \in \mathbb{F}[Z_1, Z_2]$ is well-defined for all $i \in \{0,1,2\}$ and $j \in [m_i]$, and the kernel of the $\mathbb{F}$-linear ring homomorphism $\mathbb{F}[Z_1, Z_2] \to \mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ sending Z_i to $g_i|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ contains $\langle \mathrm{res}(f_{0,1})|_{\mathbf{a}}\rangle$.*

Proof. First, for all $i \in \{0,1,2\}$ and $j \in [m_i]$, add to S_3 the denominators of all coefficients of $\mathrm{res}(f_{i,j})$. This ensures that the polynomials $\mathrm{res}(f_{i,j})|_{\mathbf{a}}$ are well-defined.

As $\mathrm{res}(f_{0,1})(g_1, g_2) \in \langle h_1, \ldots, h_{k'}\rangle$, by Lemma 3.11, we may write

$$\mathrm{res}(f_{0,1})(g_1, g_2) = \sum_{i=1}^{k'} q_i h_i,$$

where $q_i \in \mathbb{K}^{(e)}_{\leq d_0'}[\mathbf{T}]_{\leq d_0}$ for some $d_0, d_0' \in O_\delta(1)$. Add to S_3 the denominators of all the coefficients of $q_1, \ldots, q_{k'}$. Then $\mathrm{res}(f_{0,1})(g_1, g_2)|_{\mathbf{a}} = \sum_{i=1}^{k'} q_i|_{\mathbf{a}} h_i|_{\mathbf{a}}$. So $\mathrm{res}(f_{0,1})(g_1, g_2)|_{\mathbf{a}} \in \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$. It follows that the $\mathbb{F}$-linear ring homomorphism $\mathbb{F}[Z_1, Z_2] \to \mathbb{F}[\mathbf{T}]/\langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ sending Z_i to $g_i|_{\mathbf{a}} + \langle h_1|_{\mathbf{a}}, \ldots, h_{k'}|_{\mathbf{a}}\rangle$ maps $\mathrm{res}(f_{0,1})|_{\mathbf{a}}$ to zero, as desired. And by construction, we have $S_3 \subseteq \mathbb{K}^{(e)}_{\leq D_3} \cap \mathbb{F}[\mathbf{Y}^{1/p^e}]$ and $|S_3| = O_\delta(d)$, where $D_3 = O_\delta(1)$. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il