

Optimal Pseudorandom Generators for Low-Degree Polynomials over Moderately Large Fields

Ashish Dwivedi ✉ 🏠 

Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA

Zeyu Guo ✉ 🏠 

Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA

Ben Lee Volk ✉ 🏠 

Efi Arazi School of Computer Science, Reichman University, Israel

Abstract

We construct explicit pseudorandom generators that fool n -variate polynomials of degree at most d over a finite field \mathbb{F}_q . The seed length of our generators is $O(d \log n + \log q)$, over fields of size exponential in d and characteristic at least $d(d-1)+1$. Previous constructions such as Bogdanov’s (STOC 2005) and Derksen and Viola’s (FOCS 2022) had either suboptimal seed length or required the field size to depend on n .

Our approach follows Bogdanov’s paradigm while incorporating techniques from Lecerf’s factorization algorithm (J. Symb. Comput. 2007) and insights from the construction of Derksen and Viola regarding the role of indecomposability of polynomials.

2012 ACM Subject Classification Theory of computation → Pseudorandomness and derandomization

Keywords and phrases Pseudorandom Generators, Low Degree Polynomials

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.44

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2402.11915>

Funding *Ben Lee Volk*: The research leading to these results has received funding from the Israel Science Foundation (grant number 843/23).

Acknowledgements We thank Jesse Goodman and Pooya Hatami for helpful discussions. Part of this work was carried out while the first two authors were visiting the Simons Institute for the Theory of Computing at UC Berkeley. We thank the institute for its support and hospitality.

1 Introduction

The role of randomness in efficient computation is one of the central topics in complexity theory: random bits are useful for designing algorithms, but producing random bits comes at a cost and it is often desirable to reduce them or eliminate them altogether. One of the simplest yet most profound insights in this area is that efficient algorithms are, by definition, computationally limited, and cannot perform arbitrary statistical tests over their random bits. Therefore, one may hope to construct *pseudorandom* distributions that use less random bits but are able to “fool” some limited classes of tests, that cannot distinguish between them and between truly random bits.

For the pseudorandom distributions to be useful, they need to be efficiently computable themselves. This is usually modeled as a *pseudorandom generator* (PRG, for short). A PRG for a class of \mathcal{C} is an efficiently computable function $G : S \rightarrow B$ such that for every function $f \in \mathcal{C}$, the distributions $f(\mathbf{U}_B)$ and $f(G(\mathbf{U}_S))$ are close in statistical distance, where \mathbf{U}_A denotes the uniform distribution over the set A . Namely, the two experiments of applying $f(\cdot)$ to a uniformly random element of B , and applying $f(G(\cdot))$ to a uniformly random element



© Ashish Dwivedi, Zeyu Guo, and Ben Lee Volk;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 44; pp. 44:1–44:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of S , give roughly the same results. For this to be useful and non-trivial, obviously the set S needs to be significantly smaller than B . The quantity $\log |S|$ is called the *seed length* of the generator.

There has been a significant amount of work on constructing pseudorandom generators for various types of restricted distinguishers. In its most general form, where the distinguisher is allowed to be an arbitrary efficient (even non-uniform) algorithm, constructing such PRGs would imply breakthrough lower bounds in complexity theory. However, there are also unconditional constructions of PRGs for distinguishers coming from certain smaller complexity classes (see, for example, the surveys [34, 17]).

In this paper, we focus on pseudorandom generators in the algebraic setting. Here, the restriction on the distinguishers is of algebraic nature: we seek to fool distinguishers that are low-degree n -variate polynomials over finite fields.

The problem of fooling low-degree polynomials is well-studied. The most basic case is polynomials of degree one, i.e., fooling linear functions. Such generators are also known as ε -*biased sets*, and this problem was traditionally studied over \mathbb{F}_2 , although some of the constructions can be generalized to larger fields. This concept was first defined and considered by Naor and Naor [24], with various improved constructions given by [2, 14, 4], culminating in a recent nearly-optimal construction by Ta-Shma [33]. The seed length in those constructions is $O(\log n + \log q + \log(1/\varepsilon))$, where n denotes the number of variables, ε the error of the PRG, and q the field size.

While focusing on polynomials of degree one might seem a bit too restrictive, ε -biased sets have found numerous applications throughout the field of pseudorandomness and derandomization, and in the theory of computation in general.

One example relevant to this work is that ε -biased sets are in fact a basic building block in a construction of PRGs for higher-degree polynomials, using a paradigm initiated by Bogdanov and Viola [7]. They suggested constructing a generator for degree- d polynomials by summing up $\ell = \ell(d)$ independent copies of a generator for degree-one polynomials. The paper [7] proved a conditional result when the number of summands is d , assuming certain additive combinatorics conjectures. Lovett [22] showed how to prove an unconditional result at the cost of making the number of summands 2^d . Finally, Viola [35] showed (unconditionally) that in fact d summands suffice. The seed length in his construction is $O(d \log n + d2^d \log(q/\varepsilon))$. Indeed, even though the construction only sums d copies of a generator for degree-one polynomials, for the analysis to go through, the error of this generator needs to be as small as ε^{2^d} (for the final error of the generator for degree- d polynomials to be ε), which incurs a factor of 2^d in the final seed length. Improving this generator and in particular obtaining meaningful results for polynomials of degree greater than $\log n$ is an extremely important open problem in complexity theory. One reason is that such pseudorandom generators will yield pseudorandom generators for small constant-depth circuits with parity gates, since Razborov [25] and Smolensky [32] famously proved that functions computed by such circuits are approximated by low-degree polynomials.

All the constructions mentioned above work for any field. There are, however, better results when the field size q is assumed to be large (typically, at least polynomially large in d and $1/\varepsilon$). This assumption is useful since it allows one to use powerful tools from algebraic geometry, such as Weil-type estimates [36] on the number of points of varieties over finite fields.

This line of work was initiated by Bogdanov [6], who showed how to use different pseudorandom objects called *hitting set generators* for low-degree polynomials in order to construct pseudorandom generators. Bogdanov's work, followed by the later improved

constructions of hitting set generators [19, 23, 10, 16], resulted in a PRG with seed length $O(d^4 \log n + \log q)$ assuming $q \geq Cd^6/\varepsilon^2$ for a sufficiently large constant C .¹ Over fields of large enough characteristic, combining Bogdanov’s paradigm with the results of Lecerf [21] results in improving the d^4 factor to d^2 . We expand more on Bogdanov’s and Lecerf’s techniques in Section 1.2, as both are very relevant to this work.

More recently, Derksen and Viola [12] introduced fundamentally new techniques for this problem, with tools coming from invariant theory. One of their key ideas is to construct a low-degree polynomial map on a few variables that preserves the *indecomposability* property of a polynomial f when composed with it (we refer to Section 2 for more on that). Using this new tool in conjunction with other techniques, they are able to construct generators with seed length $O(d \log(dn) + \log q)$, assuming $q \geq Cd^4 n^\delta/\varepsilon^2$ (for some large constant C and small constant δ), or seed length $O(d \log n \log(d \log n) + \log q)$ for $q \geq C(d \log n)^4/\varepsilon^2$. One should note, however, that the optimal parameters in the construction of [12] are obtained after composing their construction with Bogdanov’s original construction.

A related natural question is how small the seed length can potentially be. Alon, Ben-Eliezer and Krivelevich [1] considered this question and proved a lower bound of $\Omega(d \log(n/d) + \log q + \log(1/\varepsilon))$ on the seed length. Thus, we see that the explicit constructions of [12] come very close to the optimal bound. However, unlike the result of Bogdanov [6], in the construction of Derksen and Viola [12] the minimum field size depends on the number of variables n .

1.1 Our Results

In this paper, we provide an improved construction of PRGs for low-degree polynomials, with an even shorter seed length, assuming the field size is exponentially large in d (but independent of n).

► **Theorem 1.1.** Let \mathbb{F}_q be a finite field of characteristic at least $d(d-1)+1$ and size $q \geq C(d2^d/\varepsilon + d^4/\varepsilon^2)$ (for some sufficiently large absolute constant C). Then, there exists an explicit pseudorandom generator that fools n -variate polynomials of degree at most d over \mathbb{F}_q with error ε and seed length $O(d \log n + \log q)$.

For convenience, we summarize the comparison between Theorem 1.1 and the results of Bogdanov [6], Viola [35] and Derksen and Viola [12] in the following table. All the entries in this table are given up to some constant factors, but for ease of readability, we omit $O(\cdot)$ notations.

| | Seed Length | Field Size | Characteristic |
|---------------|--|--|-----------------|
| [35] | $d \log n + d \cdot 2^d \log(q/\varepsilon)$ | Every $q \geq 2$ | Any |
| [6]+[16] | $d^4 \log n + \log q$ | d^6/ε^2 | Any |
| [6]+[16]+[21] | $d^2 \log n + \log q$ | d^6/ε^2 | $\geq d(d-1)+1$ |
| [12] | $d \log(n) + \log q$ | $d^4 n^{0.001}/\varepsilon^2$ | Any |
| [12] | $d \log n \cdot \log(d \log n) + \log q$ | $(d \log n)^4/\varepsilon^2$ | Any |
| This paper: | $d \log n + \log q$ | $d2^d/\varepsilon + d^4/\varepsilon^2$ | $\geq d(d-1)+1$ |

We also prove that, if we only want to fool polynomials of *prime* degree up to d , then the required field size in Theorem 1.1 can be improved to $O(d^4/\varepsilon^2)$, avoiding an exponential dependence on d .

¹ One should note that since q is polynomially large in $1/\varepsilon$, the seed length also implicitly depends on $\log(1/\varepsilon)$ through the $\log q$ term.

► **Theorem 1.2.** Let \mathbb{F}_q be a finite field of characteristic at least $d(d-1)+1$ and size $q \geq C(d^4/\varepsilon^2)$ (for some sufficiently large absolute constant C). Then, there exists an explicit pseudorandom generator that fools n -variate polynomials of prime degree up to d over \mathbb{F}_q with error ε and seed length $O(d \log n + \log q)$.

1.2 Proof Techniques

We start by reviewing the proof of Bogdanov’s PRG. Bogdanov’s idea is to consider restrictions of the polynomial f we are trying to fool onto planes, and to argue that “most” planes preserve the output distribution of the polynomial. Since a plane is a two-dimensional subspace, after having selected a good plane, we only need to sample two more field elements to select a random element from the plane.

The question now is how to find a good plane. Here, Bogdanov uses results by Kaltofen [18], who proved an effective version of Hilbert’s irreducibility theorem. Kaltofen demonstrated that, for every degree- d irreducible polynomial f , there exists a polynomial P of degree roughly d^4 , whose variables correspond to the parameters of the plane, such that every point at which P is nonzero corresponds to a “good” plane for f – that is, a plane that preserves the irreducibility of f . Therefore, we can use a hitting set generator for polynomials of degree d^4 to find a good plane. This results in a factor of d^4 in the final seed length.

Over fields of large characteristic (or characteristic zero), Lecerf [20, 21] obtained an improved upper bound of $O(d^2)$ on the degree of such polynomials, based on ideas of Ruppert [26, 27] and Gao [15]. However, Lecerf also presents an example where the degree of such a polynomial must be at least $\Omega(d^2)$, demonstrating that this approach alone may not suffice to achieve an improvement within Bogdanov’s framework.

Derksen and Viola circumvent this problem by using a different approach: rather than preserving irreducibility (or more generality, if a polynomial is reducible, preserving its number of irreducible factors) as in Bogdanov’s approach, they construct a map that preserves the *indecomposability* of polynomials: a polynomial $f(x_1, \dots, x_n)$ is indecomposable if it cannot be written as $f = g(h(x_1, \dots, x_n))$ where g is a univariate polynomial of degree at least 2. Over large fields, irreducibility and indecomposability both guarantee that the polynomial is roughly equidistributed, and in fact these notions are tightly connected. We refer to Section 2 for the precise definitions and description of those connections.

To prove our result, we revisit Bogdanov’s approach, while noting that by applying Lecerf’s results [21] (rather than Kaltofen’s bounds [18]) in a more careful way, and assuming the field is sufficiently large, it is in fact enough to use hitting sets generators only for polynomials of degree $O(d)$, rather than $O(d^2)$. We also incorporate the insights from [12] regarding indecomposability and equidistribution. This requires following the outline above but making sure that at each step, we only need to hit polynomials of degree $O(d)$.

Following Lecerf’s notation and terminology, suppose $F(x_1, \dots, x_n, y)$ is an irreducible polynomial of degree at most d . A point $\mathbf{a} = (a_1, \dots, a_n)$ is called a *Bertinian good* point if the bivariate polynomial $H(x, y) = F(a_1x, \dots, a_nx, y)$ remains irreducible. Lecerf proves that there exists a polynomial $\mathcal{A}(z_1, \dots, z_n)$ of degree $O(d^2)$ such that if $\mathcal{A}(a_1, \dots, a_n) \neq 0$ then \mathbf{a} is a Bertinian good point. This is achieved by transforming the question of irreducibility into a question about the rank of a solution space for a certain linear system that depends on a_1, \dots, a_n (see Section 4). This transformation naturally leads to defining \mathcal{A} as a certain minor of the matrix representing that linear system. The minor has dimensions $O(d) \times O(d)$, and each entry of the matrix is a polynomial of degree $O(d)$, which results in a total bound of $O(d^2)$ on the degree of its determinant, \mathcal{A} .

We, however, observe that Lecerf’s results actually imply a much stronger structure of this linear system: its solution space, over any field, is always spanned by vectors whose entries are in $\{0, 1\}$. In fact, Lecerf directly characterizes the relationship between the irreducible factors of $H(x, y)$ and the vectors spanning the solution space, though this detail is irrelevant for the moment.

Thus, in Lecerf’s argument, a_1, \dots, a_n are chosen such that a particular minor is nonzero, namely, a certain linear system has no non-trivial solutions. But it is enough to select a_1, \dots, a_n in a way that only guarantees that this linear system has no non-trivial $0/1$ solutions!

Fixing any vector $u \in \{0, 1\}^d$, the requirement that u is not a solution to the linear system turns out to be a condition expressible as u being a nonzero of a polynomial of degree $O(d)$, rather than $O(d^2)$. If we pick a_1, \dots, a_n from a hitting set generator with error δ smaller than 2^{-d} , we can afford to take a union bound over all vectors in $\{0, 1\}^d$ and ensure that none of them is a solution to the linear system while keeping the total error small. This is not a big price to pay in terms of the seed length of the HSG, which is $O(d \log n + \log(1/\delta))$, so requiring δ to be exponentially small in d adds an insignificant additive $O(d)$ term. Where we do pay the price for the small error is in the field size, since the explicit construction of the HSG we use requires the field size to be at least roughly d/δ . Fortunately, however, the dependence of the seed length of our generator on the field size q is also by an additive $O(\log q)$ term, which means that once more requiring q to be exponentially large in d has no adverse effects even on the total seed length of the PRG.

We briefly remark that, for technical reasons, Lecerf’s result also requires the characteristic of the underlying field to be zero or at least $d(d-1)+1$. We further elaborate on Lecerf’s techniques in Section 4.

We finally mention another important technical point. Lecerf’s irreducibility characterization [21] assumes a technical condition on the polynomial, which he called Hypothesis (H) (see Section 3). Such a “preprocessing” step, which makes the polynomial monic in a certain distinguished variable, is common to many factorization algorithms, and can usually be easily guaranteed by applying a random linear transformation to the variables. However, doing this in the naïve way would require the use of too many random bits. To solve this problem, in Section 3 we show that this part can also be derandomized by using a hitting set generator for polynomials of degree $O(d)$. As explained in Section 3, this part also uses crucially indecomposability (rather than irreducibility) in a novel way.

2 Preliminaries

We now define the basic objects studied in this paper and introduce the fundamental mathematical concepts used.

Notation

All logarithms are base 2. Denote by \mathbb{N} the set of natural numbers $\{0, 1, 2, \dots\}$. For $n \in \mathbb{N}$, define $[n] = \{1, 2, \dots, n\}$. For a finite set A , denote by \mathbf{U}_A the uniform distribution over A .

We often use symbols in bold, e.g., \mathbf{a} or \mathbf{x} , as the shorthand for a vector (a_1, \dots, a_n) or a sequence of variables x_1, \dots, x_n .

Denote by \mathbb{F}_q the finite field of size q . The algebraic closure of a field \mathbb{F} is denoted by $\overline{\mathbb{F}}$. For a commutative ring A and variables x_1, \dots, x_n , we denote by $A[[x_1, \dots, x_n]]$ or $A[[\mathbf{x}]]$ the *ring of formal power series* over A in x_1, \dots, x_n , i.e.,

$$A[[\mathbf{x}]] = \left\{ \sum_{\mathbf{e}=(e_1, \dots, e_n) \in \mathbb{N}^n} a_{\mathbf{e}} x_1^{e_1} \cdots x_n^{e_n} : a_{\mathbf{e}} \in A \right\}.$$

Pseudorandom Generators and Hitting Set Generators

► **Definition 2.1** (Pseudorandom generator, PRG). Let \mathbb{F}_q be a finite field. A *pseudorandom generator* (PRG) for n -variate polynomials of degree at most d over \mathbb{F}_q with error ε is an efficiently computable map $G : S \rightarrow \mathbb{F}_q^n$ from a finite set $S \neq \emptyset$ such that for every such polynomial f of degree at most d , the two distributions $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^n})$ are ε -close in statistical distance. That is,

$$\frac{1}{2} \sum_{a \in \mathbb{F}_q} \left| \Pr_{\mathbf{x} \in \mathbb{F}_q^n} [f(\mathbf{x}) = a] - \Pr_{\mathbf{y} \in S} [f(G(\mathbf{y})) = a] \right| \leq \varepsilon.$$

The quantity $\log |S|$ is called the *seed length* of G .

A weaker object than a PRG is a *hitting set generator*. Here, we only require that a nonzero polynomial is nonzero (with high probability) on the output of the generator.

► **Definition 2.2** (Hitting set generator, HSG). Let \mathbb{F} be a field. A *hitting set generator* (HSG) with density $1 - \delta$ for n -variate polynomials of degree at most d over \mathbb{F} is an efficiently computable map $H : S \rightarrow \mathbb{F}^n$ from a finite set $S \neq \emptyset$ such that for every such nonzero polynomial f of degree at most d ,

$$\Pr_{\mathbf{y} \in S} [f(H(\mathbf{y})) = 0] \leq \delta.$$

The quantity $\log |S|$ is called the *seed length* of G .

Building on the earlier work [19, 23] and algebraic-geometric codes, Guruswami and Xing [16] constructed explicit HSGs for low-degree polynomials with asymptotically optimal seed length and density.

► **Theorem 2.3** ([16]). There exists an absolute constant C such that for any n, d, q, δ , such that $q \geq Cd/\delta$, there exists an explicit HSG for n -variate polynomials of degree at most d over \mathbb{F}_q with density $1 - \delta$ and seed length $O(d \log n + \log(1/\delta))$.

It should also be noted that for hitting set generators, the field \mathbb{F} does not have to be finite. This generality is used in the statement of the following fact, that an HSG for a field \mathbb{F} is also a HSG for any extension field \mathbb{K} of \mathbb{F} .

► **Fact 2.4** ([6, 12]). Let $H : S \rightarrow \mathbb{F}$ be an HSG with density $1 - \delta$ for polynomials of degree at most d over a field \mathbb{F} , and let \mathbb{K} be an extension of \mathbb{F} . Then H is also an HSG with density $1 - \delta$ for polynomials of degree at most d over \mathbb{K} .

Proof. Let \mathcal{B} be a basis of \mathbb{K} over \mathbb{F} , and let f be a nonzero polynomial in $\mathbb{K}[x_1, \dots, x_n]$ of degree at most d . By expressing every coefficient $c \in \mathbb{K}$ of a monomial in f as a linear combination $c = \sum_{b \in \mathcal{B}} a_b \cdot b$ with $a_b \in \mathbb{F}$ for every $b \in \mathcal{B}$, we may write $f = \sum_{b \in \mathcal{B}} f_b \cdot b$ such that $f_b \in \mathbb{F}[x_1, \dots, x_n]$ is a polynomial of degree at most d for every $b \in \mathcal{B}$, and at least one f_b is nonzero. Thus, for any $u \in S$, $f(H(u)) = \sum_{b \in \mathcal{B}} f_b(u) \cdot b$ is nonzero unless $f_b(H(u)) = 0$ for every b , which happens with probability at most δ over the choice of $u \in S$. ◀

Indecomposable Polynomials

The *indecomposability* of a polynomial is crucially used in the analysis of the PRG construction in [12] as well as in our analysis. We first define this property.

Thus, in the case where $g = \frac{\partial f}{\partial y}$, it holds that $\text{Res}(f, g) \neq 0$ if and only if f does not have a root of multiplicity greater than one.

Hensel Lifting

Hensel lifting is a general technique for “lifting” roots or factorizations of a polynomial modulo an ideal I of a ring R to those modulo powers of I , under some mild conditions. The use of Hensel’s lifting lemma is standard in multivariate factorization algorithms, and it is available in various forms. We state one standard form, which can be derived from [13, Theorem 7.3] as a special case. This form is particularly relevant to our discussion of Lecerf’s techniques in Section 4.

► **Lemma 2.9** (Hensel’s lifting lemma). Let $f \in \mathbb{F}[x_1, \dots, x_n, y] = \mathbb{F}[\mathbf{x}, y]$ be a nonzero polynomial over a field \mathbb{F} . Suppose $\bar{\lambda} \in \mathbb{F}$ is a simple root of $f(\mathbf{0}, y) \in \mathbb{F}[y]$. Then there exists unique $\lambda \in \mathbb{F}[[\mathbf{x}]]$ such that

1. $f(\mathbf{x}, \lambda) = 0$, i.e., λ is a root of f as a univariate polynomial in y over $\mathbb{F}[[\mathbf{x}]]$, and
2. $\lambda(\mathbf{0}) = \bar{\lambda}$.

3 Hypothesis (H)

Lecerf’s papers [20, 21] on multivariate polynomial factoring assume a hypothesis about the polynomial f , which he calls Hypothesis (H). Such a hypothesis can be satisfied with high probability by applying a random linear transformation on the variables.

In this section, we discuss Lecerf’s Hypothesis (H) and show that, for our purpose, the random linear transformation can be derandomized by using a HSG for polynomials of degree $O(d)$. The fact that we are interested in the irreducibility of $f - t$ for an indeterminate t , rather than that of f , is crucial in keeping the degree linear in d .

Let \mathbb{F} be a field. First, we define Hypothesis (H).

► **Definition 3.1** (Hypothesis (H) [20, 21]). Let $f \in \mathbb{F}[x_1, \dots, x_n, y] = \mathbb{F}[\mathbf{x}, y]$ be a non-constant polynomial. We say f satisfies *Hypothesis (H)* if

1. f is monic in y and $\deg_y(f) = \deg(f)$,
2. $\text{Res}\left(f(\mathbf{0}, y), \frac{\partial f}{\partial y}(\mathbf{0}, y)\right) \neq 0$.

We also need a family of invertible linear transformations defined as follows.

► **Definition 3.2.** For $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{F}^n$, let $s_{\mathbf{a}}$ be the \mathbb{F} -linear automorphism of $\mathbb{F}[\mathbf{x}, y]$ that fixes y and sends x_i to $x_i + a_i y$.

► **Lemma 3.3.** Let $f \in \mathbb{F}[\mathbf{x}, y]$ be a nonzero polynomial of degree at most d . Then there exists a nonzero polynomial $B \in \mathbb{F}[\mathbf{x}]$ of degree at most d such that for every $\mathbf{a} \in \mathbb{F}^n$ satisfying $B(\mathbf{a}) \neq 0$, it holds that $\deg_y(s_{\mathbf{a}}(f)) = d$ and the coefficient of y^d in $s_{\mathbf{a}}(f)$ is in \mathbb{F}^\times .

Proof. Let f_d be the degree- d homogeneous part of f , so that we can write $f = f_d + g$ where $g = f - f_d$ has degree less than d . Write $f_d = \sum_{i=0}^d c_i(\mathbf{x})y^i$, where each $c_i \in \mathbb{F}[\mathbf{x}]$ is either zero or a homogeneous polynomial of degree $d - i$.

Consider $\mathbf{a} \in \mathbb{F}^n$. Note that

$$s_{\mathbf{a}}(f) = s_{\mathbf{a}}(f_d) + s_{\mathbf{a}}(g) = \sum_{i=0}^d s_{\mathbf{a}}(c_i(\mathbf{x}))y^i + s_{\mathbf{a}}(g) = \sum_{i=0}^d c_i(\mathbf{x} + y \cdot \mathbf{a})y^i + g(\mathbf{x} + y \cdot \mathbf{a}, y).$$

As $\deg(g) < d$ and each c_i is either zero or homogeneous of degree $d - i$, we have that $\deg_y s_{\mathbf{a}}(f) \leq d$, and that the coefficient of y^d in $s_{\mathbf{a}}(f)$ is $\sum_{i=0}^d c_i(\mathbf{a}) \in \mathbb{F}$. So we may choose $B = \sum_{i=0}^d c_i$, which is a nonzero polynomial of degree at most d . \blacktriangleleft

► **Lemma 3.4.** Assume $f \in \mathbb{F}[\mathbf{x}, y]$ is a polynomial of degree $d \geq 1$ that satisfies Item 1 of Hypothesis (H). Further assume that $\text{char}(\mathbb{F})$ is either zero or greater than d . Let $c \in \mathbb{F}^\times$. Then $f + ct$ is a degree- d polynomial satisfying Hypothesis (H) as a polynomial over $\mathbb{F}(t)$.

Proof. As f satisfies Item 1 of Hypothesis (H) and has degree $d \geq 1$, so does $f + ct$. So it suffices to verify Item 2. Write $f = \sum_{i=0}^d c_i y^i$, where $c_i \in \mathbb{F}[\mathbf{x}]$ and $c_d = 1$. Then $\frac{\partial(f+ct)}{\partial y} = \sum_{i=1}^d (i \cdot c_i) y^{i-1}$, which has degree $d - 1$ in y since $d \cdot c_d = d \neq 0$ by the assumption about $\text{char}(\mathbb{F})$.

Let $\bar{c}_i = c_i(\mathbf{0})$ for $i = 0, 1, \dots, d$. Let $h = \text{Res}\left((f + ct)(\mathbf{0}, y), \frac{\partial(f+ct)}{\partial y}(\mathbf{0}, y)\right)$. Then h is the determinant of the following $(2d - 1) \times (2d - 1)$ matrix:

$$\begin{pmatrix} \bar{c}_0 + ct & 0 & \cdots & 0 & \bar{c}_1 & 0 & \cdots & 0 \\ \bar{c}_1 & \bar{c}_0 + ct & \cdots & 0 & 2\bar{c}_2 & \bar{c}_1 & \cdots & 0 \\ \bar{c}_2 & \bar{c}_1 & \ddots & 0 & 3\bar{c}_3 & 2\bar{c}_2 & \ddots & 0 \\ \vdots & \vdots & \ddots & \bar{c}_0 + ct & \vdots & \vdots & \ddots & \bar{c}_1 \\ \bar{c}_d & \bar{c}_{d-1} & \cdots & \vdots & d\bar{c}_d & (d-1)\bar{c}_{d-1} & \cdots & \vdots \\ 0 & \bar{c}_d & \ddots & \vdots & 0 & d\bar{c}_d & \ddots & \vdots \\ \vdots & \vdots & \ddots & \bar{c}_{d-1} & \vdots & \vdots & \ddots & (d-1)\bar{c}_{d-1} \\ 0 & 0 & \cdots & \bar{c}_d & 0 & 0 & \cdots & d\bar{c}_d \end{pmatrix}.$$

Observe that $\deg_t h \leq d - 1$, and that the coefficient of t^{d-1} in h is $c^{d-1}(d\bar{c}_d)^d = c^{d-1}d^d \neq 0$ since only those entries on the diagonal contribute to this coefficient. This implies that $h \neq 0$, i.e., $f + ct$ satisfies Item 2 of Hypothesis (H). \blacktriangleleft

► **Corollary 3.5.** Assume that $f \in \mathbb{F}[\mathbf{x}, y]$ is a polynomial of degree $d \geq 1$ and that $\text{char}(\mathbb{F})$ is either zero or greater than d . Then there exists a nonzero polynomial $B \in \mathbb{F}[\mathbf{x}]$ of degree at most d such that for every $\mathbf{a} \in \mathbb{F}^n$ satisfying $B(\mathbf{a}) \neq 0$, $s_{\mathbf{a}}(f) - t$ equals a product $c \cdot g$ where $c \in \mathbb{F}^\times$ and $g \in \mathbb{F}(t)[\mathbf{x}, y]$ is a degree- d polynomial satisfying Hypothesis (H).

Proof. Let B be as in Lemma 3.3. Consider $\mathbf{a} \in \mathbb{F}^n$ satisfying $B(\mathbf{a}) \neq 0$. By Lemma 3.3, we may write $s_{\mathbf{a}}(f) = c \cdot \tilde{g}$ where $c \in \mathbb{F}^\times$ and \tilde{g} satisfies Item 1 of Hypothesis (H). Then $s_{\mathbf{a}}(f) - t = c \cdot \tilde{g} - t = c \cdot g$ where $g = \tilde{g} - c^{-1}t$. By Lemma 3.4, g is a degree- d polynomial satisfying Hypothesis (H). \blacktriangleleft

Thus, by choosing good $\mathbf{a} \in \mathbb{F}$ via an explicit HSG for polynomials of degree at most d and performing the transformation $f \mapsto s_{\mathbf{a}}(f)$, we may assume $f - t$ satisfies Hypothesis (H).

Satisfying Hypothesis (H) in Small Characteristics

While our final result needs $\text{char}(\mathbb{F}) > d(d - 1)$, the assumption that $\text{char}(\mathbb{F})$ is zero or large enough is not crucial for the sake of satisfying Hypothesis (H). We now sketch how to modify the proof of Lemma 3.4 when $0 < \text{char}(\mathbb{F}) \leq d$.

Let $p = \text{char}(\mathbb{F}) > 0$. For our purpose, we may assume \mathbb{F} is a perfect field and f is indecomposable over \mathbb{F} . This implies that $f \notin \mathbb{F}[x_1^p, \dots, x_n^p, y^p]$. Then it is not hard to show that there exists an integer $e > 0$ coprime to p such that for random $\mathbf{a} \in \mathbb{F}^n$, with high

44:10 Optimal PRGs for Low-Degree Polynomials over Large Fields

probability, not only is the coefficient of y^d in $s_{\mathbf{a}}(f)$ nonzero, but so is the coefficient of y^e . Choose the largest e that has this property. After replacing f by $s_{\mathbf{a}}(f)$, the polynomial $\frac{\partial f + ct}{\partial y}$ in the proof of Lemma 3.4 would have degree $e - 1$ instead of $d - 1$ in y . Then $\deg_t(h) = e - 1$ and the coefficient of t^{e-1} in h is $c^{e-1}(e\bar{c}_e)^d$, which is nonzero iff $\bar{c}_e = c_e(\mathbf{0})$ is nonzero. The latter condition can be guaranteed with high probability by performing the substitutions $x_i \mapsto x_i + b_i$ for random $\mathbf{b} = (b_1, \dots, b_n) \in \mathbb{F}^n$. Finally, it is not difficult to show that the choices of \mathbf{a} and \mathbf{b} can be derandomized by using an explicit HSG for polynomials of degree $O(d)$.

4 Lecerf's Techniques

We describe Lecerf's techniques in this section. For simplicity, our discussion is restricted to the special case where the base field is algebraically closed.

Let \mathbb{K} be an algebraically closed field, and let $f \in \mathbb{K}[\mathbf{x}, y]$ be a polynomial of degree $d \geq 1$ satisfying Hypothesis (H). Define $\bar{f} := f(\mathbf{0}, y) \in \mathbb{K}[y]$. As \mathbb{K} is algebraically closed and $\text{Res}\left(f(\mathbf{0}, y), \frac{\partial f}{\partial y}(\mathbf{0}, y)\right) \neq 0$, the univariate polynomial \bar{f} factorizes into distinct linear factors

$$\bar{f} = \prod_{i=1}^d (y - \bar{\lambda}_i)$$

where $\bar{\lambda}_i \in \mathbb{K}$ for $i \in [d]$. By Hensel's lifting lemma, the above factorization of \bar{f} over \mathbb{K} lifts to a factorization of f into distinct linear factors

$$f = \prod_{i=1}^d (y - \lambda_i),$$

where $\lambda_i \in \mathbb{K}[[\mathbf{x}]]$ and $\lambda_i(\mathbf{0}) = \bar{\lambda}_i$ for $i \in [d]$.

Now we introduce new variables $\mathbf{z} = (z_1, \dots, z_n)$ and x , and define $g := f(z_1x, \dots, z_nx, y) \in \mathbb{K}[\mathbf{z}, x, y]$. Then g factorizes into linear factors

$$g = \prod_{i=1}^d (y - \lambda_i(z_1x, \dots, z_nx))$$

where each $\lambda_i(z_1x, \dots, z_nx)$ lives in $\mathbb{K}[\mathbf{z}][[x]]$. For $i \in [d]$, let g_i be the factor $y - \lambda_i(z_1x, \dots, z_nx)$ of g , and let \hat{g}_i be its cofactor $\prod_{j \in [d] \setminus \{i\}} g_j$. So $g_i, \hat{g}_i \in \mathbb{K}[\mathbf{z}][[x]][y]$.

For $h \in A[[x]][y]$ over a commutative ring A and $(j, k) \in \mathbb{N}^2$, denote by $\text{coeff}(h, x^j y^k) \in A$ the coefficient of $x^j y^k$ in h . We are now ready to define the linear system $D_{\mathbf{z}, \sigma}$ used in [20, 21].

► **Definition 4.1** (Linear system $D_{\mathbf{z}, \sigma}$ [20, 21]). Let $\sigma \in \mathbb{N}$. Define $D_{\mathbf{z}, \sigma}$ to be the following linear system over $\mathbb{K}(\mathbf{z})$ in the unknowns ℓ_1, \dots, ℓ_d :

$$D_{\mathbf{z}, \sigma} \begin{cases} \sum_{i=1}^d \text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k\right) \cdot \ell_i = 0, & k \leq d - 1, d \leq j + k \leq \sigma - 1, \\ \sum_{i=1}^d \text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k\right) \cdot \ell_i = 0, & k \leq d - 1, j \leq \sigma - 2, d \leq j + k \leq \sigma - 1. \end{cases}$$

We have the following easy lemma.

► **Lemma 4.2.** For $(j, k) \in \mathbb{N}^2$, $\text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k\right)$, $\text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k\right) \in \mathbb{K}[\mathbf{z}]$ are polynomials of degree at most $j + 1$ and j respectively.

Proof. Consider arbitrary $i \in [d]$ and $(j, k) \in \mathbb{N}^2$. As $g_i = y - \lambda_i(z_1 x, \dots, z_n x)$ and $j \geq 0$, only terms of degree at most j in z_1, \dots, z_n contribute to the coefficient of $x^j y^k$ in g_i . Then, as the operator $\frac{\partial}{\partial x}$ is linear and sends $x^u y^v$ to $u x^{u-1} y^v$ for all $u, v \in \mathbb{N}$, one can see that only terms of degree at most $j + 1$ in z_1, \dots, z_n contribute to the coefficient of $x^j y^k$ in $\frac{\partial g_i}{\partial x}$. Also, $\frac{\partial g_i}{\partial y} = 1$ by definition.

For any $h_1, \dots, h_s \in \mathbb{K}[\mathbf{z}][[x]][y]$ and $h = \prod_{i=1}^s h_i$, we have

$$\text{coeff}(h, x^j y^k) = \sum_{\substack{j_1, \dots, j_s, k_1, \dots, k_s \in \mathbb{N} \\ \sum_i j_i = j, \sum_i k_i = k}} \prod_{i=1}^s \text{coeff}(h_i, x^{j_i} y^{k_i}). \quad (1)$$

We already know $\deg(\text{coeff}(g_i, x^j y^k)) \leq j$, $\deg(\text{coeff}(\frac{\partial g_i}{\partial x}, x^j y^k)) \leq j + 1$, and $\deg(\text{coeff}(\frac{\partial g_i}{\partial y}, x^j y^k)) = 0$ for $i \in [d]$ and $(j, k) \in \mathbb{N}^2$ by the above discussion. Choosing (h_1, \dots, h_s) to be $(g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_d, \frac{\partial g_i}{\partial x})$ and $(g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_d, \frac{\partial g_i}{\partial y})$ respectively and applying (1) proves the claim. ◀

For $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{K}^n$, we can assign a_1, \dots, a_n to z_1, \dots, z_n respectively in the polynomials $\text{coeff}(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k)$, $\text{coeff}(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k) \in \mathbb{K}[\mathbf{z}]$. This yields a linear system over \mathbb{K} , called the *specialization* of $D_{\mathbf{z}, \sigma}$ at \mathbf{a} and denoted by $D_{\mathbf{a}, \sigma}$.

► **Definition 4.3 (Specialization).** For $\sigma \in \mathbb{N}$ and $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{K}^n$, define $D_{\mathbf{a}, \sigma}$ to be the following linear system over \mathbb{K} in the unknowns ℓ_1, \dots, ℓ_d :

$$D_{\mathbf{a}, \sigma} \begin{cases} \sum_{i=1}^d \text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k\right)(\mathbf{a}) \cdot \ell_i = 0, & k \leq d - 1, d \leq j + k \leq \sigma - 1, \\ \sum_{i=1}^d \text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k\right)(\mathbf{a}) \cdot \ell_i = 0, & k \leq d - 1, j \leq \sigma - 2, d \leq j + k \leq \sigma - 1. \end{cases}$$

For $S \subseteq [d]$, define $\delta_S = (\delta_{S,1}, \dots, \delta_{S,d}) \in \mathbb{K}^d$ by

$$\delta_{S,i} = \begin{cases} 1 & i \in S, \\ 0 & i \notin S. \end{cases}$$

For every factor \tilde{f} of f , we may associate a set $S \subseteq [d]$ such that $\tilde{f} = \prod_{i \in S} (y - \lambda_i)$, i.e., S is the set of indices $i \in [d]$ such that $y - \lambda_i$ divides \tilde{f} . The irreducible factors f_1, \dots, f_r of f over \mathbb{K} are then associated with sets $S_1, \dots, S_r \subseteq [d]$, which form a partition of $[d]$. In [20, 21], Lecerf proved that, when σ is large enough, the solution space of $D_{\mathbf{z}, \sigma}$ is exactly spanned by the vectors $\delta_{S_1}, \dots, \delta_{S_r}$, and a similar statement holds for the specializations $D_{\mathbf{a}, \sigma}$. We state Lecerf's results formally as the following theorem.

► **Theorem 4.4** ([20, 21]). Assume $\text{char}(\mathbb{K})$ is zero or greater than $d(d - 1)$. Let $\sigma \geq 2d$. Let $f \in \mathbb{K}[\mathbf{x}, y]$ be a polynomial of degree $d \geq 1$ satisfying Hypothesis (H). Then:

1. Suppose $f = \prod_{i=1}^r f_i$ is the factorization of f into its irreducible factors over \mathbb{K} . For $i \in [r]$, let S_i be the set of indices $j \in [d]$ such that $y - \lambda_j$ divides f_i . Then $\delta_{S_1}, \dots, \delta_{S_r}$ form a basis of the solution space of $D_{\mathbf{z}, \sigma}$.

44:12 Optimal PRGs for Low-Degree Polynomials over Large Fields

2. Let $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{K}^n$ and $f_{\mathbf{a}} = f(a_1x, \dots, a_nx, y) \in \mathbb{K}[x, y]$. Suppose $f_{\mathbf{a}} = \prod_{i=1}^s f_{\mathbf{a},i}$ is the factorization of $f_{\mathbf{a}}$ into its irreducible factors over \mathbb{K} . For $i \in [s]$, let $S_{\mathbf{a},i}$ be the set of indices $j \in [d]$ such that $y - \lambda_j(a_1x, \dots, a_nx)$ divides $f_{\mathbf{a},i}$. Then $\delta_{S_{\mathbf{a},1}}, \dots, \delta_{S_{\mathbf{a},s}}$ form a basis of the solution space of $D_{\mathbf{a},\sigma}$.

The first item of Theorem 4.4 is explicitly stated as [21, Lemma 1]. The second item follows from [20, Theorem 1 and Lemma 4]. For a detailed analysis of the linear systems $D_{\mathbf{z},\sigma}$ and $D_{\mathbf{a},\sigma}$, and for a conceptual interpretation of these linear systems in terms of the closedness condition of differential 1-forms, we refer the reader to [20, 21], as well as to the earlier paper of Gao [15].

Bertinian Good/Bad Points

The classical Bertini irreducibility theorem [30] states, among other things, that over an algebraically closed field \mathbb{K} , the intersection of an irreducible affine variety of codimension one (i.e., a hypersurface) with a plane in general position is still irreducible. This motivates the following definition:

► **Definition 4.5** (Bertinian good/bad points [21]). Let $f \in \mathbb{K}[\mathbf{x}, y]$ be a non-constant polynomial satisfying Hypothesis (H). We say $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{K}^n$ is a *Bertinian good point* for f if for every irreducible factor \tilde{f} of f over \mathbb{K} , the bivariate polynomial $\tilde{f}_{\mathbf{a}} = \tilde{f}(a_1x, \dots, a_nx, y)$ is also irreducible over \mathbb{K} . We say $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{K}^n$ is a *Bertinian bad point* for f if it is not a Bertinian good point for f .

Lecerf [21, Theorem 6] proved that given f , there exists a nonzero polynomial $Q \in \mathbb{K}[z_1, \dots, z_n]$ of degree at most $(d-1)(2d-1)$ that vanishes at all Bertinian bad points for f , where $d = \deg(f)$. Let M be the matrix representing the linear system $D_{\mathbf{z},\sigma}$. Lecerf's proof can be sketched as follows: By Theorem 4.4, the solution space of $D_{\mathbf{a},\sigma}$ contains that of $D_{\mathbf{z},\sigma}$, and \mathbf{a} is Bertinian good as long as the two are equal. Thus, we may choose Q to be the determinant of the largest nonsingular submatrix of M . This is because for such Q , if Q does not vanish at \mathbf{a} , then $D_{\mathbf{z},\sigma}$ and $D_{\mathbf{a},\sigma}$ have the same rank, and hence their solution spaces must be equal. The bound $(d-1)(2d-1)$ on the degree of Q follows from Lemma 4.2.

In [21], Lecerf also demonstrated that the degree bound $(d-1)(2d-1)$ is asymptotically tight by providing an example for which a degree of $\Omega(d^2)$ of the polynomial Q is necessary. However, our next lemma states that, perhaps surprisingly, the degree bound can be improved to $2d-1$ if we allow the use of the zero loci of *multiple* polynomials to cover the Bertinian bad points for f . For simplicity, we state the lemma in the special case where f is irreducible, which suffices for our purpose.

► **Lemma 4.6.** Assume $\text{char}(\mathbb{K})$ is zero or greater than $d(d-1)$. Let $f \in \mathbb{K}[\mathbf{x}, y]$ be an irreducible polynomial over \mathbb{K} of degree $d \geq 1$ satisfying Hypothesis (H). Let $m = 2^{d-1} - 1$. Then there exist nonzero polynomials $Q_1, \dots, Q_m \in \mathbb{K}[\mathbf{z}] = \mathbb{K}[z_1, \dots, z_n]$ of degree at most $2d-1$ such that for every Bertinian bad point $\mathbf{a} \in \mathbb{K}^n$ for f , at least one polynomial Q_i vanishes at \mathbf{a} .

Proof. Let $\sigma = 2d$. Let N be the number of equations in $D_{\mathbf{z},\sigma}$. Let M be the $N \times d$ matrix over $\mathbb{F}(\mathbf{z})$ representing the linear system $D_{\mathbf{z},\sigma}$. Note that by Definition 4.1, the entries of M are of the form $\text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k\right)$ with $j \leq \sigma - 2$ or $\text{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k\right)$ with $j \leq \sigma - 1$. By Lemma 4.2 and the fact that $\sigma = 2d$, the entries of M are polynomials in $\mathbb{K}[\mathbf{z}]$ of degree at most $2d-1$.

There are exactly $m = 2^{d-1} - 1$ proper subsets of $[d]$ containing 1. Let S_1, \dots, S_m be an enumeration of them. Consider $i \in [m]$. As f is irreducible, by Theorem 4.4, δ_{S_i} is not in the solution space of $D_{\mathbf{z}, \sigma}$. So we can fix a row $\mathbf{r}_i = (r_{i,1}, \dots, r_{i,d})$ of M such that the inner product of \mathbf{r}_i and δ_{S_i} is nonzero, i.e., $\sum_{j=1}^d r_{i,j} \delta_{S_i,j} \neq 0$. Let $Q_i = \sum_{j=1}^d r_{i,j} \delta_{S_i,j}$, which is a nonzero polynomial in $\mathbb{K}[\mathbf{z}]$ of degree at most $2d - 1$. Choose Q_i in this way for each $i = 1, \dots, m$.

Now let \mathbf{a} be a Bertinian bad point for f . Then $f_{\mathbf{a}} = f(a_1x, \dots, a_nx, y)$ factorizes into more than one irreducible factor over \mathbb{K} . Let $\tilde{f}_{\mathbf{a}}$ be the irreducible factor of $f_{\mathbf{a}}$ divisible by $y - \lambda_1(a_1x, \dots, a_nx)$. Let S be the set of $j \in [d]$ such that $\tilde{f}_{\mathbf{a}}$ is divisible by $y - \lambda_j(a_1x, \dots, a_nx)$. Then S is a proper subset of $[d]$ containing 1. So $S = S_i$ for some $i \in [m]$. By Theorem 4.4, δ_{S_i} is in the solution space of $D_{\mathbf{a}, \sigma}$. As $D_{\mathbf{a}, \sigma}$ is the specialization of $D_{\mathbf{z}, \sigma}$ at \mathbf{a} , the vector $(r_{i,1}(\mathbf{a}), \dots, r_{i,d}(\mathbf{a}))$ is a row of the matrix representing $D_{\mathbf{a}, \sigma}$. So $\sum_{j=1}^d r_{i,j}(\mathbf{a}) \delta_{S_i,j} = 0$, i.e., $Q_i(\mathbf{a}) = 0$. ◀

The Number of Low-Degree Polynomials Needed

It is an intriguing mathematical question to us how many low-degree polynomials are needed to cover the Bertinian bad points for f . We now formalize this question.

▶ **Definition 4.7.** Let \mathbb{K} be an algebraically closed field. For positive integers d and D , define $N(d, D, \mathbb{K})$ to be the smallest $N \in \mathbb{N}$ such that the following holds: Let $f \in \mathbb{K}[\mathbf{x}, y]$ be an irreducible polynomial of degree at most d over \mathbb{K} satisfying Hypothesis (H). Then there exist N nonzero polynomials in $\mathbb{K}[\mathbf{z}]$ of degree at most D such that the union of the zero loci of these polynomials contains all Bertinian bad points for f in \mathbb{K}^n .

If such N does not exist, define $N(d, D, \mathbb{K}) = \infty$.

In our application, it suffices to consider polynomials of the special form $f + c \cdot t$, where $c \in \mathbb{F}^\times$, $f \in \mathbb{F}[\mathbf{x}, y]$ and $\mathbb{K} = \overline{\mathbb{F}(t)}$. Moreover, by performing a variable substitution $t \mapsto -c^{-1}t$, we may assume $c = -1$. This motivates us to introduce the following variant of Definition 4.7:

▶ **Definition 4.8.** Let \mathbb{F} be a field. For positive integers d and D , define $N^*(d, D, \mathbb{F})$ to be the smallest $N \in \mathbb{N}$ such that the following holds: Let $f \in \mathbb{F}[\mathbf{x}, y]$ be a polynomial of degree at most d such that $f - t$ is an irreducible polynomial over $\overline{\mathbb{F}(t)}$ satisfying Hypothesis (H). Then there exist N nonzero polynomials in $\overline{\mathbb{F}(t)}[\mathbf{z}]$ of degree at most D such that the union of the zero loci of these polynomials contains all Bertinian bad points for $f - t$ in $\overline{\mathbb{F}^n}$.

If such N does not exist, define $N^*(d, D, \mathbb{F}) = \infty$.

Lecerf's result [21, Theorem 6] can be interpreted as the statement that when $\text{char}(\mathbb{K})$ is zero or greater than $d(d - 1)$, it holds that

$$N(d, D, \mathbb{K}) = 1 \text{ for } D \geq (d - 1)(2d - 1).$$

Our Lemma 4.6 states that under the same condition, we have

$$N(d, D, \mathbb{K}) \leq 2^{d-1} - 1 \text{ for } D \geq 2d - 1.$$

In [21], Lecerf gave an example showing that the degree bound $O(d^2)$ for the smallest D satisfying $N(d, D, \mathbb{K}) = 1$ is asymptotically tight.² As one can always combine the $N(d, D, \mathbb{K})$ polynomials of degree at most D into a single polynomial of degree at most $N(d, D, \mathbb{K}) \cdot D$ by taking their product, this implies $N(d, D, \mathbb{K}) \cdot D = \Omega(d^2)$, i.e., $N(d, D, \mathbb{K}) = \Omega(d^2/D)$.

² See the example before Theorem 6 in [21].

► **Question 4.9.** Give improved upper bounds (or lower bounds) on $N(d, D, \mathbb{K})$ and $N^*(d, D, \mathbb{F})$, at least when the characteristic of \mathbb{K} or \mathbb{F} is zero or large enough.

By definition, $N^*(d, D, \mathbb{F}) \leq N(d, D, \overline{\mathbb{F}(t)})$. A subexponential upper bound on $N^*(d, D, \mathbb{F})$ for $D = O(d)$ will improve the required field size in Theorem 1.1.

Finally, it might be possible to exploit some extra structure to derive better bounds on $N^*(d, D, \mathbb{F})$ than those obtained for $N(d, D, \mathbb{K})$. For example, if we modify the definition of $N^*(d, D, \mathbb{F})$ by only considering those polynomials f of *prime* degree, then $N^*(d, 2d-1, \mathbb{F}) \leq 1$. This is because if $f_{\mathbf{a}} - t$ is reducible over $\overline{\mathbb{F}(t)}$, then $f_{\mathbf{a}}$ is decomposable over $\overline{\mathbb{F}}$ by Lemma 2.8. But as $\deg(f)$ is prime, $f_{\mathbf{a}}$ must be of the form $g(h)$ with $\deg(g) = \deg(f)$ and $\deg(h) = 1$. This in turn implies that $f_{\mathbf{a}} - t$ factorizes into $\deg(f)$ linear factors over $\overline{\mathbb{F}(t)}$. In Theorem 5.4, we use this idea to show that the required field size can be improved to $O(d^4/\varepsilon^2)$ if we only want to fool polynomials whose degrees are prime and at most d .

5 Proofs of the Main Theorems

In this section, we present our PRG construction and prove the main theorems.

Let n and d be positive integers. Let \mathbb{F}_q be a finite field of characteristic at least $d(d-1)+1$. We now present the construction of our PRG

$$G : S \rightarrow \mathbb{F}_q^{n+1}$$

for polynomials $f \in \mathbb{F}_q[\mathbf{x}, y] = \mathbb{F}_q[x_1, \dots, x_n, y]$ of degree at most d . To simplify our notation, these polynomials are assumed to be $(n+1)$ -variate rather than n -variate.

► **Construction 5.1.** The construction is as follows:

- Let $H : T \rightarrow \mathbb{F}_q^n$ be an explicit HSG for n -variate polynomials of degree at most $2d-1$ over \mathbb{F}_q with density $1-\delta$ and seed length $\log |T| = O(d \log n + \log(1/\delta))$, where $\delta = C_0(2d-1)/q$ and $C_0 > 0$ is an absolute constant. For $i \in [n]$ and $s \in T$, denote the i -th coordinate of $H(s)$ by $H(s)_i$. The existence of H is guaranteed by Theorem 2.3.
- Let $S = T \times T \times \mathbb{F}_q \times \mathbb{F}_q$. Define $G : S \rightarrow \mathbb{F}_q^{n+1}$ by

$$G(r, s, u, v) = (H(s)_1 \cdot u + H(r)_1 \cdot v, \dots, H(s)_n \cdot u + H(r)_n \cdot v, v).$$

In other words, we use random $(r, s) \in T \times T$ to pick a plane in \mathbb{F}_q^{n+1} , and use random $(u, v) \in \mathbb{F}_q \times \mathbb{F}_q$ to pick a point on the plane. The following lemma states that with high probability, a given indecomposable polynomial $f \in \mathbb{F}_q[\mathbf{x}, y]$ remains indecomposable when restricted to the plane.

► **Lemma 5.2.** Let $f \in \mathbb{F}_q[\mathbf{x}, y]$ be an indecomposable polynomial of degree at most d over \mathbb{F}_q . Let (r, s) be a random element of $T \times T$. Let $\mathbf{a} = (a_1, \dots, a_n) = H(r)$ and $\mathbf{b} = (b_1, \dots, b_n) = H(s)$. Finally, let $F = f(b_1x + a_1y, \dots, b_nx + a_ny, y) \in \mathbb{F}_q[x, y]$. Then

$$\Pr[F \text{ is indecomposable over } \mathbb{F}_q] \geq 1 - 2^{d-1}\delta.$$

Proof. Recall that $s_{\mathbf{a}}$ is the \mathbb{F}_q -linear automorphism of $\mathbb{F}_q[\mathbf{x}, y]$ that fixes y and sends x_i to $x_i + a_iy$. As f is indecomposable over \mathbb{F}_q , so is $s_{\mathbf{a}}(f)$. By Lemma 2.6, $s_{\mathbf{a}}(f)$ is also indecomposable over $\overline{\mathbb{F}_q}$. So $s_{\mathbf{a}}(f) - t$ is irreducible over $\overline{\mathbb{F}_q(t)}$ by Lemma 2.8.

By Corollary 3.5, there exists a nonzero polynomial $B \in \mathbb{F}_q[\mathbf{x}]$ of degree at most d such that if $B(\mathbf{a}) \neq 0$, then

$$s_{\mathbf{a}}(f) - t = c \cdot g \tag{2}$$

where $c \in \mathbb{F}_q^\times$ and $g \in \mathbb{F}_q(t)[\mathbf{x}, y] \subseteq \overline{\mathbb{F}_q(t)}[\mathbf{x}, y]$ is a degree- d polynomial satisfying Hypothesis (H). By the HSG property of H , the event $B(\mathbf{a}) \neq 0$ happens with probability at least $1 - \delta$. Condition on this event, so that (2) holds. As $s_{\mathbf{a}}(f) - t$ is irreducible over $\overline{\mathbb{F}_q(t)}$, so is g .

Let $m = 2^{d-1} - 1$. By Lemma 4.6, there exist nonzero polynomials $Q_1, \dots, Q_m \in \overline{\mathbb{F}_q(t)}[z_1, \dots, z_n]$ of degree at most $2d - 1$ such that the union of the zero loci of these polynomials contains all $\mathbf{b}^* = (b_1^*, \dots, b_n^*) \in \overline{\mathbb{F}_q}^n$ for which $g(b_1^*x, \dots, b_n^*x, y)$ is reducible over $\overline{\mathbb{F}_q(t)}$. By Fact 2.4, H is an HSG with density $1 - \delta$ for polynomials of degree at most $2d - 1$ over $\overline{\mathbb{F}_q(t)}$.³ Therefore, for each $i \in [m]$, the probability that $Q_i(\mathbf{b}) = 0$ is at most δ . Condition on the event $Q_1(\mathbf{b}), \dots, Q_m(\mathbf{b}) \neq 0$. Then $g(b_1x, \dots, b_nx, y)$ is irreducible over $\overline{\mathbb{F}_q(t)}$. On the other hand, note that

$$c \cdot g(b_1x, \dots, b_nx, y) \stackrel{(2)}{=} (s_{\mathbf{a}}(f))(b_1x, \dots, b_nx, y) - t = f(b_1x + a_1y, \dots, b_nx + a_ny, y) - t = F - t$$

where the second step uses the definition $s_{\mathbf{a}}(f) = f(x_1 + a_1y, \dots, x_n + a_ny, y) \in \mathbb{F}_q[\mathbf{x}, y]$. So $F - t$ is irreducible over $\overline{\mathbb{F}_q(t)}$. By Lemma 2.8, F is indecomposable over $\overline{\mathbb{F}_q}$. So it is indecomposable over \mathbb{F}_q .

The indecomposability of F over \mathbb{F}_q relies on the conditions $B(\mathbf{a}) \neq 0$ and $Q_1(\mathbf{b}), \dots, Q_m(\mathbf{b}) \neq 0$. By the union bound, these conditions are simultaneously satisfied with probability at least $1 - \delta - m\delta = 1 - 2^{d-1}\delta$, which completes the proof. ◀

Now we are ready to prove Theorem 1.1.

► **Theorem 5.3** (Theorem 1.1 restated). There exists an absolute constant $C > 0$ such that for $\varepsilon > 0$ and $q \geq C(d2^d/\varepsilon + d^4/\varepsilon^2)$ with $\text{char}(\mathbb{F}_q) \geq d(d-1) + 1$, G as in Construction 5.1 is a PRG for $(n+1)$ -variate polynomials of degree at most d over \mathbb{F}_q with error ε and seed length $O(d \log n + \log q)$.

Proof. Let $f \in \mathbb{F}_q[\mathbf{x}, y]$ be a polynomial of degree at most d . We want to prove that $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ are ε -close (in statistical distance). We may assume that f is a non-constant polynomial, i.e., $\deg(f) \geq 1$, since the claim is trivial otherwise.

Our next step is the same as in [12]: f can always be written in the form $f = g(h)$, where $g \in \mathbb{F}_q[z]$ is a univariate polynomial and $h \in \mathbb{F}_q[\mathbf{x}, y]$ is indecomposable over \mathbb{F}_q . Let $D = h(G(\mathbf{U}_S))$ and $D' = h(\mathbf{U}_{\mathbb{F}_q^{n+1}})$. Then $f(G(\mathbf{U}_S)) = g(D)$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}}) = g(D')$. If D and D' are ε -close, then $g(D)$ and $g(D')$ are also ε -close. Thus, by replacing f with h , we may assume that f is indecomposable over \mathbb{F}_q .

Let $r, s, \mathbf{a}, \mathbf{b}$ and F be as in Lemma 5.2. Then by Lemma 5.2, the probability that F is decomposable over \mathbb{F}_q over random r and s is at most $2^{d-1}\delta = C_0 2^{d-1}(2d-1)/q$, where C_0 is as in Construction 5.1. Fix r and s such that F is indecomposable over \mathbb{F}_q . Then $f(G(r, s, u, v)) = F(u, v)$ by definition. Applying Lemma 2.7 to F shows that, for such fixed r and s , the distribution of $F(u, v)$, i.e., $f(G(r, s, u, v))$, over random $u, v \in \mathbb{F}_q$ is ε' -close to $\mathbf{U}_{\mathbb{F}_q}$, where $\varepsilon' = C_1 d^2/\sqrt{q}$ and $C_1 > 0$ is an absolute constant. It follows that the statistical distance between $f(G(\mathbf{U}_S))$ and $\mathbf{U}_{\mathbb{F}_q}$ is at most $2^{d-1}\delta + \varepsilon'$.

On the other hand, as f is also indecomposable over \mathbb{F}_q , applying Lemma 2.7 to f shows that $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ is ε' -close to $\mathbf{U}_{\mathbb{F}_q}$. Therefore, the statistical distance between $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ is at most

$$(2^{d-1}\delta + \varepsilon') + \varepsilon' = 2^{d-1}\delta + 2\varepsilon' = C_0 2^{d-1}(2d-1)/q + 2C_1 d^2/\sqrt{q} \tag{3}$$

³ Note that we are applying Fact 2.4 to the infinite extension $\overline{\mathbb{F}_q(t)}/\mathbb{F}_q$. In principle, it should be possible to make the argument finitary by making some adaptations, such as considering specific values of t . However, this may increase the complexity of the proof.

44:16 Optimal PRGs for Low-Degree Polynomials over Large Fields

which is bounded by ε provided that $q \geq C(d2^d/\varepsilon + d^4/\varepsilon^2)$ and $C > 0$ is a large enough absolute constant. The seed length of G is

$$2 \log |T| + 2 \log q = O(d \log n + \log(1/\delta) + \log q) = O(d \log n + \log q)$$

as $\delta = C_0(2d - 1)/q$. ◀

We conclude this section by proving Theorem 1.2, which states that the required field size can be improved to $O(d^4/\varepsilon^2)$ if we only want to fool polynomials of prime degree.

► **Theorem 5.4** (Theorem 1.2 restated). There exists an absolute constant $C > 0$ such that for $\varepsilon > 0$ and $q \geq C(d^4/\varepsilon^2)$ with $\text{char}(\mathbb{F}_q) \geq d(d - 1) + 1$, G as in Construction 5.1 is a PRG for $(n + 1)$ -variate polynomials of *prime* degree up to d with error ε and seed length $O(d \log n + \log q)$.

Proof Sketch. Let $f \in \mathbb{F}_q[\mathbf{x}, y]$ be a polynomial whose degree d' is prime and at most d . We want to prove that $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ are ε -close (in statistical distance). Suppose f is decomposable over \mathbb{F}_q . Then $f = g(h)$ for some polynomials g, h over \mathbb{F}_q where $\deg(g) \geq 2$, and as $d' = \deg(f)$ is prime, we must have $\deg(g) = d'$ and $\deg(h) = 1$. In this case, the theorem follows by replacing f with h , which has degree one, and applying Theorem 5.3. So we may assume that f is indecomposable over \mathbb{F}_q .

The rest of the proof follows that of Theorem 5.3, except that we could bound the probability that F is decomposable over \mathbb{F}_q by 2δ , rather than by $2^{d-1}\delta$, using the following observation:

In the application of Lemma 4.6, the polynomial has the special form $g^* = f^* + ct$, where $f^* \in \mathbb{F}_q[\mathbf{x}, y]$, $c \in \mathbb{F}_q^\times$, and $\deg(f^*) = d'$. By making the substitution $t \mapsto -c^{-1}t$, we may assume $c = -1$. Consider any $\mathbf{a} \in \overline{\mathbb{F}_q}^n$ such that $g_{\mathbf{a}}^* = g^*(a_1x, \dots, a_nx, y)$ is reducible over $\overline{\mathbb{F}_q}(t)^n$. We claim that $g_{\mathbf{a}}^*$ factorizes into linear factors over $\overline{\mathbb{F}_q}(t)$. To see this, note that $f_{\mathbf{a}}^* = f^*(a_1x, \dots, a_nx, y)$ is a decomposable polynomial over $\overline{\mathbb{F}_q}$ of degree d' by Lemma 2.8 and the fact that $g_{\mathbf{a}}^* = f_{\mathbf{a}}^* - t$ is reducible over $\overline{\mathbb{F}_q}(t)$. So we may write $f_{\mathbf{a}}^* = \alpha(\beta)$ where $\alpha \in \overline{\mathbb{F}_q}[z]$, $\beta \in \overline{\mathbb{F}_q}[\mathbf{x}, y]$, and $\deg(\alpha) > 1$. As d' is prime, we must have $\deg(\alpha) = d'$ and $\deg(\beta) = 1$. As α is univariate, $\alpha - t$ factorizes into linear factors $\alpha_1, \dots, \alpha_{d'}$ over $\overline{\mathbb{F}_q}(t)$. So $g_{\mathbf{a}}^* = \alpha(\beta) - t = (\alpha - t)(\beta)$ factorizes into the linear factors $\alpha_1(\beta), \dots, \alpha_{d'}(\beta)$ over $\overline{\mathbb{F}_q}(t)$.

This observation shows that there is only one bad factorization pattern to rule out, namely, the complete factorization into linear factors. This allows us to save a factor of $2^{d-1} - 1$ and reduce the error probability in Lemma 5.2 from $(2^d - 1)\delta + \delta$ to $\delta + \delta = 2\delta$. The bound on the statistical distance between $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ in (3) now becomes $2\delta + 2\varepsilon' = 2C_0(2d - 1)/q + 2C_1d^2/\sqrt{q}$, which is bounded by ε provided that $q \geq C(d^4/\varepsilon^2)$ and $C > 0$ is a large enough absolute constant. ◀

6 Open Problems

We conclude with some open problems. The most obvious one is reducing the required field size in our construction. Using Bogdanov's [6] paradigm, it seems necessary for the field to be of size at least polynomial in d , since this argument relies on the Weil bound (and indeed, as mentioned in Section 1, the seed lengths of the known constructions over small fields like \mathbb{F}_2 are worse). Still, one could hope to obtain seed length $O(d \log n)$ with q being polynomial in d , and not exponential in d . In our construction, q is exponential in d due to the need to apply a union bound over all possible vectors in $\{0, 1\}^d$ characterizing the factorization pattern of $f_{\mathbf{a}}$. It could very well be that there is a more clever argument that rules out

multiple vectors at once. We also mention again Question 4.9. As explained in Section 4, improved upper bounds on the quantity $N^*(d, D, \mathbb{F})$ in that question would improve the field size required by our construction.

A related open problem is removing the requirement that the characteristic of \mathbb{F}_q is at least $d(d-1)+1$. This requirement comes from using Lecerf's [21] arguments (dating back to Gao [15] and Ruppert [26, 27]). We remark that our construction can be adapted to work in arbitrary characteristics $p > 0$ by employing the analysis in [8] and increasing the accuracy parameter σ of Hensel lifting to $d(d-1)+1$; however, this leads to a larger seed length of $O(d^2 \log n + \log q)$. A proof of this result will be provided in the full version of this paper.

Finally, low-degree polynomials form a natural “weak” class of polynomials. However, rather than assuming bounds on the degree of polynomials, one can also consider other weak classes of polynomials, where the restriction comes from bounding their algebraic circuit complexity. This forms another interesting avenue for generalizing the results on PRGs for low-degree polynomials. As an analogy, in the context of Boolean computation, the problem of constructing explicit PRGs for weak computational classes (such as bounded-depth circuits or read-once oblivious branching programs) is well studied (see [34]). For algebraic computational models, however, much less is known. Most of the research in this area has focused on constructing *hitting sets* of limited models of algebraic circuits (see [31, 28, 29] for some surveys on this topic), due to the relation to the famous Polynomial Identity Testing Problem. To the best of our knowledge, there is no known explicit construction of PRGs for any natural class of algebraic computation. A concrete and intriguing open problem is to explicitly construct PRGs for the class of sparse polynomials, for which, as described in the references above, there are many known explicit constructions of hitting sets.

References

- 1 Noga Alon, Ido Ben-Eliezer, and Michael Krivelevich. Small sample spaces cannot fool low degree polynomials. In *Proceedings of the 12th International Workshop on Randomization and Computation (RANDOM 2008)*, volume 5171 of *Lecture Notes in Computer Science*, pages 266–275. Springer, 2008. doi:10.1007/978-3-540-85363-3_22.
- 2 Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple construction of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. doi:10.1002/RSA.3240030308.
- 3 I. V. Arzhantsev and A. P. Petravchuk. Closed polynomials and saturated subalgebras of polynomial algebras. *Ukrainian Mathematical Journal*, 59(12):1783–1790, 2007.
- 4 Avraham Ben-Aroya and Amnon Ta-Shma. Constructing small-bias sets from algebraic-geometric codes. *Theory of Computing*, 9:253–272, 2013. doi:10.4086/TOC.2013.V009A005.
- 5 Arnaud Bodin, Pierre Debes, and Salah Najib. Indecomposable polynomials and their spectrum. *Acta Arithmetica*, 139(1):79–100, 2009.
- 6 Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 21–30. ACM, 2005. doi:10.1145/1060590.1060594.
- 7 Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM Journal on Computing*, 39(6):2464–2486, 2010. doi:10.1137/070712109.
- 8 A. Bostan, G. Lecerf, B. Salvy, É. Schost, and B. Wiebelt. Complexity issues in bivariate polynomial factorization. In *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation, ISSAC 2004*, pages 42–49. Association for Computing Machinery, 2004.
- 9 Guillaume Cheze and Salah Najib. Indecomposability of polynomials via Jacobian matrix. *Journal of Algebra*, 324(1):1–11, 2010.

- 10 Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. *Electronic Colloquium on Computational Complexity*, TR13-155, 2013. [arXiv:TR13-155](#).
- 11 David A. Cox, John B. Little, and Donal O’Shea. *Ideals, Varieties and Algorithms*. Undergraduate Texts in Mathematics. Springer, 2007. [doi:10.1007/978-0-387-35651-8](#).
- 12 Harm Derksen and Emanuele Viola. Fooling polynomials using invariant theory. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 399–406. IEEE, 2022. [doi:10.1109/FOCS54457.2022.00045](#).
- 13 David Eisenbud. *Commutative Algebra: With a View Toward Algebraic Geometry*. Springer Science & Business Media, 1995.
- 14 Guy Even, Oded Goldreich, Michael Luby, Noam Nisan, and Boban Velickovic. Efficient approximation of product distributions. *Random Structures & Algorithms*, 13(1):1–16, 1998. [doi:10.1002/\(SICI\)1098-2418\(199808\)13:1<1::AID-RSA1>3.0.CO;2-W](#).
- 15 Shuhong Gao. Factoring multivariate polynomials via partial differential equations. *Mathematics of Computation*, 72(242):801–822, 2003. [doi:10.1090/S0025-5718-02-01428-X](#).
- 16 Venkatesan Guruswami and Chaoping Xing. Hitting sets for low-degree polynomials with optimal density. In *Proceedings of the IEEE 29th Conference on Computational Complexity, CCC 2014*, pages 161–168. IEEE Computer Society, 2014. [doi:10.1109/CCC.2014.24](#).
- 17 Pooya Hatami and William Hoza. Theory of unconditional pseudorandom generators. *Electronic Colloquium on Computational Complexity*, TR23-019, 2023. [arXiv:TR23-019](#).
- 18 Erich L. Kaltofen. Effective Noether irreducibility forms and applications. *Journal of Computer and System Sciences*, 50(2):274–295, 1995. [doi:10.1006/JCSS.1995.1023](#).
- 19 Adam R. Klivans and Daniel A. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 216–223. ACM, 2001. [doi:10.1145/380752.380801](#).
- 20 Grégoire Lecerf. Sharp precision in Hensel lifting for bivariate polynomial factorization. *Mathematics of Computation*, 75(254):921–933, 2006.
- 21 Grégoire Lecerf. Improved dense multivariate polynomial factorization algorithms. *Journal of Symbolic Computation*, 42(4):477–494, 2007.
- 22 Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. *Theory of Computing*, 5(1):69–82, 2009. [doi:10.4086/TOC.2009.V005A003](#).
- 23 Chi-Jen Lu. Hitting set generators for sparse polynomials over any finite fields. In *Proceedings of the 27th Conference on Computational Complexity, CCC 2012*, pages 280–286. IEEE Computer Society, 2012. [doi:10.1109/CCC.2012.20](#).
- 24 Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. [doi:10.1137/0222053](#).
- 25 Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Matematicheskie Zametki*, 41:598–607, 1987. English translation in *Mathematical Notes of the Academy of Sci. of the USSR*, 41(4):333-338, 1987. [doi:10.1007/BF01137685](#).
- 26 Wolfgang Ruppert. Reduzibilität ebener kurven. *Journal für die reine und angewandte Mathematik*, 1986(369):167–191, 1986. [doi:doi:10.1515/crll.1986.369.167](#).
- 27 Wolfgang M. Ruppert. Reducibility of polynomials $f(x, y)$ modulo p . *Journal of Number Theory*, 77(1):62–70, 1999. [doi:10.1006/jnth.1999.2381](#).
- 28 Nitin Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS*, 99:49–79, 2009.
- 29 Nitin Saxena. Progress on polynomial identity testing-II. *Perspectives in Computational Complexity: The Somenath Biswas Anniversary Volume*, pages 131–146, 2014. [doi:10.1007/978-3-319-05446-9_7](#).
- 30 Igor Shafarevich. *Basic Algebraic Geometry 1: Varieties in Projective Space*. Springer, 1994.

- 31 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.
- 32 Roman Smolensky. On representations by low-degree polynomials. In *34th Annual Symposium on Foundations of Computer Science*, pages 130–138. IEEE Computer Society, 1993. doi:10.1109/SFCS.1993.366874.
- 33 Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 238–251. ACM, 2017. doi:10.1145/3055399.3055408.
- 34 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012. doi:10.1561/0400000010.
- 35 Emanuele Viola. The sum of D small-bias generators fools polynomials of degree D . *computational complexity*, 18(2):209–217, 2009. doi:10.1007/S00037-009-0273-5.
- 36 André Weil. Numbers of solutions of equations in finite fields. *Bulletin of the American Mathematical Society*, 55:497–508, 1949.