Electronic Colloquium on Computational Complexity, Report No. 16 (2026)

Optimal PRGs for Low-Degree Polynomials over Polynomial-Size Fields

Gil Cohen* Dean Doron† Noam Goldgraber‡

Abstract

Pseudorandom generators (PRGs) for low-degree polynomials are a central object in pseudorandomness, with applications to circuit lower bounds and derandomization. Viola's celebrated construction [Vio09] gives a PRG over the binary field, but with seed length exponential in the degree d. This exponential dependence can be avoided over sufficiently large fields. In particular, Dwivedi, Guo, and Volk [DGV24] constructed PRGs with optimal seed length over fields of size exponential in d. The latter builds on the framework of Derksen and Viola [DV22], who obtained optimal-seed constructions over fields of size polynomial in d, although growing with the number of variables n.

In this work, we construct the first PRG with *optimal seed length* for degree-d polynomials over *fields of polynomial size*, specifically $q \approx d^4$, assuming, as in [DGV24], sufficiently large characteristic. Our construction follows the framework of [DV22, DGV24] and reduces the required field size by replacing the hitting-set generator used in prior work with a new pseudorandom object.

We also observe a threshold phenomenon in the field-size dependence. Specifically, we prove that constructing PRGs over fields of sublinear size, for example $q = d^{0.99}$ where q is a power of two, would already yield PRGs for the binary field with comparable seed length via our reduction, provided that the construction imposes no restriction on the characteristic. While a breakdown of existing techniques has been noted before, we prove that this phenomenon is inherent to the problem itself, irrespective of the technique used.

*Tel Aviv University. gil@tauex.tau.ac.il. Supported by ERC starting grant 949499 and by the Israel Science Foundation grant 2989/24.

†Ben Gurion University. deand@bgu.ac.il. Supported in part by NSF-BSF grant 2022644.

‡Ben Gurion University and Tel Aviv University. goldgrab@post.bgu.ac.il. Supported by NSF-BSF grant 2022644.

ISSN 1433-8092

1 Introduction

A *pseudorandom generator* (PRG) fooling a class of functions $\mathscr{C} \subseteq \Sigma^n \to \Sigma$ is a map $G\colon \{0,1\}^s \to \Sigma^n$ that stretches a uniform seed of $s \ll n$ bits into strings in Σ^n, such that the distribution of G *fools* any function $f \in \mathscr{C}$, in the sense that $f(G(\mathbf{U}_s))$ is close, in total-variation distance, to $f(\mathbf{U}_{\Sigma^n})$.[1] Constructing explicit PRGs with short seed for various function classes $\mathscr{C}$ is central to theoretical computer science, with various applications in complexity theory (prominently derandomization and circuit lower bounds), cryptography, and algorithm design.

A fundamental and well-studied class of functions is that of *low degree polynomials*.

Definition 1.1. *We say that $G\colon \{0,1\}^s \to \mathbb{F}_q^n$ is a PRG for n-variate polynomials of total degree at most d over a finite field $\mathbb{F}_q$ with error ε if for every such polynomial f, the distributions $f(G(\mathbf{U}_s))$ and $f(\mathbf{U}_{\mathbb{F}_q^n})$ are ε-close in total variation distance. That is,*

$$\frac{1}{2}\sum_{a\in\mathbb{F}_q}\left|\Pr_{\mathbf{x}\in\mathbb{F}_q^n}[f(\mathbf{x})=a]-\Pr_{t\in\{0,1\}^s}[f(G(t))=a]\right|\le\varepsilon.$$

The seed length *of G is s, and we say that G is* explicit *if for any $t \in \{0,1\}^s$, $G(t)$ can be computed in time* $\mathrm{poly}(n,d,\log q,\log 1/\varepsilon)$.

PRGs for low-degree polynomials have been extensively studied for more than three decades, with the natural goal of *minimizing the seed length*. Moreover, over the years it has become apparent that constructing PRGs over small fields—of constant size, independent of the degree d—is more challenging than in the regime where the field size is allowed to be polynomial in d, which permits the use of deep results such as the Weil bound.

Already the case $d = 1$, which corresponds to *small-biased generators*, is extremely interesting and has found numerous applications in pseudorandomness and derandomization. In this setting, we have constructions with seed length that is optimal up to constant factors (see, e.g., [NN93, ABN+92, AGHP92, AMN98, BT13, Ta-17, CC25]), typically over any field size (although the $\mathbb{F}_2$ case is the most widely studied). For arbitrary d, one can show a lower bound of $s = \Omega\big(d\log(n/d)+\log(1/\varepsilon)+\log q\big)$ (see, e.g., [ABEK08]), and the probabilistic method guarantees the existence of a construction achieving these parameters. From now on, we refer to this as an optimal seed (ignoring constant factors). Obtaining explicit constructions is more challenging, and prior work has developed along two main strands.

PRGs over an arbitrary field. Over any finite field, and in particular over (what turned out to be) the most challenging case of $\mathbb{F}_2$, a sequence of works [LVW93, Vio07, BV10, Lov09, Vio09]

[1] Here and throughout, for an integer s, $\mathbf{U}_s$ denoted the uniform distribution over $\{0,1\}^s$, and for a set A, $\mathbf{U}_A$ denotes the uniform distribution over A.

culminated in Viola's celebrated explicit PRG with seed length $O(d \log n + d \cdot 2^d \log(q/\varepsilon))$ [Vio09]. The generators of [BV10, Lov09, Vio09] are obtained via the Bogdanov–Viola [BV10] framework: In order to fool degree-d polynomials, sum $\ell = \ell(d)$ independent copies of a small-bias generator. Viola [Vio09] proved that $\ell(d) = d$ suffices, however the error of the small-bias PRG needs to be very small, namely ε^{2^d} for a designated error ε, leading to the 2^d factor in the seed length. Note that when $d = \Omega(\log n)$, the seed length becomes trivial, and indeed, achieving any nontrivial PRGs over $\mathbb{F}_2$ for degrees greater than $\log n$ would yield breakthroughs in circuit complexity, via the Razborov–Smolensky connection between constant-depth circuits and low-degree polynomials [Raz87, Smo93].

PRGs over large fields. When $q \gg d$, better results are known, and we can handle much larger degrees with a relatively short seed length. Bogdanov [Bog05] introduced a technique for constructing PRGs from the weaker object of *hitting set generators* (HSGs; see Definition 1.5 for the formal definition). This approach is based on reducing the number of variables of the polynomial, while preserving the factorization structure of it (i.e., preserving irreducibility of its factors). Combined with subsequent improvements in HSG constructions following due to Lu [Lu12], Cohen and Ta-Shma [CT13], and Guruswami and Xing [GX14], Bogdanov's technique yields a PRG with seed length $O(d^4 \log n + \log q)$, provided that $q \geq d^6/\varepsilon^2$.

More recently, Derksen and Viola [DV22] introduced a powerful new approach based on techniques from algebraic geometry and invariant theory. For sufficiently large $q \geq (d^4 n^{0.001})/\varepsilon^2$, they achieve optimal seed length $O(d \log(dn) + \log q)$. For $q \geq (d \log n)^4/\varepsilon^2$, they obtain a sub-optimal seed length of $O(d \log n \cdot \log(d \log n) + \log q)$. We will discuss their approach, based on the preservation of indecomposability instead of irreducibility, in Section 1.2.

Recently, Dwivedi, Guo, and Volk [DGV24] were able to remove the dependence on n in the field-size requirement needed to obtain an optimal-seed PRG, albeit with an exponential dependence on d. Specifically, they achieve seed length $O(d \log n + \log q)$ whenever $q \geq d2^d/\varepsilon + d^4/\varepsilon^2$ and the field characteristic is $\Omega(d^2)$.[2] We also discuss their technique, which combines ideas from [DV22] with a derandomization approach inspired by [Bog05], in Section 1.2.

1.1 Our Result

In our work, we construct an explicit PRG with *optimal seed length* for field sizes q that are only *polynomial in* d – an exponential improvement over field size requirement in [DGV24].

[2] Interestingly, if one only cares about fooling *primes* degrees, [DGV24] show that $q = \Omega(d^4/\varepsilon^2)$ suffices.

Theorem 1.2 (see also Theorem 5.2)**.** *For every $n, d \in \mathbb{N}$, a prime power q, and $\varepsilon > 0$, satisfying $q = \Omega((d \log d)^4/\varepsilon^2)$ and $\mathrm{char}(\mathbb{F}_q) = \Omega(d^2)$, there exists an explicit PRG $G\colon \{0,1\}^s \to \mathbb{F}_q^n$ for n-variate polynomials of degree at most d over $\mathbb{F}_q$ with error ε and seed length $s = O(d \log n + \log q)$.*

Compared to [DV22] and [DGV24], our construction improves upon both works simultaneously: we achieve optimal seed length already for $q \geq \mathrm{poly}(d)$ (rather than $q \geq \exp(d)$ as in [DGV24]), and we avoid the dependence of q on n present in [DV22].

A Threshold Phenomenon for PRGs for Low-Degree Polynomials

As discussed, the study of PRGs for low-degree polynomials has split into two branches: PRGs for constant field size, most notably the binary field, and PRGs for fields whose size is sufficiently large as a function of the degree d (and, for some constructions, also of the number of variables n). These two branches rely on fundamentally different techniques. Our result falls into the second line of work: we construct an optimal-seed PRG that works whenever q is at least roughly d^4.

It was observed in [DV22] that techniques developed for large fields, such as those relying on the Weil bound, break down in the small-field regime. Perhaps surprisingly, our second contribution shows that there is, in fact, an inherent threshold phenomenon. When q is a power of two, we prove that improving the quartic dependence of q on d to a sublinear one—namely, $q = d^{1-\tau}$ for some fixed constant $\tau > 0$—would immediately yield a PRG construction over the binary field with seed length comparable to that of the PRG we started with. For our reduction to apply, the PRG must not impose any restriction on the characteristic of the field, as is the case, for example, in the Derksen–Viola construction [DV22].

This implies that there cannot be incremental progress toward constructing PRGs over the binary field: once the current quartic dependence is improved to a sublinear one, and provided there are no restrictions on the characteristic, one immediately obtains a comparable PRG over the binary field. Our reduction holds in greater generality and, in particular, applies to fields of any odd characteristic.

Theorem 1.3 (see also Proposition 6.1)**.** *Assume that for any n, d, q such that $q \geq d^{1-\eta}$ for some constant $\eta \in (0,1)$ there exists an explicit PRG for n-variate polynomials of total degree at most d over $\mathbb{F}_q$, with error 0.1 and seed length $O(d^{O(1)} \log n + \log q)$. Then, there also exists an explicit PRG for n-variate polynomials of total degree at most d over $\mathbb{F}_2$, with error 0.1, and seed length $O(d^{O(1)} \log n)$.*

Constructing G_2, the desired PRG over $\mathbb{F}_2$, is simple: Set $q = d^c$ to be a power of 2, where $c = c(\eta)$, and let G_q be our hypothesized PRG over $\mathbb{F}_q$. Then, each element of G_2 is obtained by

taking the *absolute trace* of each coordinate of an element of G_q. That is, $G_2(z)_i = \mathrm{Tr}(G_q(z)_i)$. We defer the (easy) proof to Section 6, and proceed to giving an overview of Theorem 1.2.

1.2 Proof Overview

We begin with describing the framework introduced by Derksen and Viola [DV22]. Let $\mathbf{x} = (x_1, \ldots, x_n)$. A polynomial $f \in \mathbb{F}_q[\mathbf{x}]$ is called *indecomposable* if it cannot be written as $f = g \circ h$, where $h \in \mathbb{F}_q[\mathbf{x}]$ and $g \in \mathbb{F}_q[t]$ is a univariate polynomial with $\deg(g) \geq 2$. The notion of indecomposability is interesting partially because we understand the distribution of its image. This is formalized by the following lemma.

Lemma 1.4 ([DV22], Lemma 12)**.** *There exists an absolute constant $c > 0$ such that the following holds: Suppose $f \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable over $\mathbb{F}_q$. Then, $f(\mathbf{U}_{\mathbb{F}_q^n})$ is ε-close to $\mathbf{U}_{\mathbb{F}_q}$, where $\varepsilon = c \cdot d^2/\sqrt{q}$.*

The general approach in our work, following the works [Bog05, GX14, DV22, DGV24], is to restrict f to a carefully chosen subset of $\mathbb{F}_q^n$ while preserving some algebraic property, closely related to its output distribution. Specifically, let $f \in \mathbb{F}_q[\mathbf{x}]$ be an *indecomposable* polynomial of degree at most d. Suppose we can find polynomials $p_1, \ldots, p_n \in \mathbb{F}_q[w_1, \ldots, w_\ell]$, with $\ell \ll n$, such that for every such f, the composed polynomial $f \circ (p_1, \ldots, p_n) \in \mathbb{F}_q[w_1, \ldots, w_\ell]$ is indecomposable. Let $\mathbf{w} = (w_1, \ldots, w_\ell)$, let $\mathbf{p} = (p_1, \ldots, p_n)$ be the restriction map, and let $\deg \mathbf{p} = \max_i\{\deg p_i\}$. Now, Lemma 1.4 applies both to f and to its restriction via $f \circ \mathbf{p}$. In particular, the distribution of $f(\mathbf{U}_{\mathbb{F}_q^n})$ is $O(d^2/\sqrt{q})$-close to uniform over $\mathbb{F}_q$, and similarly the distribution of $f(\mathbf{p}(\mathbf{U}_{\mathbb{F}_q^\ell}))$ is $O((d \cdot \deg \mathbf{p})^2/\sqrt{q})$-close to $\mathbf{U}_{\mathbb{F}_q}$. This means that the function $G : \mathbb{F}_q^\ell \to \mathbb{F}_q^n$ defined by

$$G(\mathbf{w}) = \mathbf{p}(\mathbf{w})$$

is a PRG for indecomposable polynomials of degree at most d with seed length $\ell \log q$, and error

$$\varepsilon = O((d \cdot \deg \mathbf{p})^2/\sqrt{q}). \tag{1.1}$$

For an *arbitrary* n-variate polynomial f of degree at most d, we can always write $f = g \circ h$, where $h \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable and $g \in \mathbb{F}_q[t]$ is univariate. h is an n-variate indecomposable polynomial of degree at most d, therefore the distributions $h(\mathbf{U}_{\mathbb{F}_q^n})$ and $h(G(\mathbf{U}_{\mathbb{F}_q^\ell}))$ are both close to $\mathbf{U}_{\mathbb{F}_q}$. Hence, the distributions $f(\mathbf{U}_{\mathbb{F}_q^n}) = g(h(\mathbf{U}_{\mathbb{F}_q^n}))$ and $f(G(\mathbf{U}_{\mathbb{F}_q^\ell})) = g(h(G(\mathbf{U}_{\mathbb{F}_q^\ell})))$ are both close to $g(\mathbf{U}_{\mathbb{F}_q})$. This shows that G is actually a PRG for all n-variate polynomials of degree at most d.

The challenge is, of course, to find low-degree $p_1, \ldots, p_n$ that preserve indecomposability for all indecomposable n-variate polynomials f of degree at most d.

1.2.1 The Derksen-Viola restriction map

A main contribution of [DV22] is the explicit construction of such restriction polynomials $p_1, \ldots, p_n$, which proceeds as follows. Let $M_1, \ldots, M_n$ be distinct monomials in m variables. Consider ℓ independent copies of these variables, and denote by $M_i^{[j]}$ the monomial M_i evaluated on the variables from the j-th copy. Then define

$$p_i \;=\; M_i^{[1]} + \cdots + M_i^{[\ell]}\,.$$

Using tools from invariant theory, [DV22] show that for a suitable choice of parameters and monomials, the resulting substitution $\mathbf{p} = (p_1, \ldots, p_n)$ preserves indecomposability. This allows them to construct a PRG with optimal seed length $O(d\log(dn) + \log q)$, assuming $q = \Omega(d^4 n^{0.001}/\varepsilon^2)$, or alternatively a PRG with seed length $O(d\log n \cdot \log(d\log n) + \log q)$, assuming $q = \Omega((d\log n)^4/\varepsilon^2)$.

1.2.2 Constructing the restriction map via hitting set generators

In contrast to [DV22], approach, Bogdanov [Bog05] and Dwivedi, Guo, and Volk [DGV24] adopt a derandomization approach. Rather than constructing a single restriction map $\mathbf{p}$ that preserves indecomposability for all polynomials of degree at most d, they use a carefully designed distribution over restriction maps $\mathbf{p}$. This distribution has the property that for every indecomposable polynomial f of degree at most d, the composition $f \circ \mathbf{p}$ remains indecomposable with high probability. In both works, the restriction polynomials are linear, and the distribution is designed using a pseudorandom object called a *hitting set generator* (HSG).

Definition 1.5. *A function $H\colon T \to \mathbb{F}_q^n$ is a hitting set generator (HSG) with density δ for n-variate polynomials of degree at most d over $\mathbb{F}_q$ if for every such $f \neq 0$,*

$$\Pr_{t\in T}[f(H(t)) \neq 0] \geq 1 - \delta.$$

The seed length *of the HSG is $\log_2 |T|$, and we say that H is* explicit *if for any $t \in T$, $H(t)$ can be computed in time* $\mathrm{poly}(n, d, \log q, \log 1/\delta)$.

Since our construction builds primarily on [DGV24], we proceed by describing their construction in more detail.

Let

$$p_i(x,y) = \begin{cases} \beta_i x + \alpha_i y & 1 \leq i \leq n-1 \\ y & i = n \end{cases},$$

where the vectors $\alpha, \beta \in \mathbb{F}_q^n$ are sampled from a HSG for $(n-1)$-variate polynomials of degree at most $O(d)$ (where the implicit constant is absolute). This transformation can also be written

as follows: Let s_α be the ring automorphism of $\mathbb{F}_q[\mathbf{x}]$ such that

$$s_\alpha(f(\mathbf{x})) = f(x_1 + \alpha_1 x_n, \ldots, x_{n-1} + \alpha_{n-1} x_n, x_n).$$

Let $r_\beta : \mathbb{F}_q[\mathbf{x}] \to \mathbb{F}_q[x,y]$ be the homomorphism such that $r_\beta(f(\mathbf{x})) = f(\beta_1 x, \ldots, \beta_{n-1} x, y)$. Then, for all $f \in \mathbb{F}_q[\mathbf{x}]$ we have

$$f \circ \mathbf{p} = r_\beta \circ s_\alpha \circ f.$$

It is shown in [DGV24] that for every degree-d $f \in \mathbb{F}_q[\mathbf{x}]$, there exists a polynomial $B \in \mathbb{F}_q[x_1, \ldots, x_{n-1}]$ of degree at most d with the following property. For any α satisfying $B(\alpha) \neq 0$, the polynomial

$$s_\alpha(f) - t \in \mathbb{F}_q(t)[\mathbf{x}]$$

satisfies *Hypothesis (H)* (up to multiplication by an element of $\mathbb{F}_q^\times$). Hypothesis (H) is a condition required for applying *Lecerf's technique*; see Definition 3.1 for details. Thus, by picking α using a HSG, it is promised that the polynomial $s_\alpha(f) - t$ satisfies Hypothesis (H) with high probability.

They proceed by showing, building on results of Lecerf [Lec06, Lec07], that if $g \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable and $g - t$ satisfies Hypothesis (H), then the restriction $r_\beta(g)$ remains indecomposable with high probability.

The proof proceeds roughly as follows: Let $\mathbb{K} = \overline{\mathbb{F}_q(t)}$ be the algebraic closure of $\mathbb{F}_q(t)$. Let $g \in \mathbb{K}[\mathbf{x}]$ be a polynomial satisfying Hypothesis (H). For such a polynomial g, assuming that $\operatorname{char} \mathbb{F}_q = \operatorname{char} \mathbb{K} > d(d-1)$, Lecerf constructed a linear system of equations D_g in d variables, with the coefficients being $(n-1)$-variate polynomials in $\mathbb{K}[z_1, \ldots, z_{n-1}]$ of degree at most $2d-1$, with the following properties:

1. The number of irreducible factors of g equals the dimension of the space of solutions of D_g.
2. For $\beta \in \mathbb{K}^{n-1}$, let D_g^β be the system of equations D_g after evaluating $(z_1, \ldots, z_{n-1}) = \beta$. The number of irreducible factors of $r_\beta(g)$ equals the dimension of the space of solutions of D_g^β.
3. There are sets $S_1 \subseteq S_2 \subseteq \{0,1\}^d$ such that the space of solutions of D_g is spanned by S_1, and the space of solutions of D_g^β is spanned by S_2.

The condition that a polynomial $f \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable is equivalent to $g = f - t \in \mathbb{K}[\mathbf{x}]$ being irreducible (see Lemma 2.9). Hence, preserving indecomposability of f is equivalent to preserving the space of solutions of D_g in the reduction to D_g^β.

Let $f \in \mathbb{F}_q[\mathbf{x}]$ be an indecomposable polynomial satisfying Hypothesis (H). Then $g = f - t \in \mathbb{K}[\mathbf{x}]$ is irreducible, and by Items 1 and 3 the solution space of D_g is spanned by a single vector

$v \in \{0,1\}^d$. For every $v \neq w \in \{0,1\}^d$, there is an equation in D_g such that w does not satisfy, i.e., there are polynomials $P_1^w, \ldots, P_d^w \in \mathbb{K}[z_1, \ldots, z_{n-1}]$ of degree at most $2d-1$ such that

$$Q_w := \sum_{i=1}^{d} w_i \cdot P_i^w \neq 0.$$

Notice that $Q_w \in \mathbb{K}[z_1, \ldots, z_{n-1}]$ is a polynomial of degree at most $2d-1$ as well. By Item 2, if $Q_w(\beta) \neq 0$ for all $v \neq w \in \{0,1\}^d$, then the polynomial $r_\beta(g)$ remains irreducible, which means that the polynomial $r_\beta(f)$ is indecomposable.

This implies that for every indecomposable polynomial that satisfies Hypothesis (H), there are $2^d - 1$ polynomials in $\mathbb{K}[z_1, \ldots, z_{n-1}]$ of degree at most $2d-1$ such that if each one of them does not vanish at β, then $r_\beta(f)$ remains indecomposable. Hence, by picking β using an HSG for $n-1$-variate polynomials of degree at most $2d-1$ with small enough density, the probability that $r_\beta(f)$ remains indecomposable is high. Thus, using the union bound, the probability that that the substitution $f \circ \mathbf{p}$ preserves indecomposability is at least $1 - (2^d - 1)\delta$.

Any HSG for polynomials of degree at most $2d-1$ over $\mathbb{F}_q$ has density at most $1-\delta$ with $\delta = \Theta(d/q)$, as the density of zeroes of a degree-d polynomial can be $\Omega(d/q)$. Consequently, the analysis yields a meaningful guarantee only when $q = \Omega(d\,2^d)$, which is where the exponential dependence in d shows up in the field size.

1.2.3 Our construction via polynomial hitting set generators

In this work, we further refine this approach to achieve optimal seed-length PRGs for much smaller field sizes. Our key idea is to choose β not as a vector over $\mathbb{F}_q$, but as a *vector of polynomials* of bounded degree in a small number of variables. Let A denote the size of the set from which each polynomial coordinate is chosen. By the Schwartz–Zippel lemma, any nonzero degree-d polynomial vanishes on at most a d/A fraction of this distribution. This suggests that one can derandomize this construction – analogously to standard HSGs over $\mathbb{F}_q$ – to obtain significantly higher-density hitting sets. To achieve this, we introduce the notion of a *polynomial hitting set generator* (PHSG).

Definition 1.6. *Let $\mathbb{F}$ be a finite field. A* polynomial hitting set generator (PHSG) *with density $1-\delta$ for n-variate polynomials of degree at most d over $\mathbb{F}$ with ℓ-variate polynomial evaluation points of degree at most h is a map*

$$H \colon T \to (\mathbb{F}^{\leq h}[w_1, \ldots, w_\ell])^n$$

from a finite set $T \neq \emptyset$ such that for every such nonzero polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most d,

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

The quantity $\log|T|$ *is called the seed length of* H.

In Section 4, we show that any HSG for polynomials over a field extension $\mathbb{F}_{q^k}$ can be simply turned into a PHSG over $\mathbb{F}_q$. Consequently, existing constructions of HSGs immediately yield PHSGs.[3] So indeed, in our construction we choose each p_i to be of the form

$$p_i(x,y,w_1,\ldots,w_\ell) = b_i(w_1,\ldots,w_\ell)\cdot x + \alpha_i \cdot y,$$

where $\alpha \in \mathbb{F}_q^n$ is sampled from an HSG $H_1 \colon T \to \mathbb{F}_q^n$ with density $1-\delta_1$, and $\mathbf{b} \in \mathbb{F}_q^{\leq h}[w_1,\ldots,w_\ell]^n$ is sampled from a PHSG $H_2 \colon T_2 \to (\mathbb{F}^{\leq h}[w_1,\ldots,w_\ell])^n$ with density $1-\delta_2$.

Adapting the analysis of [DGV24], assuming $\operatorname{char}\mathbb{F}_q > d(d-1)$ we show that the probability that $\mathbf{p}$ preserves indecomposability is at least

$$1-\delta_1-(2^{d-1}-1)\delta_2.$$

The analysis extends naturally to the setting where the evaluation points are themselves polynomials rather than field elements, i.e., are taken from a field extension. In this case, one can show that the polynomial $r_{\mathbf{b}}(f)-t$ is irreducible as an element of $\overline{\mathbb{F}_q(t,w_1,\ldots,w_\ell)}[x,y]$. We then use the structure of the polynomial and apply Gauss's lemma to deduce that this polynomial is in fact irreducible over $\overline{\mathbb{F}_q(t)}[w_1,\ldots,w_\ell,x,y]$. This implies that the composed polynomial $f\circ\mathbf{p} \in \mathbb{F}_q[w_1,\ldots,w_\ell,x,y]$ is indecomposable.

In Section 4, we construct PHSGs with $\ell = h =: \log k$, seed length $d\log n + k\log q$ and density $1-\delta$, where $\delta \geq d/q^k$. Choosing $\delta_2 = O(d/q^k)$ with $k = d/\log q$, and $\delta_1 = O(d/q)$ gives

$$\mathbf{Pr}[f\circ\mathbf{p} \text{ is indecomposable}] = 1 - O(d/q).$$

This way we eliminate the requirement $q \geq \Omega\left(d2^d\right)$. This comes at the cost of increasing the number of variables in $\mathbf{p}$ from 2 to ℓ, and using a PHSG in addition to the HSG. Those increase the seed length only by a constant factor: the PHSG instantiated with our parameters requires seed length of $O(d\log n + \ell\log q)$, and the uniform distribution over $\mathbb{F}_q^\ell$ requires additional $\ell\log q$ random bits. As $\ell\log q = O(d+\log q)$, the resulting PRG seed length is $O(d\log n + \log q)$.

The requirement $q \geq (d\log d)^4/\varepsilon^2$ follows from Equation (1.1), since

$$\deg\mathbf{p} \leq h = \ell = O(\log d).$$

[3] An obstacle in that approach is that explicitly constructing an HSG over a field extension requires an explicit representation of $\mathbb{F}_{q^k}$, which in turn necessitates the construction of irreducible polynomials over $\mathbb{F}_q$, for which the best-known deterministic algorithm would be too costly. To overcome this issue, we pick the required irreducible polynomials at random, and use a sampler to amplify the success probability, incurring only a small additional randomness cost. Fortunately, this incurs only a constant-factor increase in the overall seed length.

As explained earlier, the requirement $\mathrm{char}(\mathbb{F}_q) > d(d-1)$ follows from the use of Lecerf's technique.

To sum up, our pseudorandom generator $G : T_1 \times T_2 \times \mathbb{F}_q \times \mathbb{F}_q \times \mathbb{F}_q^{\ell} \to \mathbb{F}_q^n$ is defined by

$$G(r,s,u,v,\mathbf{t}) = (H_1(r)_1(\mathbf{t}) \cdot v + H_2(s)_1 \cdot u, \ldots, H_1(r)_n(\mathbf{t}) \cdot v + H_2(s)_n \cdot u, u), \tag{1.2}$$

where $\mathbf{t} = (t_1, \ldots, t_\ell)$, and $u, v \in \mathbb{F}_q$.

2 Preliminaries

We denote by $\mathbb{N}$ the set of nonnegative integers, including 0. Throughout the paper, boldface letters denote vectors; for example, $\mathbf{x} = (x_1, \ldots, x_n)$. For a multi-index $\mathbf{i} = (i_1, \ldots, i_n) \in \mathbb{N}^n$, we use the standard notation $\mathbf{x}^{\mathbf{i}} := x_1^{i_1} \cdots x_n^{i_n}$.

Resultants. Let R be a commutative ring and let $f(x) = \sum_{i=0}^{d_1} a_i x^i$ and $g(x) = \sum_{i=0}^{d_2} b_i x^i$ be polynomials in $R[x]$, with $a_{d_1} \neq 0$ and $b_{d_2} \neq 0$. Suppose that $d_1 + d_2 \geq 1$.

Definition 2.1 (Sylvester matrix)**.** *The* Sylvester matrix $\mathrm{Syl}(f,g)$ *of f and g is the following $(d_1+d_2) \times (d_1+d_2)$ matrix defined over R:*

$$\begin{pmatrix}
a_0 & & & & b_0 & & & \\
a_1 & a_0 & & & b_1 & b_0 & & \\
a_2 & a_1 & \ddots & & b_2 & b_1 & \ddots & \\
\vdots & & \ddots & a_0 & \vdots & & \ddots & b_0 \\
 & \vdots & & a_1 & b_{d_2} & \vdots & & b_1 \\
a_{d_1} & & & & & b_{d_2} & & \\
 & a_{d_1} & & \vdots & & & \vdots & \\
 & & \ddots & & & & \ddots & \\
 & & & a_{d_1} & & & & b_{d_2}
\end{pmatrix}.$$

Definition 2.2 (resultant)**.** *The* resultant *of f and g, denoted* $\mathrm{Res}(f,g)$*, is defined as*

$$\mathrm{Res}(f,g) := \det\big(\mathrm{Syl}(f,g)\big) \in R.$$

The resultant satisfies the following property.

Lemma 2.3. *Let R be an integral domain with field of fractions $\mathbb{L}$. Then* $\mathrm{Res}(f,g) = 0$ *if and only if f and g have a common root in $\overline{\mathbb{L}}$.*

Definition 2.4 (formal power series)**.** *Let $\mathbb{F}$ be a field and let $\mathbf{x} = (x_1, \ldots, x_n)$ be indeterminates. The ring of* formal power series $\mathbb{F}[[\mathbf{x}]]$ *consists of all infinite sums*

$$\sum_{\mathbf{i} \in \mathbb{N}^n} a_{\mathbf{i}} \mathbf{x}^{\mathbf{i}} \qquad (a_{\mathbf{i}} \in \mathbb{F}),$$

with addition and multiplication defined formally.

2.1 Prior Results

We will use the optimal HSG (over large fields) of Guruswami and Xing.

Theorem 2.5 ([GX14], Theorem 5.1)**.** *There exists an absolute constant c such that for any n, d, q, δ, for which $q \geq c \cdot d/\delta$, there exists an efficiently computable HSG for n-variate polynomials of degree at most d over $\mathbb{F}_q$ with density $1 - \delta$ and seed length $O(d \log n + \log(1/\delta))$.*

We will also make use of the following lemma.

Lemma 2.6 ([DGV24], Fact 2.4)**.** *Let H 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 $\mathbb{K}$.*

2.2 Indecomposable Polynomials

The analysis of our construction use the notion of *indecomposable* polynomials and some of their properties.

Definition 2.7. *A non-constant polynomial $f \in \mathbb{F}[\mathbf{x}]$ is said to be* decomposable *over $\mathbb{F}$ if there exist $h \in \mathbb{F}[\mathbf{x}]$ and a univariate polynomial $g \in \mathbb{F}[z]$ such that $\deg(g) \geq 2$ and $f = g \circ h$. Otherwise, f is said to be* indecomposable.

We will use the following equivalences.

Lemma 2.8 ([BDN09], Theorem 4.2)**.** *A polynomial $f \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable over $\mathbb{F}$ if and only if it is indecomposable over $\overline{\mathbb{F}}$.*

Lemma 2.9 ([CN10], Lemma 7)**.** *Let $f \in \mathbb{F}[\mathbf{x}]$ be a non-constant polynomial over a field $\mathbb{F}$. Then f is indecomposable over $\overline{\mathbb{F}}$ iff $f - t$ is irreducible over $\overline{\mathbb{F}(t)}$, where t is a new formal variable.*

Using the Weil bound, Derksen and Viola have showed in [DV22] that indecomposable polynomials are approximately equidistributed over large enough fields.

Lemma 2.10 ([DV22], Lemma 12)**.** *There exists an absolute constant $c > 0$ such that the following holds: Suppose $f \in \mathbb{F}_q[\mathbf{x}]$ is indecomposable over $\mathbb{F}_q$. Then $f(\mathbf{U}_{\mathbb{F}_q^n})$ is ε-close to $\mathbf{U}_{\mathbb{F}_q}$, where $\varepsilon = c \cdot d^2/\sqrt{q}$.*

2.3 Gauss's Lemma

Gauss's lemma provides a fundamental link between irreducibility over an integral domain and irreducibility over its field of fractions.

Definition 2.11. *Let R be a unique factorization domain (UFD). Let $f(x) = \sum_{i=0}^{d} c_i x^i \in R[x]$. The content of f, denoted by $c(f)$, is the greatest common divisor of $c_0, \ldots, c_n$; the content is well defined up to invertible elements in R. f is called primitive if $c(f) = 1$.*

Lemma 2.12 (Gauss)**.** *Let R be a UFD, and let $\mathbb{L}$ be its field of fractions. A non-constant polynomial $f \in R[x]$ is irreducible in $R[x]$ if and only if it is both irreducible in $\mathbb{L}[x]$ and primitive in $R[x]$.*

The following corollary follows by Lemma 2.12, together with the fact that the polynomial ring $\mathbb{F}[x_1, \ldots, x_n]$ is a UFD for every n and field $\mathbb{F}$.

Corollary 2.13. *Let $\mathbb{F}$ be a field and let $R = \mathbb{F}[x_1, \ldots, x_n]$, $\mathbb{L} = \mathbb{F}(x_1, \ldots, x_n)$. Then, a multivariate polynomial $f \in \mathbb{F}[\mathbf{x}, y]$ such that $f \notin \mathbb{F}[\mathbf{x}]$ is irreducible in $\mathbb{F}[\mathbf{x}, y]$ if and only if it is irreducible in $\mathbb{F}(\mathbf{x})[y]$ and $c(f) = 1$.*

3 Evaluations that Preserve Irreducibility

In this section, we follow the main outline of [DV22] and [DGV24], and show that their approach extends to the setting where the evaluation points are taken from a larger field.

3.1 Hypothesis (H)

In order to invoke Lecerf's technique [Lec07], we need to assume that our polynomial meets two standard conditions. Lecerf called these conditions Hypothesis (H). For an arbitrary polynomial, [DGV24] showed that one can apply a suitable linear transformation to obtain a polynomial that satisfy Hypothesis (H), provided a certain algebraic condition holds.

Definition 3.1 (Hypothesis (H), [Lec06, Lec07])**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n, y] = \mathbb{F}[\mathbf{x}, y]$ 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. $\operatorname{Res}\left(f(\mathbf{0}, y), \frac{\partial f}{\partial y}(\mathbf{0}, y)\right) \neq 0$.

Definition 3.2 ([DGV24])**.** *For $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{F}^n$, let $s_{\mathbf{a}}$ be the $\mathbb{F}$-linear automorphism of $\mathbb{F}[\mathbf{x}, y]$ that fixes y and sends x_i to $x_i + a_i y$.*

The following lemma provides an algebraic condition on $\mathbf{a}$ under which $s_{\mathbf{a}}(f) - t$ yields a polynomial that satisfies Hypothesis (H).

Lemma 3.3 ([DGV24], Corollary 3.5)**.** *Assume that $f \in \mathbb{F}[\mathbf{x}, y]$ is a polynomial of degree $d \geq 1$ and that* $\mathrm{char}(\mathbb{F})$ *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 $B(\mathbf{a}) \neq 0$, $s_{\mathbf{a}}(f) - t = c \cdot g$ where $c \in \mathbb{F}^\times$ and $g \in \mathbb{F}(t)[\mathbf{x}, y]$ is a degree-d polynomial satisfying Hypothesis (H).*

3.2 Lecerf's Technique

This subsection closely follows [DGV24, Section 4], with minor adaptations to the case of evaluating $f \in \mathbb{K}[\mathbf{x}]$ at points drawn from a larger field $\mathbb{L}/\mathbb{K}$. As discussed in Section 1.2.2 and Section 1.2.3, our goal is to formulate an algebraic condition under which an irreducible polynomial in $\overline{\mathbb{F}_q(t)}[\mathbf{x}]$ remains irreducible after evaluation at elements from a field extension. We refer to such evaluation points as *Bertinian points*.

Definition 3.4. *Let $f \in \mathbb{K}[\mathbf{x}, y]$ be a non-constant polynomial satisfying Hypothesis (H), and let $\mathbb{L}/\mathbb{K}$ be a field extension. We say $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{L}^n$ is a* Bertinian good point *for f if for every irreducible factor $\tilde{f}$ of f over $\mathbb{L}$, the bivariate polynomial $\tilde{f}_{\mathbf{a}}(x, y) = \tilde{f}(a_1 x, \ldots, a_n x, y)$ is also irreducible over $\mathbb{L}$. Otherwise, $\mathbf{a}$ is called a* Bertinian bad point.

We now begin to describe Lecerf's condition for a point to be Bertinian.

Lemma 3.5 ([DGV24], Lemma 2.9, Hensel's Lifting)**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n, y] = \mathbb{K}[\mathbf{x}, y]$ be a nonzero polynomial. Suppose $\bar{\lambda} \in \mathbb{K}$ 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 $\mathbb{K}[\mathbf{x}]$, and*
2. *$\lambda(\mathbf{0}) = \bar{\lambda}$.*

Let $\mathbb{L}/\mathbb{K}$ be an extension of algebraically closed fields. 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 $\mathrm{Res}\left(f(\mathbf{0}, y), \frac{\partial f}{\partial y}(\mathbf{0}, y)\right) \neq 0$, the univariate polynomial $\bar{f}$ factorizes into distinct linear factors

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

where $\bar{\lambda}_i \in \mathbb{K}$ for all $i \in [d]$. By Lemma 3.5, the above factorization of $\bar{f}$ over $\mathbb{K}$ lifts to a factorization of f into distinct linear factors

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

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

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

$$g(\mathbf{z}, x, y) = \prod_{i=1}^{d} (y - \lambda_i(z_1 x, \ldots, z_n x))$$

where each

$$\lambda_i(z_1 x, \ldots, z_n x) \in \mathbb{K}[[\mathbf{z}]][[x]].$$

For $i \in [d]$, let g_i be the factor $y - \lambda_i(z_1 x, \ldots, z_n x)$ 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 $\operatorname{coeff}\left(h, x^j y^k\right) \in A$ the coefficient of $x^j y^k$ in h. We are now ready to define the linear system $D_{\mathbf{z},\sigma}$ used in Lecerf's papers.

Definition 3.6 (The linear system $D_{\mathbf{z},\sigma}$)**.** *Let $\sigma \in \mathbb{N}$. Define $D_{\mathbf{z},\sigma}$ to be the following linear system over $\mathbb{K}(\mathbf{z})$ in the unknowns $\ell_1, \ldots, \ell_d$:*

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

We have the following lemma.

Lemma 3.7 ([DGV24], Lemma 4.2)**.** *For $(j,k) \in \mathbb{N}^2$,*

$$\operatorname{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial x}, x^j y^k\right), \operatorname{coeff}\left(\hat{g}_i \frac{\partial g_i}{\partial y}, x^j y^k\right) \in \mathbb{K}[\mathbf{z}]$$

are polynomials of degree at most $j+1$ and j respectively.

Observe that the construction of $D_{\mathbf{Z},\sigma}$ carries over when we view $f \in \mathbb{L}[\mathbf{x}, y]$ (as opposed to $\mathbb{K}[\mathbf{x}, y]$). Indeed, the elements $\bar{\lambda}_i \in \mathbb{K}$ remain unchanged. By the uniqueness of λ in Lemma 3.5, the lifting to

$$\lambda_i \in \mathbb{K}[\mathbf{x}] \subseteq \mathbb{L}[\mathbf{x}]$$

is therefore unchanged, and consequently $D_{\mathbf{z},\sigma}$ remains the same. Thus, the entire construction can be regarded as in the case of $f \in \mathbb{L}[\mathbf{x}, y]$. Moreover, the coefficients of $D_{\mathbf{z},\sigma}$ lie in $\mathbb{K}[\mathbf{z}]$. Applying [DGV24, Lemma 4.6] to f viewed in $\mathbb{L}[\mathbf{x}, y]$, and noting that the polynomials Q_i are sums of entries of $D_{\mathbf{z},\sigma}$, we obtain the following theorem.

Lemma 3.8. *Let $\mathbb{K}$ be an algebraically closed field with* $\operatorname{char}\mathbb{K} = 0$ *or greater than $d(d-1)$, and let $\mathbb{L}/\mathbb{K}$ be some extension which is algebraically closed. Let $f \in \mathbb{K}[\mathbf{x}, y]$ be an irreducible polynomial of degree $d \geq 1$ satisfying Hypothesis (H). Let $m = 2^{d-1} - 1$. Then, there exist nonzero polynomials $Q_1, \ldots, Q_m \in \mathbb{K}[\mathbf{z}] = \mathbb{K}[z_1, \ldots, z_n]$ of degree at most $2d-1$ such that for every Bertinian bad point $\mathbf{a} \in \mathbb{L}^n$ for f over $\mathbb{L}$, at least one polynomial Q_i vanishes at $\mathbf{a}$.*

4 HSG with Polynomial Evaluation Points

In this section we construct our *polynomial hitting set generators*, and specifically, show that *any* hitting set generator for polynomials over a field extension $\mathbb{F}_{q^k}$ can be simply turned into a polynomial hitting set generator (PHSG) over $\mathbb{F}_q$. Then we use samplers to construct a field extension of $\mathbb{F}_q$ efficiently using only a small amount of random bits.

The main advantage of PHSGs over HSGs is their ability to achieve much higher density, namely $1 - \Theta(d/q^k)$ rather than $1 - \Theta(d/q)$.

We first recall the definition of a PHSG.

Definition 4.1. *Let $\mathbb{F}$ be a finite field. A* polynomial hitting set generator (PHSG) *with density $1-\delta$ for n-variate polynomials of degree at most d over $\mathbb{F}$ with ℓ-variate polynomial evaluation points of degree at most h is a map*

$$H \colon T \to (\mathbb{F}^{\leq h}[w_1, \ldots, w_\ell])^n$$

from a finite set $T \neq \emptyset$ such that for every such nonzero polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most d,

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

The quantity $\log |T|$ *is called the seed length of H.*

Our approach for constructing PHSGs is using a HSG for a field extension

$$\mathbb{E} = \mathbb{F}[w_1, \ldots, w_\ell]/P,$$

where $P \lhd \mathbb{F}[w_1, \ldots, w_\ell]$ is a maximal ideal. Assume that in every equivalence class $g + P \in \mathbb{E}$ there exists an element $g' \in \mathbb{F}[w_1, \ldots, w_\ell]$ of total degree at most h. Let $\{g_1 + P, \ldots, g_k + P\} \subseteq \mathbb{E}$ be a basis of $\mathbb{E}/\mathbb{F}$, and for all $1 \leq i \leq k$ let $g_i' \in \mathbb{F}[w_1, \ldots, w_\ell]$ be an element in the equivalence class $g_i + P$ of total degree at most h. Let

$$\varphi \colon \mathbb{E} \longrightarrow \mathbb{F}^{\leq h}[w_1, \ldots, w_\ell]$$

be the $\mathbb{F}$-linear map defined uniquely by $\varphi(g_i + P) = g'_i$ for all $1 \le i \le k$. Let $\pi\colon \mathbb{F}[w_1, \ldots, w_\ell] \to \mathbb{E}$ be the ring homomorphism $a \mapsto a \mod P$. For all $a \in \mathbb{E}$ we have

$$\pi(\varphi(a)) = a.$$

With this notation in place, we show that an HSG over $\mathbb{E}$ can be regarded as a PHSG. Whenever we apply φ or π to a vector, we apply the map coordinate-wise.

Claim 4.2. *Let $\widehat{H}\colon S \to \mathbb{E}^n$ be a HSG with density $1-\delta$ for n-variate polynomials over $\mathbb{E}$ of degree at most d. Then,*

$$H = \varphi \circ \widehat{H}\colon S \to (\mathbb{F}^{\le h}[w_1, \ldots, w_\ell])^n$$

is a PHSG with density $1-\delta$.

Proof. Let $f \in \mathbb{E}[x_1, \ldots, x_n]$. Let $s \in S$ such that $f(\widehat{H}(s)) \neq 0$. Since π is a ring homomorphism, we obtain

$$0 \neq f(\widehat{H}(s)) = f(\pi(H(s))) = \pi(f(H(s))),$$

and in particular $f(H(s)) \neq 0$. □

Remark 4.3. *We note that $\mathbb{E}$ can equivalently be represented using a single irreducible polynomial $p \in \mathbb{F}[w]$ of degree k, by setting $\mathbb{E} = \mathbb{F}[w]/(p(w))$. Alternatively, we will use several irreducible polynomials of degree 2 in order to optimize the requirement on the field size q. As we will see in Section 5, the field-size requirement depends on the degrees of the polynomial representations of elements of $\mathbb{E}$. Constructing $\mathbb{E}$ via a single degree-k irreducible polynomial would lead to a requirement of $q \ge d^8/\varepsilon^2$, whereas our construction yields the improved bound $q \ge (d \log d)^4/\varepsilon^2$.*

4.1 Sampling Irreducible Polynomials

Toward constructing PHSGs, we require an efficient method for building the extension field $\mathbb{E}$; in particular, we must construct irreducible polynomials over $\mathbb{F}_q$. At present, no unconditional deterministic algorithm is known for constructing an irreducible polynomial of degree d over a field of size q in time $\mathrm{poly}(\log q, d)$. Relevant progress includes the work of Shoup [Sho90], who gave an algorithm running in time $O((\sqrt{p} + \log^2 q)\, d^4)$, where $p = \mathrm{char}(\mathbb{F}_q)$, and the work of Adleman and Lenstra [AL86], who provided a deterministic $\mathrm{poly}(\log q, d)$-time algorithm assuming the Extended Riemann Hypothesis. We overcome this difficulty by using randomness, which we can consider as part of the construction's seed.

Consider sampling a random degree-a monic polynomial, which requires $O(a \log q)$ random bits. It is well known that with probability $\Theta(1/a)$, it will be irreducible. We wish to amplify this probability to $1-\delta$ for an arbitrary $\delta > 0$. To do it randomness-efficiently, we use randomness samplers.

Definition 4.4 (sampler). *Let n,m,t be some positive integers, and let $\varepsilon,\delta>0$. A function $S\colon \{0,1\}^n \to (\{0,1\}^m)^t$ is a (δ,ε) averaging sampler with t samples using n random bits if for every function $f\colon \{0,1\}^m \to [0,1]$ we have*

$$\Pr_{(Z_1,\ldots,Z_t)\sim S(U_n)}\left[\left|\frac{1}{t}\sum_i f(Z_i)-\mathbb{E}[f]\right|\le\varepsilon\right]\ge 1-\delta.$$

We use the recent nearly-optimal sampler of [XZ25] (although some earlier constructions, such as [Zuc97] and [RVW00], would work just as well for our use up to constant factors).

Theorem 4.5 ([XZ25]). *For every constant $\beta>0$, and for every $0<\delta\le\varepsilon<1$ there exists an averaging sampler for the domain $\{0,1\}^m$ that uses $r=m+O(\log(1/\delta))$ random bits and $t=O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\delta}\right)^{1+\beta}$ random samples. This sampler can be constructed in time $\mathrm{poly}(t,r)=\mathrm{poly}(m,\log(1/\delta),1/\varepsilon)$.*

After the sampling step, we will have to test if the sampled polynomials are irreducible. For this, we will use the following well known algorithm.

Algorithm 4.6. *Let $\mathbb{F}_q$ be a finite field. There exists a deterministic algorithm that on input a polynomial $f\in\mathbb{F}_q[w]$ of degree a, decides whether f is irreducible in time $\mathrm{poly}(\log q,a)$.*

Proof (sketch). The algorithm is based on the classical characterization of irreducible polynomials over finite fields. For each integer $i=1,\ldots,a-1$, compute

$$\gcd\bigl(f(x),x^{q^i}-x\bigr).$$

Since every irreducible polynomial of degree i over $\mathbb{F}_q$ divides $x^{q^i}-x$, the polynomial f is irreducible over $\mathbb{F}_q$ if and only if all these greatest common divisors are equal to 1.

Each computation of $x^{q^i} \bmod f$ and the corresponding gcd can be carried out in time polynomial in $\log q$ and a ([vzGG13]), and since the number of iterations is $a-1$, the overall running time is $\mathrm{poly}(\log q,a)$. □

We are ready to construct $\mathbb{E}$ efficiently.

Theorem 4.7. *Let $\mathbb{F}=\mathbb{F}_q$, let $k=2^\ell$ be some positive power of 2 and let $\delta>0$. There exists a probabilistic algorithm that uses $O(k\log q+\log(1/\delta))$ random bits and runs in time $\mathrm{poly}(k\log q,\log(1/\delta))$, such that with probability at least $1-\delta$ outputs ℓ polynomials $p_i\in\mathbb{F}[w_1,\ldots,w_\ell]$, such that $P=(p_1,\ldots,p_\ell)\lhd\mathbb{F}[w_1,\ldots,w_\ell]$ is maximal and $\mathbb{E}=\mathbb{F}[w_1,\ldots,w_\ell]/P$ is a field extension of degree k over $\mathbb{F}$. Moreover, for all $1\le i\le\ell$ we have $p_i=w_i^2-h_i$, where $h_i\in\mathbb{F}_q^{\le i-1}[w_1,\ldots,w_{i-1}]$. Otherwise, the algorithm declares failure.*

Proof. We begin by setting the following notation.

- $R_i = \{f \in \mathbb{F}_q[w_1,\ldots,w_i] : \deg_{w_j} f \leq 1 \text{ for all } 1 \leq j \leq i\}$. Notice that $|R_i| \leq q^{2^i}$.
- $D = \prod_{1\leq i\leq \ell}(R_i \setminus \{0\})$.
- $m = \log|D| = O(k \log q)$.
- $\varepsilon = \frac{1}{2k}$.

For $\mathbf{a} = (\alpha_1,\ldots,\alpha_\ell) \in D$, let $\mathbf{p}(\mathbf{a}) = (p_1,\ldots,p_\ell) = (w_1^2 - \alpha_1,\ldots,w_\ell^2 - \alpha_\ell)$. Let $\mathbb{F}_0 = \mathbb{F}$. For $0 \leq i < \ell$, let

$$\mathbb{F}_{i+1} = \mathbb{F}_i[w_{i+1}]/p_{i+1}(w_{i+1}).$$

Let $\mathbb{E} = \mathbb{F}_\ell$, and let $P = (p_1,\ldots,p_\ell) \lhd \mathbb{F}[w_1,\ldots,w_\ell]$. Notice that $\mathbb{F}_\ell$ is a field if and only if $P \lhd \mathbb{F}_q[w_1,\ldots,w_\ell]$ is maximal, which happens if and only if $p_i \in \mathbb{F}_{i-1}[w_i]$ is irreducible for all $1 \leq i \leq \ell$.

We will use a sampler to find $\mathbf{a} \in D$ such that $\mathbb{E}$ is a field. Let $S\colon \{0,1\}^r \to D^t$ be the (δ,ε) sampler given by Theorem 4.5 with $\beta = 1$, $r = m + O(\log\frac{1}{\delta})$ and $t = O\left(\frac{1}{\varepsilon^2}\log(1/\delta)\right)^2$. Let $f\colon D \to \{0,1\}$ be the indicator function such that $f(\mathbf{a}) = 1$ if and only if $\mathbf{p}(\mathbf{a})$ is maximal. It is well knownt that for every finite field of characteristic different than 2, the probability that an invertible element is a square equals $1/2$ [4]. Hence, if $(\alpha_1,\ldots,\alpha_\ell) \in D$ is chosen uniformly, the probability that all polynomials $p_i \in \mathbb{F}_{i-1}[w_i]$ are irreducible is $1/2^\ell = 1/k$. Hence, $\mathbb{E}[f] = 1/k$.

Using the sampler S we pick $t = O(k^2\log(1/\delta))^2$ values $(\mathbf{a}_1,\ldots,\mathbf{a}_t) \in D^t$ using $r = m + O(\log(1/\delta))$ random bits in time $\mathrm{poly}(k\log(q),\log(1/\delta))$, such that

$$\Pr_{(Z_1,\ldots,Z_t)\sim S(U_r)}\left[\left|\frac{1}{t}\sum_i f(Z_i) - \mathbb{E}[f]\right| \leq \varepsilon\right] \geq 1-\delta.$$

Since $\mathbb{E}[f] = \frac{1}{k}$ and $\varepsilon = \frac{1}{2k}$, this implies in particular that with probability at least $1-\delta$ over $(Z_1,\ldots,Z_t) \sim S(U_r)$, $f(Z_i) = 1$ for at least one $i \in [t]$.

On input a sampler seed $z \in \{0,1\}^r$, we compute $S(z) = (z_1,\ldots,z_t)$, and for each $z_i = (\alpha_1,\ldots,\alpha_\ell) \in D$, we test if $f(z_i) = 1$, as follows:

- For all i from 1 to ℓ, do:
  - Run Algorithm 4.6 to check if $p(w_i) \in \mathbb{F}_{i-1}[w_i]$ is irreducible. If not, return false.
- Return true.

[4] In characteristic 2, one may replace the sampling of polynomials of the form $w_i^2 - \alpha_i$ by sampling arbitrary monic quadratic polynomials $w_i^2 + \alpha_i w_i + \beta_i$. By the prime polynomial theorem, such a polynomial is irreducible with probability at least $1/2 - 1/q$, and an analogous analysis applies. We omit this case, as our results are already taking place only for large characteristic.

If we have found that indeed $f(z_i) = 1$ for some $i \in [t]$, return $(p_1, \ldots, p_\ell)$. Otherwise, the algorithm declares failure.

Recall that the runtime of each irreducibility test is at most poly $(\log |\mathbb{F}_{i-1}|) = \text{poly}(k \log q)$. Thus, the overall runtime of our algorithm is $\text{poly}(k, \log q, \log(1/\delta))$ Finally, recall that if $f(\mathbf{a}) = 1$ then each $p_i \in \mathbb{F}_{i-1}[w_i]$ is an irreducible polynomial of degree 2, and hence $[\mathbb{F}_i : \mathbb{F}_{i-1}] = 2$. Therefore,

$$[\mathbb{E} : \mathbb{F}] = \prod_{1 \le i \le \ell} [\mathbb{F}_i : \mathbb{F}_{i-1}] = 2^\ell = k.$$

This completes the proof. □

4.2 The PHSG Construction

In this subsection we prove the following theorem.

Theorem 4.8. *Let $\mathbb{F} = \mathbb{F}_q$ be a finite field, d a positive integer and $\delta > 0$. Let $k = 2^\ell$ be a power of 2. Then, there exists an absolute constant c such that if $\delta \ge c \cdot \frac{d}{q^k}$, there exists a PHSG $H : T \to (\mathbb{F}^{\le \ell}[w_1, \ldots, w_\ell])^n$ with density $1 - \delta$ for polynomials of degree at most d over $\mathbb{F}$, which can be constructed in time $\text{poly}(d, n, k, \log q)$, and has seed length $O(d \log n + k \log q)$.*

Proof. We start by picking $p_1, \ldots, p_\ell \in \mathbb{F}[w_1, \ldots, w_\ell]$ using the algorithm from Theorem 4.7 with $\delta' = \delta/2$. This algorithm uses $O(k \log q + \log(1/\delta)) = O(k \log q)$ random bits; i.e. there exists an efficiently computable map

$$A : T_1 \to (\mathbb{F}_q^{\le \ell}[w_1, \ldots, w_\ell])^\ell$$

such that with probability at least $1 - \delta'$ over $t \in T_1$ we have that

$$A(t) = (p_1, \ldots, p_\ell) \in \mathbb{F}[w_1, \ldots, w_\ell]$$

satisfies that $P = (p_1, \ldots, p_\ell) \lhd \mathbb{F}[w_1 \ldots, w_\ell]$ is a maximal ideal, and $p_i = w_i^2 - h_i$ where $h_i \in \mathbb{F}_q^{\le 1}[w_1, \ldots, w_{i-1}]$, and $\mathbb{E} = \mathbb{F}[w_1 \ldots, w_\ell]/P$ is a finite field with q^k elements. Moreover, $\log |T_1| = O(k \log q)$.

If the algorithm from Theorem 4.7 did not declare failure, we can construct the field

$$\mathbb{E} = \mathbb{F}[w_1, \ldots, w_\ell]/(p_1, \ldots, p_\ell).$$

Let $\widehat{H} : T_2 \to \mathbb{E}^n$ be the [GX14] HSG given by Theorem 2.5 for n-variate polynomials of degree at most d over $\mathbb{E}$, set with $\delta' = \delta/2$. Recall that $\widehat{H}$ has seed length $O(d \log n + \log(1/\delta)) = O(d \log n + k \log q)$. Let

$$\varphi : \mathbb{E} \to \mathbb{F}^{\le \ell}[w_1, \ldots, w_\ell]$$

be the $\mathbb{F}$-linear map such that $\deg(\varphi(\alpha))_{w_i} \leq 1$ for all $i \in [\ell]$. By Claim 4.2,

$$H' = \varphi \circ \widehat{H} \colon T_2 \to (\mathbb{F}_q^{\leq \ell}[w_1, \ldots, w_\ell])^n$$

is a PHSG for n-variate polynomials of degree at most d with density $1 - \delta'$.

Let

$$H \colon T_1 \times T_2 \to (\mathbb{F}_q[w_1, \ldots, w_\ell])^n$$

be the map such that for $(t_1, t_2) \in T_1 \times T_2$, if $A(t_1)$ succeeds then $H(t_1, \cdot) = H'$, and otherwise, for concreteness, we set $H(t_1, \cdot) = 0$. Now, fix some nonzero $f \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most d. With probability at least $1 - \delta'$ over $t_1 \in T_1$, $A(t_1)$ succeeds, and conditioned on that, with probability at least $1 - \delta'$ over $t_2 \in T_2$,

$$f(H(t_1, t_2)) = f(H'(t_2)) \neq 0.$$

Therefore,

$$\Pr_{(t_1, t_2) \in T_1 \times T_2}[f(H(t_1, t_2)) \neq 0] \geq 1 - 2\delta' = 1 - \delta,$$

which completes the proof. □

5 Our PRG Construction

Let n, d be positive integers, let $q = p^u$ be a prime power with $p \geq d(d-1)+1$ and $q \geq C((d \log d)^4 / \varepsilon^2)$, for some universal constant C to be determined later on, and let $\varepsilon > 0$. We now present the construction of our ε-error PRG

$$G \colon S \longrightarrow \mathbb{F}_q^{n+1}$$

for polynomials $f \in \mathbb{F}_q[\mathbf{x}, y] = \mathbb{F}_q[x_1 \ldots, x_n, y]$ of degree at most d. Let c be an absolute constant, larger than the constants appear in Theorem 2.5, Theorem 4.8 and Lemma 2.10. We will need one PHSG and one HSG for the construction.

1. Let $k = 2^\ell$ be a power of 2 such that $\left\lceil \frac{d}{\log q} \right\rceil + 1 < k \leq 2 \left\lceil \frac{d}{\log q} \right\rceil + 2$. Let

   $$H_1 \colon T_1 \longrightarrow (\mathbb{F}_q^{\leq \ell}[w_1, \ldots, w_\ell])^n$$

   be the PHSG given by Theorem 4.8 for n-variate polynomials over $\mathbb{F}_q$ of degree at most $2d-1$, with $\delta_1 = c \cdot d / q^k \leq c \cdot \frac{d}{2^d q}$, and seed length $O(d \log n + \log q)$.

2. Let

   $$H_2 \colon T_2 \longrightarrow \mathbb{F}_q^n$$

   be the HSG given by Theorem 2.5 for n-variate polynomials over $\mathbb{F}_q$ of degree at most d, with $\delta_2 = c \cdot d / q$ and seed length $O(d \log n + \log q)$.

We are now ready to define G. Let $S = T_1 \times T_2 \times \mathbb{F}_q^\ell \times \mathbb{F}_q \times \mathbb{F}_q$. Define $G\colon S \to \mathbb{F}_q^{n+1}$ by

$$G(r,s,\mathbf{t},u,v) = (H_1(r)_1(\mathbf{t})\cdot v + H_2(s)_1 \cdot u, \ldots, H_1(r)_n(\mathbf{t})\cdot v + H_2(s)_n \cdot u, u), \tag{5.1}$$

where $\mathbf{t} = (t_1,\ldots,t_\ell)$.

Note that the running time of G on input $s \in S$ is $\mathrm{poly}(n,d,\log q)$. Indeed, by Theorems 2.5 and 4.8, the HSG and PHSG can be computed within this time bound, and the additional step of evaluating $\ell = O(\log d)$ variables in n polynomials of degree at most ℓ requires only $\mathrm{poly}(n,\log d,\log q)$ time.

We proceed by showing that if f is indecomposable, then the random restricted polynomial $F = f \circ \mathbf{p}$ remains indecomposable with high probability.

Proposition 5.1. *Let $f \in \mathbb{F}_q[\mathbf{x},y]$ be an indecomposable $(n+1)$-variate polynomial of degree at most d over $\mathbb{F}_q$. Let (r,s) be a random element of $T_1 \times T_2$. Let $H_2(s) = (a_1,\ldots,a_n) = \mathbf{a}$, let $H_1(r) = (b_1(\mathbf{w}),\ldots,b_n(\mathbf{w})) = \mathbf{b}(\mathbf{w})$ for $\mathbf{w} = (w_1,\ldots,w_\ell)$, and finally, denote*

$$F = f(b_1(\mathbf{w})x + a_1 y, \ldots, b_n(\mathbf{w})x + a_n y, y) \in \mathbb{F}_q[x,y,\mathbf{w}].$$

Then,

$$\mathbf{Pr}[F \text{ is indecomposable over } \mathbb{F}_q] \geq 1 - \delta_2 - (2^{d-1}-1)\delta_1.$$

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 $x_i + a_i y$. As f is indecomposable over $\mathbb{F}_q$, so is $s_\mathbf{a}(f)$. By Lemma 2.8, $s_\mathbf{a}(f)$ is also indecomposable over $\overline{\mathbb{F}}_q$. By Lemma 2.9, we further have that $s_\mathbf{a}(f) - t$ is irreducible over $\overline{\mathbb{F}_q(t)}$.

By Lemma 3.3, 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

$$s_\mathbf{a}(f) - t = c \cdot g \tag{5.2}$$

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). Since H_2 is a HSG, the event $B(\mathbf{a}) \neq 0$ happens with probability at least $1-\delta_2$. Condition on this event, so Equation (5.2) holds. As $s_\mathbf{a}(f) - t$ is irreducible over $\overline{\mathbb{F}_q(t)}$, so is g.

Let $\mathbb{K} = \overline{\mathbb{F}_q(t)}$ and let $\mathbb{L} = \overline{\mathbb{F}_q(t,\mathbf{w})}$ be such that $w_1,\ldots,w_\ell$ are new variables. Let $m = 2^{d-1}-1$. As $g \in \mathbb{K}[\mathbf{x},y]$, by Lemma 3.8, there exist nonzero polynomials $Q_1,\ldots,Q_m \in \mathbb{K}[z_1,\ldots,z_n]$ of degree at most $2d-1$ such that the union of the zero loci of these polynomials contains all $\mathbf{b}^* = (b_1^*,\ldots,b_n^*) \in \mathbb{L}^n$ for which $g(b_1^* x,\ldots,b_n^* x, y)$ is reducible over $\mathbb{L}$. H_1 is a PHSG with density $1-\delta_1$ for polynomials of degree at most $2d-1$ over $\overline{\mathbb{F}_q(t)}$. Therefore, for each $i \in [m]$, the probability that $Q_i(\mathbf{b}(\mathbf{w})) = 0$ is at most δ_1.

Condition on the event that $Q_i(\mathbf{b}(\mathbf{w})) \neq 0$ for all $i \in [m]$. Then, $g(b_1(\mathbf{w})x,\dots,b_n(\mathbf{w})x,y)$ is *irreducible* over $\mathbb{L}$. On the other hand, note that

$$\begin{aligned} c \cdot g(b_1(\mathbf{w})x,\dots,b_n(\mathbf{w})x,y) &\overset{(5.2)}{=} (s_{\mathbf{a}}(f))(b_1(\mathbf{w})x,\dots,b_n(\mathbf{w})x,y) - t \\ &= f(b_1(\mathbf{w})x + a_1 y,\dots,b_n(\mathbf{w})x + a_n y, y) - t = F - t, \end{aligned}$$

where the second step uses the definition $s_{\mathbf{a}}(f) = f(x_1 + a_1 y,\dots,x_n + a_n y, y) \in \mathbb{F}_q[\mathbf{x},y]$.

Thus, $F - t$ is irreducible over $\mathbb{L}$, and hence as an element in $\overline{\mathbb{F}_q(t)}(\mathbf{w})[x,y]$. By Lemma 2.12, it is then irreducible as an element in $\overline{\mathbb{F}_q(t)}(\mathbf{w},x)[y]$. Note that the coefficient of y^d in $F - t$ is and element in $\mathbb{F}_q$. Thus, as an element in $\left(\overline{\mathbb{F}_q(t)}[\mathbf{w},x]\right)[y]$, the content of $F - t$ is $c(F) = 1$. Hence, by Corollary 2.13, $F - t$ is irreducible as an element in

$$\overline{\mathbb{F}_q(t)}[\mathbf{w},x][y] = \overline{\mathbb{F}_q(t)}[x,y,\mathbf{w}].$$

By Lemma 2.9, F is indecomposable over $\overline{\mathbb{F}}_q$. So it is indecomposable over $\mathbb{F}_q$.

Overall, note that the indecomposability of F over $\mathbb{F}_q$ relies on the conditions $B(\mathbf{a}) \neq 0$ and $Q_1(\mathbf{b}(\mathbf{w})) \neq 0,\dots,Q_m(\mathbf{b}(\mathbf{w})) \neq 0$. By the union bound, these conditions are simultaneously satisfied with probability at least $1 - \delta_2 - m\delta_1 = 1 - \delta_2 - (2^{d-1} - 1)\delta_1$, which completes the proof. □

Theorem 5.2. *There exists an absolute constant $C > 0$ such that for all $\varepsilon > 0$ and $q \geq C\frac{(d\log d)^4}{\varepsilon^2}$ with $\mathrm{char}(\mathbb{F}_q) \geq d(d-1)+1$, G as defined in Equation (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 $C = 4 \cdot 16^2 \cdot c^2$ for the absolute constant c defined earlier. 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 non constant, since the claim is trivial otherwise.

Our first step is the same as in [DV22] and [DGV24]: f can always be written in the form $f = g \circ h$, 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 Proposition 5.1. Then, by Proposition 5.1, the probability that F is decomposable over $\mathbb{F}_q$ over a random choice of r and s is at most $2^{d-1}\delta_1 + \delta_2 \leq 2c \cdot d/q$. Fix r and s such that F is indecomposable over $\mathbb{F}_q$ and note that by definition,

$$f(G(r,s,\mathbf{t},u,v)) = F(v,u,\mathbf{t}).$$

Applying Lemma 2.10 to F shows that, for such fixed r and s, the distribution of $F(\mathbf{t},u,v)$, i.e., $f(G(r,s,\mathbf{t},u,v))$, over random $t_i,u,v \in \mathbb{F}_q$ is ε'-close to $\mathbf{U}_{\mathbb{F}_q}$, where $\varepsilon' \leq c \cdot (\deg F)^2/\sqrt{q}$. Notice

that

$$\deg F \le d \cdot \ell \le d \cdot 4 \log d.$$

This implies that the statistical distance between $f(G(\mathbf{U}_S))$ and $\mathbf{U}_{\mathbb{F}_q}$ is at most

$$2c \cdot \frac{d}{q} + 16c \cdot \frac{(d \log d)^2}{\sqrt{q}} \le \varepsilon/2.$$

On the other hand, as f is also indecomposable over $\mathbb{F}_q$, applying Lemma 2.10 to f shows that $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ is ε'-close to $\mathbf{U}_{\mathbb{F}_q}$, where $\varepsilon' = c \cdot d^2/\sqrt{q} \le \varepsilon/2$. Therefore, the statistical distance between $f(G(\mathbf{U}_S))$ and $f(\mathbf{U}_{\mathbb{F}_q^{n+1}})$ is at most ε.

The seed length of G is

$$\log |T_1| + \log |T_2| + \ell \log q + 2 \log q = O(d \log n + \log q),$$

which completes the proof. □

6 An Approach Towards Smaller Fields

For $q = p^a$, p prime, denote by $\mathrm{Tr}_{q \to p} \colon \mathbb{F}_q \to \mathbb{F}_p$ the absolute field trace. Assume that we are given a PRG $G \colon S \to \mathbb{F}_q^n$ for n-variate polynomials of total degree at most d over $\mathbb{F}_q$ with error ε. Importantly, assume that q must be such that

$$q \ge \tau(d, \varepsilon)$$

for some threshold function τ. (In our construction, $\tau = \frac{C(d \log d)^4}{\varepsilon^2}$ for some universal constant C.) Also, assume that we have no lower bound on the characteristic p (which is not the case for our construction). A natural attempt to construct a PRG for polynomials over smaller fields is to take traces, namely, $G' \colon S \to \mathbb{F}_p^n$, where

$$G'(s) = (\mathrm{Tr}_{q \to p} \circ G)(s) = \left(\mathrm{Tr}_{q \to p}(G(s)_1), \ldots, \mathrm{Tr}_{q \to p}(G(s)_n)\right).$$

It turns out that this simple approach works, as long as τ mild enough! For concreteness, we fix $\varepsilon > 0$ to some constant, and concentrate on the dependence on d. Moreover, we assume that our "base" PRG G has a seed of length $O(d^{O(1)} \log n + \log q)$, but the proof can easily be adapted to handle other seed lengths.

Proposition 6.1. *Fix some constant $\varepsilon_0 \in (0,1)$. Assume that for any n, d, q such that $q \ge \tau(d, \varepsilon_0) = \tau_0(d) = d^{1-\eta}$ for some $\eta \in (0,1)$, there exists an explicit PRG for n-variate polynomials of total degree at most d over $\mathbb{F}_q$, with error ε_0, and seed length $O(d^c \log n)$, where c is some absolute constant.*

Then, for any n, d, p where $p \le d$ is prime, there exists an explicit PRG for n-variate polynomials of total degree at most d over $\mathbb{F}_p$, with error ε_0, and seed length $O((d/p)^{O(1/\eta)} \log n)$.

Proof. Let $f \in \mathbb{F}_p^{\leq d}[x_1, \ldots, x_n]$. Let q be a power of p soon to be determined. Since $\deg(\mathrm{Tr}_{q\to p}) = q/p$, we have that $h = f \circ \mathrm{Tr}_{q\to p} \colon \mathbb{F}_q^n \to \mathbb{F}_p$, where we apply traces to each field element individually, has degree at most $d' = (dq)/p$ as a polynomial over $\mathbb{F}_q$. Notice that

$$f(G'(s)) = h(G(s)),$$

so G' fools f with error ε_0 whenever $q \geq \tau_0((dq)/p)$. This amounts to

$$q \geq \left(\frac{d}{p}\right)^{\frac{1-\eta}{\eta}}.$$

Invoking G with a suitable q,[5] and degree d', the seed length becomes

$$O(d' \log n + \log q) = O\left(\left(\frac{d}{p}\right)^{O(1/\eta)} \cdot \log n\right),$$

as desired. □

In particular, if such a PRG G exists with any constant η, then we would get a PRG for $\mathbb{F}_2$-polynomials with seed length $d^{O(1)} \cdot \log n$, beating Viola's PRG [Vio09] in the regime where $d = \Omega(\log \log n)$.

References

[ABEK08] Noga Alon, Ido Ben-Eliezer, and Michael Krivelevich. Small sample spaces cannot fool low degree polynomials. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 266–275. Springer, 2008.

[ABN+92] Noga Alon, Jehoshua Bruck, Joseph Naor, Moni Naor, and Ron M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *Information Theory, IEEE Transactions on*, 38(2):509–516, 1992.

[AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

[AL86] Leonard M. Adleman and Hendrik W. Lenstra. Finding irreducible polynomials over finite fields. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 350–355. ACM, 1986.

[5] More precisely, we need q to be a power of p, but the first such q satisfies $q \leq p \cdot (d/p)^{(1-\eta)/\eta}$, and the extra p factor will not change the parameters.

[AMN98] Yossi Azar, Rajeev Motwani, and Joseph Naor. Approximating probability distributions using small sample spaces. *Combinatorica*, 18(2):151–171, 1998.

[BDN09] Arnaud Bodin, Pierre Dèbes, and Salah Najib. Indecomposable polynomials and their spectrum. *Acta Arithmetica*, 139(1):79–100, 2009.

[Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Annual ACM Symposium on Theory of Computing (STOC)*, page 21–30. ACM, 2005.

[BT13] Avraham Ben-Aroya and Amnon Ta-Shma. Constructing small-bias sets from algebraic-geometric codes. *Theory of Computing*, 9(5):253–272, 2013.

[BV10] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM Journal on Computing*, 39(6):2464–2486, 2010.

[CC25] Gil Cohen and Itay Cohen. Wide replacement products meet Gray codes: Toward optimal small-bias sets. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2025.

[CN10] Guillaume Chèze and S. Najib. Indecomposability of polynomials via Jacobian matrix. *Journal of Algebra*, 324(1):1–111, 2010.

[CT13] Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2013.

[DGV24] Ashish Dwivedi, Zeyu Guo, and Ben Lee Volk. Optimal pseudorandom generators for low-degree polynomials over moderately large fields. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 44:1–44:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

[DV22] Harm Derksen and Emanuele Viola. Fooling polynomials using invariant theory. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 399–406. IEEE, 2022.

[GX14] Venkatesan Guruswami and Chaoping Xing. Hitting sets for low-degree polynomials with optimal density. In *Conference on Computational Complexity (CCC)*, pages 161–168. IEEE, 2014.

[Lec06] Grégoire Lecerf. Sharp precision in Hensel lifting for bivariate polynomial factorization. *Mathematics of Computation*, 75(254):921–933, 2006.

[Lec07] Grégoire Lecerf. Improved dense multivariate polynomial factorization algorithms. *Journal of Symbolic Computation*, 42(4):477–494, 2007.

[Lov09] Shachar Lovett. Unconditional pseudorandom generators for low-degree polynomials. *Theory of Computing*, 5(1):69–82, 2009.

[Lu12] Chi-Jen Lu. Hitting set generators for sparse polynomials over any finite fields. In *Computational Complexity Conference (CCC)*, pages 280–286. IEEE, 2012.

[LVW93] Michael Luby, Boban Velickovic, and Avi Wigderson. Deterministic approximate counting of depth-2 circuits. In *Israel Symposium on Theory and Computing Systems (ISTC)*, pages 18–24. IEEE, 1993.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.

[Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Matematicheskie Zametki*, 41(4):598–607, 1987.

[RVW00] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 3–13. IEEE, 2000.

[Sho90] Victor Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation*, 54(189):435–447, 1990.

[Smo93] Roman Smolensky. On representations by low-degree polynomials. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 130–138. IEEE, 1993.

[Ta-17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 238–251. ACM, 2017.

[Vio07] Emanuele Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM Journal on Computing*, 36(5):1387–1403, 2007.

[Vio09] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *computational complexity*, 18(2):209–217, 2009.

[vzGG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, Cambridge, 3rd edition, 2013.

[XZ25] Zhiyang Xun and David Zuckerman. Near-optimal averaging samplers and matrix samplers. In *Computational Complexity Conference (CCC)*, pages 6:1–6:28. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025.

[Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Structures & Algorithms*, 11(4):345–367, 1997.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il