

Hitting Set Generators for low-degree $\mathbb{F}_2$ -polynomials with almost optimal density

Tal Yankovitz*

Abstract

We construct explicit hitting set generators for $\mathbb{F}_2$ -polynomials in n variables of degree at most d with an almost optimal density, and seed length $O(d \log n + d^2)$, which is optimal for $d = O(\log n)$.

Second, we address large-field hitting set generators. The influential state-of-the-art construction by Guruswami and Xing (CCC '14) - which is based on machinery from algebraic geometric codes - had an almost optimal seed length and density. Our first contribution there is a construction with an improved density. Our second large-field contribution is a comparable construction which is based on elementary machinery only.

*University of Toronto. tal.yankovitz@utoronto.ca. The research leading to the results in this paper was done while the author was at UT Austin and supported by NSF Grant CCF-2312573.

1 Introduction

Pseudorandom generators, and hitting set generators, for different classes of functions have long been objects of study in the theory of complexity, pseudorandomness, and derandomization. One such important class is that of low-degree polynomials: n -variate polynomials over a finite field $\mathbb{F}_q$, with total degree at most d . They are the focus of this work, and more specifically the problem of constructing efficient and low-randomness hitting set generators for their class. The main two regimes under which this problem has been studied is the small-field regime, and most importantly the binary $q = 2$ case; and the large-field regime, in which one assumes $q > d$. We start by recalling two basic facts about the number of roots that bounded-degree polynomials may have.

Fact 1.1 (Root bound for the $q = 2$ case). *Let $p \in \mathbb{F}_q[x_1, \dots, x_n]$ be a polynomial of total degree at most d , such that $p \bmod \langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle \not\equiv 0$. Then*

$$\Pr_{\alpha_1 \sim U(\mathbb{F}_q), \dots, \alpha_n \sim U(\mathbb{F}_q)} [p(\alpha_1, \dots, \alpha_n) = 0] \leq 1 - \frac{1}{2^d}.$$

Fact 1.2 (Root bound for the $d < q$ case). *Let $p \in \mathbb{F}_q[x_1, \dots, x_n]$ be a polynomial of total degree at most d . Then*

$$\Pr_{\alpha_1 \sim U(\mathbb{F}_q), \dots, \alpha_n \sim U(\mathbb{F}_q)} [p(\alpha_1, \dots, \alpha_n) = 0] \leq \frac{d}{q}.$$

The challenge of constructing hitting sets generators for low-degree polynomials concerns attempting to find explicit distributions, with a much smaller support than the uniform distribution over $\mathbb{F}_q^n$, which admit a number-of-root bound similar to [Fact 1.1](#) or [Fact 1.2](#).

Definition 1.3 (Hitting set generator). *For a set S , and $h : S \rightarrow \mathbb{F}_q^n$, we say that h is a (d, δ) -HSG (hitting set generator) if for every $p \in \mathbb{F}_q[x_1, \dots, x_n]$ such that $p \bmod \langle x_1^q - x_1, \dots, x_n^q - x_n \rangle \not\equiv 0$, $\Pr_{\alpha \sim U(S)} [p(h(\alpha)) \neq 0] \geq \delta$. We say that h has seed length ℓ to mean that $\log |S| \leq \ell$. The parameter δ is called the density of the hitting set generator.*

We say that a family $\{h_i : S_i \rightarrow \mathbb{F}_{q_i}^{n_i}\}_{i \in I}$ of HSGs, parameterized by infinite set I , is explicit if on input $i \in I$ and given $\lceil \log |S_i| \rceil$ uniformly random bits, one can output an element $y \sim h_i(U(S_i))$, in time $\text{poly}(n, d, \log q)$.

We make several remarks. First, in most cases and prior works, one usually considers values of d much smaller than n . A simple lower bound (see [\[BV10\]](#)) shows in particular that when $d < \sqrt{n}$, if h has seed length ℓ , then $\ell = \Omega(d \log n + \log(\frac{1}{1-\delta}))$.

Second, *pseudorandom generators (PRGs)* for low degree polynomials are a notion tightly related to hitting set generators. In their case, one has a small-support distribution

$\mathcal{D}$ on inputs in $\mathbb{F}_q^n$ such that for a low-degree polynomial p , the distribution $p(\mathcal{D})$ is close in statistical-distance to $p(U(\mathbb{F}_q))$. A pseudorandom generator implies a hitting set generator, where the density δ depends on the statistical-distance guarantee.

Past constructions. In the large-field regime, Klivans and Spielman [KS01] construed explicit HSGs with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$, assuming that the field size satisfies $q > (\frac{nd}{1-\delta})^6$. Bogdanov [Bog05] improved upon them by constructing explicit HSGs with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$ assuming $q > (d \log n / (1 - \delta))^{O(1)}$. Cohen and Ta-Shma [CTS13] constructed HSGs for $q = \Omega(d^6)$ with seed length $O(d^4 \log n)$, which (while not explicit per Definition 1.3) are computable in time $n^{\text{poly}(d)}$. Guruswami and Xing [GX14] achieved explicit HSGs with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$ mildly assuming $q > 2(d/(1-\delta))$. We importantly mention that by the framework devised by [Bog05], in the large-field case, HSGs yield PRGs. Further works importantly achieved PRGs with optimal seed length, assuming that the underlying field is large enough [DV22, DGV24, CDG26].¹

In the small field regime, we will focus on past results for the binary case. The work of Luby, Veličković and Wigderson [LVW93] yielded a HSG with a seed length $\exp(O(d \cdot \sqrt{\log n}))$. Bogdanov [Bog05] described a construction of HSGs with seed length $O(d^2 \log n)$, with density $\delta \geq n^{-2d^2}$, which is computable in time $m^{O(d^2)}$. The independent works of Viola [Vio09] (building on [BV10, Lov08]) on $\mathbb{F}_2$ -PRGs and Dvir and Shpilka [DS08] on noisy interpolating sets gave similar constructions. The HSGs arising from the analysis in [Vio09] have seed length $O(d \log n + \text{poly}(d)2^d)$ (if $\delta > 0$),² and arising from [DS08] are explicit HSGs such that, for some constant $\varepsilon < \frac{1}{4}$, have seed length $d(\log n + O(1))$ and density $\varepsilon^d < \frac{1}{4^d}$. See also specialized constructions, by Lu [Lu12] - for sparse polynomials; and, by Doron, Ta-Shma and Tell [DTST22] - for polynomials that vanish Harley.

Recent independent progress. Very recently, exciting breakthrough progress on the problem of PRGs for low-degree $\mathbb{F}_2$ -polynomials has been obtained: Chattopadhyay, Hatami, Lee, Lovett, Tal and Viola [CHL⁺26] achieved seed length $O(d^2 \log^2 n \log(1/\varepsilon))$ (where ε is the designated statistical-distance); following which, Cohen [Coh26] achieved seed length $O(d^4 \log n)$ or alternatively $\tilde{O}(d^3 \log n)$ for constant ε , and the exact dependence on ε is by a multiplicative dependence on a factor about linear in $\log(1/\varepsilon)$. Our work reported here on HSGs is independent of these recent works. We remark that the

¹Another related problem is (randomized or non-randomized) Polynomial Identity Testing, however a solution for it does not necessarily yield a δ -HSG as there one is allowed more than one sample.

²The analysis in [Vio09] allows with seed length $O(d \log n + \text{poly}(d)2^d)$ reaching density $\delta = (1 - \alpha)2^{-d}$ for α exponentially small in d .

HSG seed length in our construction is polynomially shorter compared to what is induced from these recent PRG results (arising both from the dependence in n or d , and in ε the designated statistical-distance).

1.1 Our results

Our main result is an explicit construction $\mathbb{F}_2$ -hitting set generators with a small seed length (in particular optimal for $d = O(\log n)$) and an almost optimal density compared to the ideal $\frac{1}{2^d}$. More formally,

Theorem 1.4. *For every n and d , and $\varepsilon > 0$ there is a (d, δ) -HSG $h : S \rightarrow \mathbb{F}_2^n$ with $\delta := (1 - \varepsilon)^{\frac{1}{2^d}}$ and seed length $O(d \log n + d^2/\varepsilon)$.³ Furthermore, given any n, d, ε , an output element $y \in \mathbb{F}_2^n$ of their corresponding h can be computed in time $\text{poly}(n, d, 1/\varepsilon)$, given $\lceil \log |S| \rceil$ uniformly random bits, such that $y \sim h(U(S))$.*

Secondly, we address the large-field regime. The construction by Guruswami and Xing [GX14] achieved density $\delta = 1 - \frac{d}{q} - \varepsilon$ where $\varepsilon = O(\frac{d}{q^2})$; when d and q are small or constant, this “error” becomes more substantial. They asked (see [GX14], section VI) if a construction with seed length $O(d \log n + \log \frac{1}{\varepsilon})$ and density $\delta = 1 - \frac{d}{q} - \varepsilon$ is possible for smaller values of ε (notice that δ corresponds to the distance of an induced punctured Reed-Muller code). We show this for $\varepsilon = O((\frac{\log \log n}{\log n})^{1/3})$, in particular yielding a construction in which also for constant d, q , the density δ (or the induced code distance) approaches $1 - \frac{d}{q}$. We prove the following.

Theorem 1.5. *For any q a prime power, and every d , there is a (d, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with $\delta = 1 - \frac{d}{q} - \varepsilon$ and seed length $O(d \log n + \log(\frac{1}{1-\delta}))$, where $\varepsilon = O((\frac{\log \log n}{\log n})^{1/3})$ (for a more general obtained dependance, see Theorem 2.5 in the technical section). Furthermore, given any q, n, d , an output element $y \in \mathbb{F}_q^n$ of their corresponding h_q can be computed in time $\text{poly}(n, d, \log 1/\delta)$, given $\lceil \log |S| \rceil$ uniformly random bits, such that $y \sim h_q(U(S))$.*

Lastly, we also provide a construction of large-field HSGs, which has an optimal (up-to constant factors) seed length, and density tradeoff - as in the [GX14] construction. This construction is of elementary nature, and opposed to the [CTS13] and [GX14] HSGs, does not rely on algebraic geometric codes machinery. See the account in Section 2.4 and the obtained Theorem 2.11 there for the details.

1.2 Proof Overview

In this section we give an overview of our techniques.

³Notice that the notation ε here plays a different part than ε in a PRG’s definition.

1.3 The $\mathbb{F}_2$ -Hitting Set Generator

We turn to describe our hitting set generator for the binary setting. Recall that our goal is to construct a HSG $h: S \rightarrow \mathbb{F}_2^n$ for degree- d n -variate polynomials over $\mathbb{F}_2$ with seed of length $\log |S| = O(d \log n + d^2)$ and density $\delta = \Omega(2^{-d})$. Our proof shows how to convert any HSG over a large field into an HSG over a small field with a small seed loss: in fact we will use the large-field HSG as a tool to reduce the number of variables.

In particular, we will use an instantiation of the large-field HSG with the parameters obtained by [GX14], $h_q: S_q \rightarrow \mathbb{F}_q^{dn}$, for degree- d dn -variate polynomials over $\mathbb{F}_q$, which has seed of length $O(d \log n + \log q)$ and density $\delta_q = 1 - \frac{2d}{q}$, where q is a power of 2. We will choose $q = O(d)$ such that $\delta_q \approx 1/2$.

We start with describing our HSG h . For simplicity, we will describe how to sample from the distribution $\mathcal{D}(h)$ induced by h ,⁴ rather than the mapping from seed to output which can easily be deduced. To sample $\beta \sim \mathcal{D}(h)$:

1. Sample a random seed $s \leftarrow S_q$ for h_q . Set $\alpha = (\alpha_{j,i})_{j \in [d], i \in [n]} = h_q(s)$, where $\alpha_{j,i} \in \mathbb{F}_q$.
2. For every $j \in [d]$, sample a uniformly random function $\phi_j: \mathbb{F}_q \rightarrow \mathbb{F}_2$.
3. Output $\beta = (\beta_i)_{i \in [n]}$ for $\beta_i = \sum_{j \in [d]} \phi_j(\alpha_{j,i})$, where the summation is over $\mathbb{F}_2$.⁵

First, notice that to sample an output of h , we need to sample a seed for h_q together with the d functions ϕ_j . Thus the seed length of h is $\log S = \log S_q + dq = O(d \log n + d^2)$. We are therefore left to bound the density of h .

Towards this, fix a non-zero degree d n -variate polynomial p over $\mathbb{F}_2$. For the sake of this overview, we will assume that the degree of p is exactly d . In the technical part we show that this assumption is without loss of generality, due to a simple argument. Our proof has two steps. First, we identify a property of a vector α , such that given this property, the probability over random choice of $\phi_1, \dots, \phi_d$ that $p(\beta)$ is non-zero is at least 2^{-d} . We will say that a vector α that has this property is *good* for p . Next, we show that h_q must output good α with probability at least $\delta_q \approx 1/2$, which concludes our proof.

⁴Namely, $\mathcal{D}(h)$ is the output distribution of h on a random input sampled from S .

⁵We note that taking the sum of d elements is reminiscent of [Vio09] and [DS08] (though the sum is used differently in those works, and the constructions are otherwise different).

Defining “good” α . Fix a vector $\alpha = (\alpha_{j,i})_{j \in [d], i \in [n]}$, and consider a polynomial $r^\alpha(\mathbf{y})$ over dq variables $\mathbf{y} = (y_{j,\alpha})_{j \in [d], \alpha \in \mathbb{F}_q}$, defined by

$$r^\alpha(\mathbf{y}) := p \left(\sum_{j \in [d]} y_{j,\alpha(j,1)}, \dots, \sum_{j \in [d]} y_{j,\alpha(j,n)} \right).$$

We are interested in identifying a criterion under which, for a random assignment $\mathbf{y} \leftarrow \mathbb{F}_2^{dq}$, $r^\alpha(\mathbf{y}) \neq 0$ with probability at least 2^{-d} . Recalling that

$$p(\beta) = p \left(\sum_{j \in [d]} \phi_j(\alpha_{j,1}), \dots, \sum_{j \in [d]} \phi_j(\alpha_{j,n}) \right),$$

and since a random choice of the functions $\phi_1, \dots, \phi_d$ corresponds exactly to a random assignment for $\mathbf{y}$, we have that under the same criterion, and conditioned on the particular choice of α , $p(\beta)$ has probability at least 2^{-d} to be non-zero – what we aspire towards.

We say that α is good for p if r^α is not the zero polynomial modulo the ideal $I := \langle y_{j,\alpha}^2 - y_{j,\alpha} \mid j \in [d], \alpha \in \mathbb{F}_q \rangle$. To see why this definition is useful, observe that modulo I the polynomial r^α is a degree- d multilinear polynomial over dq variables. Thus, if $r^\alpha \bmod I$ is not the zero polynomial, then it has probability at least 2^{-d} to be non-zero over a random assignment $\mathbf{y} \leftarrow \mathbb{F}_2^{dq}$.

Bounding the probability for “good” α . We next want to show that a random α sampled from h_q is good for p with probability at least δ_q . For this, we will rely on the fact that h_q is a HSG. In particular, we will construct a non-zero degree- d polynomial g over dn variables, such that whenever $g(\alpha)$ is non-zero, α is good for p , which is enough to conclude the proof.

To define g , we first consider a polynomial h over dn variables $\mathbf{z} = (z_{j,i})_{j \in [d], i \in [n]}$, defined by $h(\mathbf{z}) = p(\sum_j z_{j,1}, \dots, \sum_j z_{j,n})$. Writing

$$p(\mathbf{x}) = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \prod_{i \in S} x_i.$$

we have that

$$h(\mathbf{z}) = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \prod_{i \in S} \sum_{j \in [d]} z_{j,i},$$

which, by opening the summation can be written as

$$h(\mathbf{z}) = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} \sum_{f: S \rightarrow [d]} c_S \prod_{i \in S} z_{f(i),i}.$$

We next separate the above sum into two parts: the first, that have $|f(S)|$ exactly d , and the other summands. That is, we write

$$h(\mathbf{z}) = \underbrace{\sum_{\substack{S \subseteq [n] \\ |S|=d}} c_S \sum_{\substack{f: S \rightarrow [d] \\ |f(S)|=d}} \prod_{i \in S} z_{f(i),i}}_{:=h^{d\text{-colors}}(\mathbf{z})} + \underbrace{\sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \sum_{\substack{f: S \rightarrow [d] \\ |f(S)| < d}} \prod_{i \in S} z_{f(i),i}}_{:=h^{< d\text{-colors}}(\mathbf{z})}.$$

Finally, we set $g(\mathbf{z}) = h^{d\text{-colors}}(\mathbf{z})$, and observe that since p is a degree- d polynomial, g is indeed non-zero (see the proof in the technical parts for more details). We thus left to show that α is good for p whenever $g(\alpha) \neq 0$. We proceed with a few observations.

First, observe that for every $\alpha \in \mathbb{F}_q^{nd}$, $r^\alpha(\mathbf{y})$ is equal to the polynomial $h(\mathbf{z})$ under the projection which for every $j \in [d]$ and $i \in [n]$ sets $z_{j,i} = y_{j,\alpha(j,i)}$, which we denote by $h[z_{j,i} \leftarrow y_{j,\alpha(j,i)}]$.⁶ Thus, α is good for p if

$$h^{d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}] + h^{< d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}] \neq 0 \pmod{I}.$$

Next, we show in the proof that the polynomials $h^{d\text{-colors}}$ and $h^{< d\text{-colors}}$ have disjoint sets of monomials (and this is true also modulo I), which implies that they cannot cancel each other. Thus, to show that α is good for p it is enough to show that

$$h^{d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}] \neq 0 \pmod{I}.$$

Since $h^{d\text{-colors}}$ is already a multi-linear polynomial, we can equivalently show that

$$h^{d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}] \neq 0,$$

for which it is enough to observe any assignment $\mathbf{y}$ on which $h^{d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}](\mathbf{y}) \neq 0$. We construct such an assignment $\mathbf{y}$ from α by setting $y_{j,\alpha(j,i)} = \alpha(j,i)$ for every $j \in [d]$ and $i \in [n]$. We observe that for this assignment, it holds that

$$h^{d\text{-colors}}[z_{j,i} \leftarrow y_{j,\alpha(j,i)}](\mathbf{y}) = g(\alpha),$$

since we set $z_{j,i} = y_{j,\alpha(j,i)} = \alpha(j,i)$ and by the choice of $g = h^{d\text{-colors}}$. This in particular implies that α is good whenever $g(\alpha) \neq 0$, as we wanted to show.

1.3.1 On the large-field constructions

We briefly overview our two large-field constructions - the one which improves upon the obtained density, and the one which reproduces the [GX14] result. They are achieved using

⁶Note that $h[z_{j,i} \leftarrow y_{j,\alpha(j,i)}]$ is a polynomial in the variables $\mathbf{y}$.

similar techniques, with some additions. The construction improving the density works in a similar way to the described binary construction, by first increasing the underlying field, and then performing a step that can be viewed as reducing the number of variables – this is done by using a hitting set generator for the larger field; the fact that this field is larger allows the better density. The elementary construction of optimal-density hitting set generators starts with a hitting set generator for large fields (however with non-optimal density, negatively affected by n), due to [Bog05], and observes that it can be made to work on field sizes polynomially smaller compared to the [Bog05] requirement.⁷ After reducing the field size requirement by using this transformation, we further observe that in the case of fields which are large enough, one gets a density that is close to optimal. It remains to handle fields of intermediate size, which we do in a way similar to the other two constructions – by increasing the field size, and then reducing the number of variables. We refer to Section 2.3 and Section 2.4 for the details.

2 Technical part

2.1 Preliminaries

Notation. $\mathbb{N} = \{1, 2, \dots\}$ is the set of natural numbers. All logarithms are taken base 2. $\mathbb{F}_q$ denotes a field of size q . $\mathbb{F}_q[x_1, \dots, x_n]$ is the ring of polynomials in variables $x_1, \dots, x_n$. For a set A , $U(A)$ is the uniform distribution on A . We use bold letters to distinguish between vectors and values, for example $\mathbf{x} = (x_1, \dots, x_n)$.

A polynomial is *multi-linear* if the individual degree of each variable is at most 1.

For a function $h : A \rightarrow B$ we denote by $\mathcal{D}(h)$ its corresponding distribution, obtained by sampling a uniformly random $a \in A$ and outputting $h(a)$, that is $h(U(A))$. Any construction of a hitting set generator $h : S \rightarrow \mathbb{F}_q^n$ will be described through a procedure for sampling an element $\beta \sim h(U(S))$, from which a definition of an appropriate S and h will naturally follow. We will say, for short, that h is *computable in time t* to mean that sampling an element $\beta \sim h(U(S))$ can be done in time t , given access to $\lceil |S| \rceil$ random bits.

2.2 The $\mathbb{F}_2$ -HSG construction

We state our main result.

⁷In this part we make use of a trick whose inspiration we drew after considering the [CDG26] PRG construction.

Theorem 2.1. *Let q be a power of 2. For every n and d there is a (d, δ) -HSG $h : S \rightarrow \mathbb{F}_2^n$ with $\delta := (1 - \frac{2d}{q}) \cdot \frac{1}{2^d}$ and seed length $O(d \log n + dq)$. Furthermore, h can be computed in time $\text{poly}(n, d, q)$.*

Before proving [Theorem 2.1](#) we conclude the following corollary of it.

Corollary 2.2. *For every n and d , and $\varepsilon > 0$ there is a (d, δ) -HSG $h : S \rightarrow \mathbb{F}_2^n$ with $\delta := (1 - \varepsilon) \frac{1}{2^d}$ and seed length $O(d \log n + d^2/\varepsilon)$. Furthermore, h can be computed in time $\text{poly}(n, d, 1/\varepsilon)$.*

Proof. Follows immediately from [Theorem 2.1](#) by letting q be the minimal power of 2 larger than $\frac{2d}{\varepsilon}$. $\square$

Next, we move to prove [Theorem 2.1](#). We will need hitting set generators for large fields.

Theorem 2.3 ([\[GX14\]](#); See also [Theorem 2.11](#)). *For any q a prime power, and every d , there is a (d, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with $\delta = 1 - \frac{2d}{q}$ and seed length $O(d \log n + \log(q/d))$. Furthermore, h_q can be computed in time $\text{poly}(n, d, \log q)$.*

Prior to describing in full the proof for [Theorem 2.1](#), we start by presenting the specific construction of our $\mathbb{F}_2$ -HSG, h .

Construction of h , a (d, δ) -HSG for n -variate polynomials over $\mathbb{F}_2$; assuming h_q is a (d, δ_q) -HSG for dn -variate polynomials over $\mathbb{F}_q$.

To sample $\beta \sim \mathcal{D}(h)$:

1. Sample $\alpha = (\alpha_{j,i})_{j \in [d], i \in [n]} \sim \mathcal{D}(h_q)$, where $\alpha_{j,i} \in \mathbb{F}_q$.
2. For every $j \in [d]$, sample a uniformly random function

$$\phi_j : \mathbb{F}_q \rightarrow \mathbb{F}_2.$$

3. Output $\beta = (\beta_i)_{i \in [n]}$ such that

$$\beta_i = \sum_{j \in [d]} \phi_j(\alpha_{j,i}),$$

where the summation is over $\mathbb{F}_2$.

Figure 1: The construction of the HSG h . We describe h in the language of how to sample from its corresponding distribution $\mathcal{D}(h)$ (without explicitly referring to the seed set S , such an equivalent description naturally follows). The seed length is what is required by h_q plus additional dq bits.

We now proceed with the proof for the main theorem.

Proof for Theorem 2.1. Let h_q be a (d, δ_q) -HSG for dn -variate polynomials over $\mathbb{F}_q$, guaranteed by Theorem 2.3, with $\delta_q := 1 - \frac{2d}{q}$. We take h , our supposed (d, δ) -HSG, to be the function defined in Figure 1. Observe that the seed length for it is bounded by the seed length for h_q together with the seed needed in order to sample $\phi_1, \dots, \phi_q$, that is, at most

$$O(d \log(dn) + \log(q/d)) + dq = O(d \log n + dq),$$

where we used the fact that without loss of generality, $d < n$.

We turn to show the correctness of h . We fix a non-zero polynomial

$$p = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \prod_{i \in S} x_i.$$

We claim that

$$\Pr_{\beta \sim \mathcal{D}(h)} [p(\beta) \neq 0] \geq \delta, \tag{2.1}$$

where recall, $\delta := (1 - \frac{2d}{q}) \cdot \frac{1}{2^d}$. We start with assuming that the degree of p is exactly d . Later, we will show how to handle lower degree polynomials. To show the above, we

start by identifying a condition on the α chosen by the sampling process of h , Figure 1, conditioned upon which, the probability that $p(\beta) \neq 0$ is at least $1/2^d$.

Assume $\alpha = (\alpha(j, i))_{j \in [d], i \in [n]} \in \mathbb{F}_q^{nd}$ has been sampled by h in Step 1. In order to inspect the output of p as a function of the sampled $\phi_1, \dots, \phi_d$, we define for every $j \in [d]$ and $\alpha \in \mathbb{F}_q$ the random variable $y_{j, \alpha} := \phi_j(\alpha)$, and also denote the tuple $\mathbf{y} := (y_{j, \alpha})_{j \in [d], \alpha \in \mathbb{F}_q}$. It is immediate from the definition of h that in this case, the output value $p(\beta)$ is given by the polynomial $r^\alpha(\mathbf{y})$,

$$r^\alpha(\mathbf{y}) := p \left(\sum_{j \in d} y_{j, \alpha(j, 1)}, \dots, \sum_{j \in d} y_{j, \alpha(j, n)} \right).$$

Since $(y_{j, \alpha})_{j \in [d], \alpha \in \mathbb{F}_q}$ take uniformly random independent values in $\mathbb{F}_2$, then by Fact 1.1, if the following property of p, α

- *Condition $C(\alpha)$* : r^α is not the zero polynomial modulo the ideal $I := \langle y_{j, \alpha}^2 - y_{j, \alpha} \mid j \in [d], \alpha \in \mathbb{F}_q \rangle$

holds, then

$$\Pr_{\mathbf{y}}[r^\alpha(\mathbf{y}) \neq 0] \geq 1/2^d,$$

and $p(\beta) \neq 0$ with the same probability.

Thus, our goal is to bound from below the probability that h samples an α for which Condition $C(\alpha)$ holds. Towards that, we first inspect the following polynomial $h(\mathbf{z})$, this time over variables $\mathbf{z} = (z_{j, i})_{j \in [d], i \in [n]}$,

$$h(\mathbf{z}) = p \left(\sum_{j \in d} z_{j, 1}, \dots, \sum_{j \in d} z_{j, n} \right) \quad (2.2)$$

It is immediate that for every $\alpha \in \mathbb{F}_q^{nd}$, $r^\alpha(\mathbf{y})$ (which is a polynomial in dq variables) is equal to the polynomial $h(\mathbf{z})$ (which is a polynomial in dn variables) under the projection which sets $\forall j \in [d], i \in [n] : z_{j, i} = y_{j, \alpha(j, i)}$, which we denote by $h(z_{j, i} \leftarrow y_{j, \alpha(j, i)})$. Furthermore, we have that, by the definition of h and p ,

$$\begin{aligned} h(\mathbf{z}) &= \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \prod_{i \in S} \sum_{j \in [d]} z_{j, i} \\ &= \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \sum_{f: S \rightarrow [d]} \prod_{i \in S} z_{f(i), i} \\ &= \underbrace{\sum_{\substack{S \subseteq [n] \\ |S| = d}} c_S \sum_{\substack{f: S \rightarrow [d] \\ |f(S)| = d}} \prod_{i \in S} z_{f(i), i}}_{:= h^{d\text{-colors}}(\mathbf{z})} + \underbrace{\sum_{\substack{S \subseteq [n] \\ |S| \leq d}} c_S \sum_{\substack{f: S \rightarrow [d] \\ |f(S)| < d}} \prod_{i \in S} z_{f(i), i}}_{:= h^{< d\text{-colors}}(\mathbf{z})}. \end{aligned}$$

Recall that we wish to show that $r^\alpha(\mathbf{y}) \bmod I \neq 0$ with a high probability, or equivalently

$$h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)}) + h^{< d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)}) \neq 0 \pmod{I}$$

We start with the observation that the set of monomials of $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is disjoint of the set of monomials of $h^{< d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$, which means the first cannot be canceled out by the latter. Indeed, by the definition of $h^{d\text{-colors}}$, each monomial in the support of $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is indexed by a set $S = \{i_1, \dots, i_d\} \subseteq [n]$, and a function $f : S \rightarrow [d]$ where we assume without loss of generality that $f(i_1) = 1, \dots, f(i_d) = d$. In this case, the monomial is of the form $y_{1,\alpha(1,i_1)} \cdots y_{d,\alpha(d,i_d)}$, and has degree exactly d . In contrast, every $(S' = \{i'_1, \dots, i'_{|S'|}\}, f')$ -monomial in the support of $h^{< d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is of the form $y_{f'(i'_1),*} \cdots y_{f'(i'_{|S'|}),*}$, where $\{f'(i'_1), \dots, f'(i'_{|S'|})\} \neq [d]$, since $|f(S')| < d$. Thus every monomial of $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is not supported on the monomials of $h^{< d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$, which means the first cannot be canceled out by the latter.

Moreover, observe that this disjoint-monomial-support fact is preserved when taking modulo I , and that $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)}) \bmod I = h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$, since every variable in each monomial of $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ has degree at most 1 (that is, it is multi-linear).

From the above, we conclude that it suffices to show that with a high probability over the choice of α , $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is not the zero polynomial, as in which case it is implied that $r^\alpha(\mathbf{y}) \bmod I$ is not the zero polynomial.

Towards showing that, we first observe that $h^{d\text{-colors}}(\mathbf{z})$ (as is, prior to being projected $z_{j,i} \leftarrow y_{j,\alpha(j,i)}$) is not the zero polynomial. Indeed, by the assumption that $\deg p = d$, there is some monomial $S = \{i_1, \dots, i_d\}$ where $i_1 < \dots < i_d$, such that $c_S \neq 0$, and so the monomial $z_{1,i_1} \cdots z_{d,i_d}$ which corresponds to $f : S \rightarrow [d]$ such that $f(i_j) = j$, appears in the definition of $h^{d\text{-colors}}(\mathbf{z})$ with coefficient $c_S \neq 0$. Moreover, it cannot be canceled out by monomials arising from $S' \neq S$, or monomials with the same S but a different mapping $f' \neq f$, since they are of different form. Therefore indeed $h^{d\text{-colors}}(\mathbf{z}) \neq 0$. Secondly, observe that if there is some evaluation (to $\mathbf{y}$) on which $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is non-zero, then it follows that $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is not the zero polynomial (even if the evaluation point comes from an extension field). Thus, it is enough to observe such an evaluation. For this, given some α , consider a point in $\mathbb{F}_q^{nq}$ such that

$$\forall j \in [d], i \in [n] : y_{j,\alpha(j,i)} \leftarrow \alpha(j,i).$$

Evaluated on this point, $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$ is equal to $h^{d\text{-colors}}(\alpha)$. However, α chosen by h is sampled by an $\mathbb{F}_q$ -HSG, while $h^{d\text{-colors}}(\mathbf{z})$ is a degree- d non-zero polynomial. So, with probability at least $(1 - \frac{2d}{q})$, $h^{d\text{-colors}}(\alpha) \neq 0$. We conclude that $h^{d\text{-colors}}(z_{j,i} \leftarrow y_{j,\alpha(j,i)})$

is not the zero polynomial with probability at least the same, in which case Condition $C(\alpha)$ holds. This establishes the claimed bound on the density, Equation (2.1).

It remains only to address the efficiency of h : besides sampling an element of h_q , which is done in time $\text{poly}(n, d, \log(q))$, sampling $\phi_1, \dots, \phi_d$ can be done in time $\text{poly}(dq)$, thus h has the claimed time complexity. Thus the theorem follows.

Handling lower degrees. Finally, we explain how the proof can be generalized to handle the case that the degree of p is $d' < d$. In this case, let $m = x_{i_1} \cdots x_{i_{d'}}$ be any maximal-degree monomial in the support of p . Let $x_{i_{d'+1}}, \dots, x_{i_d}$ be any $d - d'$ variables that do not appear in m . Consider the polynomial

$$\hat{p}(x) := x_{i_{d'+1}} \cdots x_{i_d} \cdot p(x).$$

It is of degree exactly d , and therefore by the proof, is non-zero on the output of h with probability at least δ . Moreover, on any evaluation on which p is zero, so is $\hat{p}$. The generalized claim follows. $\square$

2.3 A large-field construction with improved density

In this section we improve the density guarantee for large-field HSGs. We first state the HSGs achieved by [GX14] with a more exact density expression.

Theorem 2.4 ([GX14]). *For any q a prime power, and every d , there is a (d, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with $\delta = 1 - \frac{d}{q} - \varepsilon$ and seed length $O(d \log n + \log(q/d))$, where $\varepsilon = O(\frac{d}{q^2})$. Furthermore, h_q can be computed in time $\text{poly}(n, d, \log q)$.*

In this part we prove the following.

Theorem 2.5. *For any q a prime power, and every d , there is a (d, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with $\delta = 1 - \frac{d}{q} - \varepsilon$ and seed length $O(d \log n + \log(q/d))$, where $\varepsilon = O(\min(\sqrt{\frac{d \log q}{\log n}}, \frac{d}{q^2}))$ (in particular, if $d < q$, $\varepsilon = O((\frac{\log \log n}{\log n})^{1/3})$). Furthermore, h_q can be computed in time $\text{poly}(n, d, \log q)$.*

Construction of h_q , a (d, δ) -HSG for n -variate polynomials over $\mathbb{F}_q$.

- Let q' be the minimal power of q , such that $q' \geq \sqrt{\frac{d \log n}{\log q}}$.
- Let $h_{q'}$ be a $(d, \delta_{q'})$ -HSG for n -variate polynomials over $\mathbb{F}_{q'}$.

To sample $\beta \sim \mathcal{D}(h_q)$:

1. Sample $\alpha \sim \mathcal{D}(h_{q'})$, where $\alpha_i \in \mathbb{F}_{q'}$ for $i \in [n]$.
2. Sample a uniformly random function

$$\phi : \mathbb{F}_{q'} \rightarrow \mathbb{F}_q.$$

3. Output β such that

$$\beta_i = \phi(\alpha_i).$$

Figure 2: The construction of the HSG h_q . We describe h_q in terms of how to sample from its corresponding distribution $\mathcal{D}(h_q)$.

Proof. First, if $q > \sqrt{\frac{d \log n}{\log q}}$, then we can directly use [Theorem 2.4](#), and get $\varepsilon = O(\frac{\log q}{\log n})$.

We thus proceed under the assumption that $q \leq \sqrt{\frac{d \log n}{\log q}}$. We analyze the construction of h_q given in [Figure 2](#), where we take $h_{q'}$ to be a hitting set generator from [Theorem 2.4](#) for n variables and field $\mathbb{F}_{q'}$, with density $\delta_{q'} = 1 - O(\frac{d}{q'})$. Observe that q' as defined in h_q must satisfy

$$q' \leq q \cdot \sqrt{\frac{d \log n}{\log q}} \leq \frac{d \log n}{\log q}. \quad (2.3)$$

The seed length needed for the described construction is that of $h_{q'}$ plus additional $q' \log q$ bits, that is

$$O(d \log(n) + \log(q'/d)) + q' \lceil \log q \rceil = O(d \log(n) + \log(\log n / \log q)) + O(d \log n) = O(d \log(n))$$

many bits, by [Equation \(2.3\)](#).

We analyze the success probability. We fix a polynomial $p \in \mathbb{F}_q[x_1, \dots, x_n]$ with total degree at most d .

Assume $\alpha = (\alpha(i))_{i \in [n]} \in \mathbb{F}_{q'}^n$ has been sampled by h_q in [Step 1](#). We define for every $\gamma \in \mathbb{F}_{q'}$ the variable y_γ , and also denote the tuple $\mathbf{y} := (y_\gamma)_{\gamma \in \mathbb{F}_{q'}}$. We consider the polynomial

$$r^\alpha(\mathbf{y}) := p(y_{\alpha_1}, \dots, y_{\alpha_n}).$$

We first observe that the probability that $r^\alpha(\mathbf{y})$ is not the zero polynomial is high. Indeed, on the event that $p(\alpha) \neq 0$, it must be that $r^\alpha(\mathbf{y})$ is not the zero polynomial, since if let $r^\alpha(y_\gamma \leftarrow \gamma)$ denote the value of r^α under the evaluation which sets every y_γ to γ , then

$$r^\alpha(y_\gamma \leftarrow \gamma) = p(\alpha).$$

We thus see that as $p(\alpha)$ is non-zero with probability at least $\delta_{q'}$, r^α is not the zero-polynomial with at least the same probability.

If $r^\alpha(\mathbf{y})$ is not the zero-polynomial, then it is an immediate implication of [Fact 1.2](#) that over a uniformly random $\phi : \mathbb{F}_{q'} \rightarrow \mathbb{F}_q$, the probability that $r^\alpha((\phi(\gamma))_{\gamma \in \mathbb{F}_{q'}}) = p(\phi(\alpha_1), \dots, \phi(\alpha_n))$ is non-zero at least $1 - \frac{d}{q}$. We conclude that the success probability of h_q is bounded below by

$$\begin{aligned} \delta_{q'} \cdot \left(1 - \frac{d}{q}\right) &= \left(1 - O\left(\frac{d}{q'}\right)\right) \cdot \left(1 - \frac{d}{q}\right) \\ &= 1 - \frac{d}{q} - O\left(\frac{d}{q'}\right) \\ &= 1 - \frac{d}{q} - O\left(\sqrt{\frac{d \log q}{\log n}}\right), \end{aligned}$$

by the choice of q' .

It remains to address the efficiency of $h_{q'}$: the step it performs consist of finding q' , sampling from $h_{q'}$, finding and applying ϕ . Sampling from $h_{q'}$ is done in time $\text{poly}(n, d, \log q') = \text{poly}(n, d)$ and rest of the steps can evidently be performed in time $\text{poly}(n, d, q') = \text{poly}(n, d)$, by [Equation \(2.3\)](#). Hence h_q has the claimed efficiency. In the large q case, the efficiency bound $\text{poly}(n, d, \log q)$ from [Theorem 2.4](#) applies. The theorem follows. $\square$

2.4 Elementary large-field HSGs

As mentioned, [\[GX14\]](#) constructed large-field HSGs with close to optimal seed and density, based on algebraic codes machinery. In this section we do the same, but based on an elementary construction. The conclusion of this part will be given in [Theorem 2.11](#) below, which will correspond to [Theorem 2.3](#) quoted above (and also in the improvement following it, [Theorem 2.12](#), which will correspond to [Theorem 2.4](#) quoted above), but prior to describing the construction, we will require a few prerequisites.

We first recall the following result of [\[Bog05\]](#) (in itself proved in an elementary way).

Theorem 2.6 ([Bog05]). *For any n, d , and $q > (d \log n / (1 - \delta))^{32}$ a prime power, there is an efficient (d, q, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$.⁸*

Using a trick, we will deduce from Theorem 2.6 an improvement of it, reducing the required field size.⁹ We do so in Lemma 2.8, but prior to describing it, we will require the following fact.

Fact 2.7 (See also [Sho90, CDG26]). *Let $\mathbb{F}_q$ be a field. There is a randomized algorithm which finds a degree 2 irreducible polynomial $r(y) \in \mathbb{F}_q[y]$ with probability at least $1 - \frac{1}{q^2}$, or outputs “fail”, and runs in time $\text{poly}(\log q)$ and by using $O(\log q)$ random bits.¹⁰*

Lemma 2.8. *Assume that for some functions f, s , for any n, d, δ and $q > f(n, d, \delta)$ a prime power, there is an efficient (d, δ) -HSG for n -variate polynomials over $\mathbb{F}_q$, with seed length $s(q, n, d)$. Then, for any $n, d, \delta > 0$ and $q > \sqrt{f(n, d, \delta)}$ a prime power, there is an efficient $(d, \delta - \frac{d}{q})$ -HSG for n -variate polynomials over $\mathbb{F}_q$, with seed length $s(q^2, n, d) + O(\log q)$.*

Proof. Let

$$q > \sqrt{f(n, d, \delta)}$$

be a prime power. We describe how (sampling from) a $(d, \delta - \frac{d}{q})$ -HSG for n -variate polynomials over $\mathbb{F}_q$ proceeds.

Set

$$q' := q^2 > f(n, d, \delta) \tag{2.4}$$

Use Fact 2.7 to find an irreducible polynomial $r(y) \in \mathbb{F}_q[y]$ of degree 2, or if the algorithm for finding one fails, halt and output $(0, \dots, 0) \in \mathbb{F}_q^n$. Use the polynomial $r(y)$ to represent $\mathbb{F}_{q'} \cong \mathbb{F}_q[y]/(r(y))$; and so we represent every element $\gamma \in \mathbb{F}_{q'}$ as $\gamma_1 y + \gamma_2$ for $\gamma_1, \gamma_2 \in \mathbb{F}_q$. Use the HSGs for n -variate degree at most d polynomials over $\mathbb{F}_{q'}$ of the hypothesis to sample $\alpha_1, \dots, \alpha_n \in \mathbb{F}_{q'}$; we can do so by Equation (2.4). For every $i \in [n]$ let $\alpha_i(y)$

⁸[Bog05] does not specify the explicit constant in the exponent of the field expression, but it can be seen to be at most 18, which we rounded to the next power of 2. We note that [Bog05] also has a theorem where this exponent is just 2, but there the seed dependence on d is non-optimal.

⁹Our inspiration for this trick came after we read the [CDG26] PRG construction.

¹⁰Proof sketch: an easy counting shows that a random degree 2 polynomial is irreducible with probability at least $1/4$; moreover, one can efficiently and deterministically check if such a sampled polynomial $r(y)$ is irreducible, by checking its GCD with $(y^q - y) \bmod r(y)$. Repeating this with $O(\log q)$ independent samples gives the desired success probability, but costs $O(\log^2 q)$ random bits. To achieve the same with $O(\log q)$ random bits, one can instead iterate over $\text{poly}(\log q)$ dependent samples - sampled by invoking a Randomness Sampler - for example the one by [Zuc97], requiring only $O(\log q)$ random bits to achieve the desired success probability. There are also other, more direct, ways to achieve the same, see [Sho90]. See also a discussion in [CDG26].

denote the polynomial representing α_i . The defined HSG's sampling process concludes by sampling $\beta \sim U(\mathbb{F}_q)$, and outputting $(\alpha_1(\beta), \dots, \alpha_n(\beta)) \in \mathbb{F}_q^n$.

For the correctness, assume that an irreducible $r(y)$ was successfully obtained, and let $p \in \mathbb{F}_q[x_1, \dots, x_n]$ be of degree at most d . With probability at least δ , $p(\alpha_1(y), \dots, \alpha_n(y))$ is not the zero polynomial: otherwise $\mathbb{F}_{q'} \ni p(\alpha_1, \dots, \alpha_n) = 0$. In which case it also follows, as $\alpha_1(y), \dots, \alpha_n(y)$ are linear univariate polynomials, that $p(\alpha_1(y), \dots, \alpha_n(y))$ is a non-zero univariate polynomial of degree at most d . Thus when evaluated on a uniformly random $\beta \in \mathbb{F}_q$ it is non-zero with probability at least $1 - \frac{d}{q}$. The success probability is thus at least $\delta - \frac{d}{q} - \frac{1}{q^2}$, also accounting for the probability that $r(y)$ was not successfully found. The seed length is that required by the hypothesis's HSG, plus the number of random bits required in order to sample $r(y)$: $O(\log q)$ bits, and additional $\log q$ bits are used in order to sample $\beta \sim U(\mathbb{F}_q)$, and so $s(q^2, n, d) + O(\log q)$ in total. The lemma follows. $\square$

We conclude the following theorem, whose proof is by starting from [Theorem 2.6](#), and applying [Lemma 2.8](#) 8 times.

Theorem 2.9. *For any n, d , and $q > (d \log n / (1 - \delta))^{1/8}$ a prime power, there is an efficient $(d, q, \delta - \frac{8d}{q} - \frac{8}{q^2})$ -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$.*

We also obtain the following corollary.

Corollary 2.10. *For any n, d , if $q > (\log n)^{1/7}$ is a prime power, there is an efficient $(d, q, \delta := 1 - \frac{9d}{q} - \frac{8}{q^2})$ -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$.*

Proof. Define $\delta' := 1 - \frac{d}{q} > \delta$ (without loss of generality, $d < q$). Notice that

$$\left(\frac{d \log n}{1 - \delta'} \right)^{1/8} = (q \log n)^{1/8},$$

therefore $q > (d \log n / (1 - \delta'))^{1/8} \iff q > (\log n)^{1/7}$. Since we assume the right hand side, we can apply [Theorem 2.9](#) with respect to δ' . We therefore get that there is an efficient $(d, q, \delta' - \frac{8d}{q} - \frac{8}{q^2})$ -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta'}))$. It remains to notice that $\delta' - \frac{8d}{q} - \frac{8}{q^2} = 1 - \frac{9d}{q} - \frac{7}{q^2} = \delta$, and that $\log(\frac{1}{1-\delta'}) = \log(\frac{1}{1-\delta}) + O(1)$, and the corollary follows. $\square$

We turn to state and prove the main theorem of this part. Following it, we will also discuss an improvement.

Theorem 2.11. *For any n, d, δ and $q > 17d/(1 - \delta)$ there is an efficient (d, q, δ) -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$.*

Construction of h_q , a (d, δ) -HSG for n -variate polynomials over $\mathbb{F}_q$.

- Let q' be the minimal power of q such that $q' > (\log n)^{1/7}$.
- Let $h_{q'}$ be a $(d, 1 - O(\frac{d}{q'}))$ -HSG for n -variate polynomials over $\mathbb{F}_{q'}$.

To sample $\beta \sim \mathcal{D}(h_q)$:

1. Sample $\alpha \sim \mathcal{D}(h_{q'})$, where $\alpha_i \in \mathbb{F}_{q'}$ for $i \in [n]$.
2. Sample a uniformly random function

$$\phi : \mathbb{F}_{q'} \rightarrow \mathbb{F}_q.$$

3. Output β such that

$$\beta_i = \phi(\alpha_i).$$

Figure 3: The construction of the HSG h_q . We describe h_q in the language of how to sample from its corresponding distribution $\mathcal{D}(h_q)$.

Proof. First, if already $q > (\log n)^{1/7}$, the theorem immediately follows by [Corollary 2.10](#). We therefore proceed under the assumption that

$$q \leq (\log n)^{1/7}. \quad (2.5)$$

We analyze the construction of h_q given in [Figure 3](#), where we take $h_{q'}$ to be a hitting set generator from [Corollary 2.10](#) for n variables over the larger field, of size q' .

Observe that q' found by h_q must satisfy

$$q' \leq q \cdot (\log n)^{1/7} \leq (\log n)^{2/7}.$$

The seed length needed for the described construction is that of $h_{q'}$ plus additional $q' \log q$ bits, that is

$$O(d \log(n) + \log(q'/d)) + q' \lceil \log q \rceil$$

many bits. Observe that $q' \lceil \log q \rceil = O((\log n)^{2/7} \log \log n)$, and that $\log(q'/d) < \log(q') = O(\log \log n)$. Hence, the total seed length (in our case, assuming [Equation \(2.5\)](#)) is $O(d \log n)$.

We analyze the success probability. Fix a polynomial $p \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree at most d . Assume $\alpha = (\alpha(i))_{i \in [n]} \in \mathbb{F}_{q'}^n$ has been sampled by h_q in [Step 1](#). We define for every $\gamma \in \mathbb{F}_{q'}$ the variable y_γ , and also denote the tuple $\mathbf{y} := (y_\gamma)_{\gamma \in \mathbb{F}_{q'}}$. We consider the polynomial

$$r^\alpha(\mathbf{y}) := p(y_{\alpha_1}, \dots, y_{\alpha_n}).$$

We first observe that the probability that $r^\alpha(\mathbf{y})$ is not the zero polynomial is high. Indeed, on the event that $p(\alpha) \neq 0$, it must be that $r^\alpha(\mathbf{y})$ is not the zero polynomial, since if let $r^\alpha(y_\gamma \leftarrow \gamma)$ denote the value of r^α under the evaluation which sets ever each y_γ to γ , then

$$r^\alpha(y_\gamma \leftarrow \gamma) = p(\alpha).$$

We thus see that as $p(\alpha)$ is non-zero with probability at least $1 - \frac{9d}{q'} - \frac{8}{q^2}$, r^α is not the zero-polynomial with at least the same probability.

Finally if r^α is not the zero-polynomial, then it is an immediate implication of [Fact 1.2](#) that over a uniformly random $\phi : \mathbb{F}_{q'} \rightarrow \mathbb{F}_q$, the probability that $r^\alpha((\phi(\gamma))_{\gamma \in \mathbb{F}_{q'}}) = p(\phi(\alpha_1), \dots, \phi(\alpha_n))$ is non-zero at least $1 - \frac{d}{q}$. We conclude that the success probability of h_q is bounded below by

$$\left(1 - \frac{9d}{q'} - \frac{8}{q^2}\right) \cdot \left(1 - \frac{d}{q}\right) \geq 1 - \frac{17d}{q} > \delta,$$

as required. $\square$

In [Theorem 2.11](#) the achieved density tradeoff we get is $\delta \geq 1 - \frac{17d}{q}$ with an optimal seed length, while in [\[GX14\]](#) the guarantee $\delta \geq 1 - \frac{2d}{q}$, and more precisely, $\delta \geq 1 - \frac{d}{q} - O(\frac{d}{q^2})$. However, we can also obtain the same.

Theorem 2.12. *For any n, d , and $q > d$ there is an efficient $(d, q, \delta = 1 - \frac{d}{q} - O(\frac{d}{q^2}))$ -HSG $h_q : S \rightarrow \mathbb{F}_q^n$ with seed length $O(d \log n + \log(\frac{1}{1-\delta}))$.*

Proof. Apply [Lemma 2.8](#) with the function

$$f(n, d, \delta') := \frac{17d}{1 - \delta'},$$

and with the HSGs of [Theorem 2.11](#). We get for any n, d, δ' and $q > \sqrt{f(n, d, \delta')}$ there is an efficient $(d, \delta' - \frac{d}{q})$ -HSG for n -variate polynomials over $\mathbb{F}_q$, with seed length $O(d \log n + \log q) = O(d \log n + \log q - \log d) = O(d \log n + \log(\frac{1}{1-\delta'}))$. Notice that

$$q > \sqrt{f(n, d, \delta')} = \sqrt{\frac{17d}{1 - \delta'}}$$

holds if and only if

$$\delta' < 1 - \frac{17d}{q^2}.$$

Hence, if we set $\delta' = 1 - \frac{18d}{q^2}$, we get for any n, q, d there is an efficient $(d, 1 - \frac{18d}{q^2} - \frac{d}{q})$ -HSG, with the argued seed length. As desired. $\square$

Acknowledgment

The author thanks David Zuckerman for helpful discussions.

AI statement. No AI model was used in obtaining any of the ideas or in the writing of the paper (the main ideas were obtained prior to the release of ChatGPT 5.6-sol).

References

- [Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 21–30, 2005.
- [BV10] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM Journal on Computing*, 39(6):2464–2486, 2010.
- [CDG26] Gil Cohen, Dean Doron, and Noam Goldgraber. Optimal prgs for low-degree polynomials over polynomial-size fields. *arXiv preprint arXiv:2602.10030*, 2026.
- [CHL⁺26] Eshan Chattopadhyay, Pooya Hatami, Chin Ho Lee, Shachar Lovett, Avishay Tal, and Emanuele Viola. Exponential correlation bounds for polynomials. *arXiv preprint arXiv:2609.28839*, 2026.
- [Coh26] Gil Cohen. Fooling low-degree polynomials: Prgs with optimal dependence on input length and expander walks. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2026. TR26-243.
- [CTS13] Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 20, page 155, 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 (AP-PROX/RANDOM, Leibniz International Proceedings in Informatics (LIPIcs))*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

- [DS08] Zeev Dvir and Amir Shpilka. Noisy interpolating sets for low degree polynomials. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 140–148. IEEE, 2008.
- [DTST22] Dean Doron, Amnon Ta-Shma, and Roei Tell. On hitting-set generators for polynomials that vanish rarely. *computational complexity*, 31(2):16, 2022.
- [DV22] Harm Derksen and Emanuele Viola. Fooling polynomials using invariant theory. In *2022 IEEE 63rd 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 *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 161–168. IEEE, 2014.
- [KS01] Adam R Klivans and Daniel Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 216–223, 2001.
- [Lov08] Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 557–562, 2008.
- [Lu12] Chi-Jen Lu. Hitting set generators for sparse polynomials over any finite fields. In *2012 IEEE 27th Conference on Computational Complexity*, pages 280–286. IEEE, 2012.
- [LVW93] Michael Luby, Boban Velickovic, and Avi Wigderson. Deterministic approximate counting of depth-2 circuits. In *[1993] The 2nd Israel Symposium on Theory and Computing Systems*, pages 18–24. IEEE, 1993.
- [Sho90] Victor Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of computation*, 54(189):435–447, 1990.
- [Vio09] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d . *Computational Complexity*, 18(2):209–217, 2009.
- [Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Structures & Algorithms*, 11(4):345–367, 1997.