


New Baselines for Local Pseudorandom Number Generators by Field Extensions

Akin Ünal 

Department of Computer Science
ETH Zurich
Zürich, Switzerland
`akin.uenal@inf.ethz.ch`

May 26, 2023

Abstract. We will revisit recent techniques and results on the cryptanalysis of local pseudorandom number generators (PRGs). By doing so, we will achieve a new attack on PRGs whose time complexity only depends on the algebraic degree of the PRG. Concretely, for PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$, we will give an algebraic algorithm that distinguishes between random points and image points of F , whose time complexity is bounded by

$$\exp(O(\log(n)^{\deg F / (\deg F - 1)} \cdot n^{1 - e / (\deg F - 1)}))$$

and whose advantage is at least $1 - o(1)$ in the worst case.

To the best of the author’s knowledge, this attack outperforms current attacks on the pseudorandomness of local random functions with guaranteed noticeable advantage and gives a new baseline algorithm for local PRGs. Furthermore, this is the first subexponential attack that is applicable to polynomial PRGs of constant degree over fields of any size with a guaranteed noticeable advantage.

We will extend this distinguishing attack further to achieve a search algorithm that can invert a uniformly random constant-degree map $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$ with high advantage in the average case. This algorithm has the same runtime complexity as the distinguishing algorithm.

Keywords: PRGs · NC^0 · Local Random Functions · Polynomial Equation Systems · Algebraic Attacks · Subexponential · Lower Bounds

1 Introduction

Pseudorandom Number Generators. A pseudorandom number generator (PRG) is a deterministic algorithm $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ that stretches a seed consisting of n bits to a longer string of m bits. Ideally, a PRG guarantees that, if the seed has been sampled uniformly at random, then the bitstring output by the PRG is indistinguishable from a truly random bitstring of length m for a

certain class of algorithms. PRGs are part of the foundations of modern theoretical cryptography [HILL99, App12]. In the real world, keystream generators (which can be seen as a more advanced form of PRGs) are a popular method to construct fast and reliable symmetric encryption between two parties.

Local PRGs. While it is widely accepted that PRGs do exist from a practical point of view, it is in theory not clear which computational complexity a PRG must have to be reasonably secure while maintaining a high (i.e. polynomial) output length. The most simple but plausibly secure PRGs are so-called *local PRGs*, which have been introduced by Goldreich [Gol11] (also known as *random local functions*, *Goldreich’s PRGs* and *PRGs in \mathbf{NC}^0*). For each of its output bits, a local PRG only needs to look up a constant number of its input bits. This makes local PRGs from a practical and a theoretical point of view highly efficient. In fact, local PRGs have several applications in theory:

1. Ishai, Kushilevitz, Ostrovsky & Sahai [IKOS08] showed that local poly-stretch PRGs (i.e. $m \geq n^{1+e}$ for $e > 0$ constant) together with oblivious transfer imply highly efficient secure communication protocols for two parties evaluating a circuit on private data. In the semi-honest model, where the users abide to the protocol but try to learn as much about the other party’s data as possible, the authors could construct a protocol where the computational complexity for both parties is linear in the size of the circuit to be evaluated. In the malicious model, where users may deviate from the protocol, the authors constructed a protocol where the computational complexity of both parties is slightly superlinear in the size of the circuit. The first protocol has been extended to arithmetic circuits assuming local poly-stretch PRGs $F : k^n \rightarrow k^m$ over a field k [ADI⁺17]. In a recent work, it could even be improved such that it remains secure in the malicious model while maintaining only a constant computational overhead [AK23].
2. Further, local PRGs are interesting for multiparty-computation (MPC) protocols and fully-homomorphic encryption (FHE) schemes. These primitives suffer strongly from computing circuits of large depths. While the complexity of MPC protocols rise with the depth of the to be evaluated circuit, the noise of FHE ciphers grows substantially with the multiplicative depth of the circuit that is to be evaluated on encrypted data. In the case of lattice-based FHE, this forces one to weaken the underlying learning with errors (LWE) assumption to support the evaluation of deeper circuits.
3. Another important application of local PRGs are indistinguishability obfuscation (iO) schemes. Jain, Lin & Sahai [JLS21, JLS22] gave recently new iO schemes whose security is reduced to a number of assumptions, including local PRGs of polynomial stretch.

Subexponential Security. In particular, the recent iO constructions of Jain, Lin & Sahai [JLS21, JLS22] need that the local PRGs they use have stronger security guarantees than usual: the advantage of each poly-time adversary of distinguishing the output of the PRG from uniform randomness must not only be

negligible, but smaller than the inverse of some subexponential functional. They dub this *subexponential security*, and, in fact, require subexponential security of all assumptions they use.

This raises the interest in attack algorithms on local PRGs whose runtime are beyond poly-time and whose advantage is below negligible. While some efficient attacks on local PRGs have been known for quite some time, there has been a recent interest in subexponential attacks on PRGs [BQ09, OW14, AL16, CDM⁺18, Üna23a, Üna23b]. While efficient attacks are only applicable when a certain stretch of the PRG is given, subexponential attacks are applicable even for small polynomial stretches, and degrade gracefully with decreasing stretch of PRGs. Hence, when estimating the concrete security of Goldreich’s PRGs, existing subexponential attacks must be taken into consideration [CDM⁺18].

Faster Baselines for Local PRGs. In [Üna23b], the author gave new algebraic attacks on local PRGs and PRGs of constant degree over large fields. He mistakenly claimed that his attacks would give new baselines for local PRGs. This claim was erroneous as there are so-called *shrinking-set attacks* [Zic17, AIK08] that provably break local PRGs of poly-stretch in subexponential time and are by a log-factor in the exponent faster than the attacks given in [Üna23b].

However, as we will show in this paper, by improving the techniques of [Üna23b] we can substantially lower the time complexity of algebraic attacks on PRGs s.t. their runtimes only depend on the *degree* of PRGs rather than their *locality* and still retain a high advantage. Since the degree of a PRG must always be smaller than its locality, this yields a new algebraic attack on local PRGs that surpasses the shrinking-set attack.

1.1 Contribution

In this work, we will revisit and improve the techniques of [Üna23a, Üna23b]. In Section 3, our improvements will yield a distinguishing attack on PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of constant *degree* d over \mathbb{Z}_2 and poly-stretch $m \geq n^{1+e}$ of subexponential time complexity $2^{O(\log(n)^{d/(d-1)} \cdot n^{1-e/(d-1)})}$ and high advantage $1 - O\left((\log(n)/n^e)^{1/(d-1)}\right)$. We will give a generalized version of this attack that works on constant-degree PRGs $F : k^n \rightarrow k^m$ over *any* field k .

Further, for maps $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ where each output bit is computed by a uniformly random polynomial $f_i \in \mathbb{Z}_2[X_1, \dots, X_n]$ of degree d , we will give an inversion algorithm in Section 4. The inversion algorithm has the same time complexity as the distinguishing attack of Section 3. Its advantage depends on the randomness of F and is high in the average case.

Finally, in Section 6, we will use our new techniques to give reductions between learning parity with noise-problems over fields of different sizes. Our results will show that the learning parity with noise-problem over \mathbb{Z}_2 and over its extension field $\mathbb{F}_{2^{\lceil \log n \rceil}}$ are – up to logarithmic factors in the parameters – equivalent.

To the best of the author’s knowledge, this paper gives the first distinguishing attack on binary constant-degree poly-stretch PRGs that is provably subexponential and has a provable non-negligible advantage, even in the worst case. Since the time complexity of this attack is independent of the locality of the PRG, this attack outperforms other subexponential attacks on local PRGs that have a provably non-negligible advantage in the worst case and gives us new baselines for distinguishing the output of random local functions from true randomness.

1.2 Technical Overview

For a PRG F , we will in the following denote by $\deg F$ its (algebraic) degree over its base field, and by $\text{loc } F$ its locality.

Note that the attacks against (local) PRGs in [Üna23b] are based on two simple tricks: a basic algebraic distinguishing algorithm and a hashing resp. reduction technique. The basic algebraic algorithm has a subexponential runtime and works well against PRGs $F : k^n \rightarrow k^m$ of constant degree over a large field k , however it struggles with PRGs over small fields, e.g. $k = \mathbb{Z}_2 = \{0, 1\}$.

The hashing technique compensates this problem by turning a binary PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ into a PRG $G : k^n \rightarrow k^{m'}$ with $m' \approx m$ s.t. the degree of G equals the locality of F .

We will revisit both techniques in more detail:

The Basic Algebraic Attack. This attack appears in [Üna23b] as well as in [Zic17]. Assume we are given a PRG $F : k^n \rightarrow k^m$ of constant degree over a field k . Denote by $f_1, \dots, f_m \in k[X] := k[X_1, \dots, X_n]$ the polynomials that compute the output values of F .

Now assume that we were aware of an algebraic relationship h between the polynomials f_1, \dots, f_m . I.e., a new polynomial $h \in k[Y] := k[Y_1, \dots, Y_m]$ s.t. $h(Y_1, \dots, Y_m)$ is not the zero polynomial, but $h(f_1(X), \dots, f_m(X))$ is the zero polynomial in $k[X]$. The key observation is that we can use h to check if a point $y \in k^m$ lies in the image of F . Indeed, if y equals $F(x)$ for some $x \in k^n$, then we must have

$$h(F(x)) = (h \circ F)(x) = (h(f_1, \dots, f_m))(x) = (0)(x) = 0. \quad (1)$$

On the other side, because of the famous Schwartz-Zippel Lemma we have for a uniformly random $y \leftarrow k^m$

$$\Pr_{y \leftarrow k^m} [h(y) \neq 0] \geq 1 - \frac{\deg h}{\#k}. \quad (2)$$

A distinguisher that uses h to play the pseudorandomness game for the PRG F would therefore have an advantage of $1 - \frac{\deg h}{\#k}$, which converges towards 1 if the size of k grows faster than the degree of h .

Hence, an algebraic relationship of low degree among the output values of F helps substantially in distinguishing the outputs of F from true randomness.

However, how can we compute such a relationship and, more importantly, how can we bound the degree of h in the worst case?

In [Üna23b], it was shown that h will be sublinear if F is of poly-stretch and of constant degree. In fact, we have $\deg h \in O(n^{1-\frac{e}{a-1}})$ if $m \geq n^{1+e}$ and $d = \deg F$. This was shown by considering the morphism of rings

$$\phi : k[Y_1, \dots, Y_m] \longrightarrow k[X_1, \dots, X_n] \quad (3)$$

$$g(Y_1, \dots, Y_m) \longmapsto g(f_1(X), \dots, f_m(X)) \quad (4)$$

that substitutes each variable Y_i by the polynomial f_i . Now, h is a non-zero kernel element of ϕ of minimal degree. To estimate the degree of h , we restrict ϕ to the subspace

$$k[Y]^{\leq D} := k[Y_1, \dots, Y_m]^{\leq D} := \{g \in k[Y] \mid \deg g \leq D\} \quad (5)$$

of polynomials of degree $\leq D$. When restricted to $k[Y]^{\leq D}$, the image of ϕ is contained in

$$k[X]^{\leq dD} := k[X_1, \dots, X_n]^{\leq dD} := \{g \in k[X] \mid \deg g \leq dD\}. \quad (6)$$

In fact, when we plug the degree- d polynomials f_1, \dots, f_m in $g(Y_1, \dots, Y_m)$ the degree of the resulting polynomial $\phi(g) = g(f_1(X), \dots, f_m(X))$ will be stretched by a factor of at most d . This means, by restricting ϕ to $k[Y]^{\leq D}$ we get a linear map

$$\phi_D : k[Y_1, \dots, Y_m]^{\leq D} \longrightarrow k[X_1, \dots, X_n]^{\leq dD} \quad (7)$$

of spaces of finite dimensions. According to the dimension formula for linear maps we now have

$$\dim \ker \phi_D \geq \dim k[Y_1, \dots, Y_m]^{\leq D} - \dim k[X_1, \dots, X_n]^{\leq dD} \quad (8)$$

$$= \binom{m+D}{D} - \binom{n+dD}{dD}. \quad (9)$$

I.e., we can guarantee the existence of an algebraic relationship of degree $\leq D$ whenever $\binom{m+D}{D} > \binom{n+dD}{dD}$.

In [Üna23a], it was shown that Eq. (9) does hold for some sublinear D . In fact, we have:

Lemma 1 ([Üna23a] Lemma 7). *Let $d \in \mathbb{N}, d \geq 2$. Let $m : \mathbb{N} \rightarrow \mathbb{N}$ be a function with $m(n) \geq 2^{2d-1} \cdot d^{d-1} \cdot n$.*

Then, we have for all integers $n \geq 2d$

$$\binom{m(n)+D(n)}{D(n)} > \binom{n+dD(n)}{dD(n)} \quad (10)$$

where $D(n) = \left\lceil \left(\frac{(2n)^d}{m} \right)^{\frac{1}{d-1}} \right\rceil$. For $m(n) \geq n^{1+e}$, we in particular have for n large enough

$$D(n) = \left\lceil 2^{\frac{d}{d-1}} \cdot n^{1-\frac{e}{d-1}} \right\rceil \in O(n^{1-\frac{e}{d-1}}). \quad (11)$$

Now, how do we find $h \in \ker \phi$ of degree $\leq D \in O(n^{1-\frac{\epsilon}{d-1}})$? Since we know that h is a non-zero element of $\ker \phi_D$, it suffices to write down a matrix representation M of $\phi_D : k[Y]^{\leq D} \rightarrow k[X]^{\leq dD}$. M is a matrix of shape $\binom{m+D}{D} \times \binom{n+dD}{dD}$ over k . The algebraic relation h corresponds to non-zero kernel vector of M . Hence, we can use Gaussian elimination to find h . Eliminating the rows of M costs $O\left(\binom{m+D}{D}^2 \cdot \binom{n+dD}{dD}\right)$ arithmetic operations. Therefore, the memory complexity of finding h is $\binom{m+D}{D} \cdot \binom{n+dD}{dD} \in n^{O(D)} = 2^{O(\log(n) \cdot n^{1-\epsilon/(d-1)})}$ and the time complexity lies in $O\left(\binom{m+D}{D}^2 \cdot \binom{n+dD}{dD}\right) \leq 2^{O(\log(n) \cdot n^{1-\epsilon/(d-1)})}$. Now, evaluating h on a point $y \in k^m$ costs $\deg h \cdot \binom{m+\deg h}{m} \in 2^{O(\log(n) \cdot n^{1-\epsilon/(d-1)})}$ arithmetic operations in k .

In total, we get a distinguishing algorithm with space and time complexities in $2^{O(\log(n) \cdot n^{1-\epsilon/(d-1)})}$ and noticeable advantage for fields k of size $\geq n$. Formally, it was shown:

Theorem 1 ([Üna23a] Theorem 2). *Let $m \in \omega(n)$. Let $F : k^n \rightarrow k^m$ be a PRG of constant degree $\deg F$ over k .*

Then, there is an algebraic attack on F whose time and space complexities are bounded from above by $n^{O((n^{\deg F}/m)^{\frac{1}{\deg F-1}})}$. Further, the advantage of this attack in the pseudorandomness game of F (Definition 4) is lower bounded by

$$\text{adv}_F(\mathcal{A}) \geq 1 - O\left(\left(n^{\deg F}/m\right)^{\frac{1}{\deg F-1}} \cdot \frac{1}{\#k}\right). \quad (12)$$

For $m \geq n^{1+\epsilon}$ and $\#k \geq n$, the attack is a subexponential algorithm with time and space complexities in $2^{O(\log(n) \cdot n^{1-\epsilon/\deg F-1})}$ and high advantage $\geq 1 - O(n^{-\epsilon/(\deg F-1)})$.

Hashing to Larger Fields. Note that the above attack does not fare well against PRGs of constant degree over small fields. In the reign of, let's say, $\mathbb{Z}_2 = \{0, 1\}$, one can alter the above algorithm to find an algebraic relationship h that is reduced modulo the field equations $Y_1^2 - Y_1, \dots, Y_m^2 - Y_m$ and prove that the probability of h vanishing on a random point $y \leftarrow \{0, 1\}^m$ is bounded by

$$\Pr_{y \leftarrow \{0, 1\}^m} [h(y) = 0] \leq 1 - 2^{-\deg h}. \quad (13)$$

However, this will only yield a distinguishing attack with advantage $\geq 2^{-\deg h}$. Since the degree of h will be sublinear in the worst case, we can only guarantee a subexponentially small advantage of this attack, which is very unsatisfactory.

To solve this problem for local PRGs, a simple hashing resp. reduction technique was introduced in [Üna23b]: if we are given a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and a point $y \in \{0, 1\}^m$, the idea is to convert F and y to a new PRG $G : \{0, 1\}^n \rightarrow k^{m'}$ and a new point $y \in k^{m'}$ s.t. the following things hold:

1. The size of the field k is large enough i.e. $\#k \geq n$.
2. m' is only by a small factor smaller than m .

3. y' lies in the image of G if y lies in the image of F .
4. y' is close to being uniformly random over $k^{m'}$ if y is sampled uniformly random from $\{0, 1\}^m$.

To achieve this, the field $k = \mathbb{Z}_p$ for a prime $p \geq n$ was considered in [Üna23b]. Set $m' := \lfloor m/(3 \lceil \log p \rceil) \rfloor \in O(m/\log n)$ and note that the distribution of random matrices $A \leftarrow \mathbb{Z}_p^{m' \times m'}$ gives us a universal family of hash functions of type $\{0, 1\}^m \rightarrow \mathbb{Z}_p^{m'}$. According to the leftover hash lemma, the matrix-vector product $A \cdot y$ is statistically close to being uniform for $y \leftarrow \{0, 1\}^m$. In fact, we can bound its statistical distance from $z \leftarrow \mathbb{Z}_p^{m'}$ by

$$\Delta((A, Ay), (A, z)) \leq 2^{-n}. \quad (14)$$

This justifies to set $y' := A \cdot y \in \mathbb{Z}_p^{m'}$.

Now, we define $G : \{0, 1\}^n \rightarrow \mathbb{Z}_p^{m'}$ as the concatenation $G := A \circ F$. I.e., G first evaluates F normally on its seed and then applies the matrix A to the binary output of F (over \mathbb{Z}_p). It is clear that $y' = A \cdot y$ must lie in $G(\{0, 1\}^n) = A \cdot F(\{0, 1\}^n)$ if y lies in $F(\{0, 1\}^n)$.

We are now in a situation where we can apply the algorithm of Theorem 1 on G and y' to decide if y is an image of F or uniformly random. However, to bound the advantage and runtime of this algorithm, we need to know the algebraic degree of G or, more formally, of a representation of G by constant-degree polynomials. It can be shown that G can be computed by polynomials of degree d when F is of *locality* d . I.e., the degree of G (which influences the performance of our algorithm) equals the locality of F , in general.

In total, we get a distinguishing algorithm for the PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $m \geq n^{1+e}$, with time and space complexities in $2^{O(\log(n)^{\text{loc } F / (\text{loc } F - 1)} \cdot n^{1-e / (\text{loc } F - 1)})}$ and advantage $\geq 1 - O\left(\frac{\log(n)}{n^e}\right)^{1/(\text{loc } F - 1)}$ where $\text{loc } F$ is the *locality* of F (the additional $\log(n)^{1/(\text{loc } F - 1)}$ factor in the time complexity and the advantage of our attack stems from the fact that m' is by a logarithmic factor smaller than m).

A Bad Trade-Off. Note, that this attack has a high advantage, since $n^{e/(\text{loc } F - 1)}$ grows substantially faster than $\log(n)^{1/(\text{loc } F - 1)}$. However, the runtime of our attack worsened: while the time and space complexities of the attack in Theorem 1 only depend on the degree of the PRG, the complexities of the attack that uses the old hashing technique in [Üna23b] depend on its locality. In particular, this attack is not applicable to PRGs that are computed by dense polynomials of constant degree over the binary numbers.

In the following, we will improve the hashing technique and, as a result, gain a new subexponential attack algorithm whose runtime is independent of the locality of the PRG.

The New Extension Technique. The problem when hashing from \mathbb{Z}_2 to \mathbb{Z}_p is that the embedding $\mathbb{Z}_2 \hookrightarrow \mathbb{Z}_p$ that maps zero to zero and one to one is not

homomorphic. While this inclusion preserves multiplication, it does not preserve addition. Indeed, we have $1 + 1 = 0$ in \mathbb{Z}_2 but not in \mathbb{Z}_p if $p > 2$. To get a homomorphic inclusion we need to consider extension fields of \mathbb{Z}_2 .

Set $N := 2^{\lceil \log n \rceil}$ and denote by \mathbb{F}_N the Galois field that has N elements. Up to isomorphism, \mathbb{F}_N is uniquely determined by the number of its elements and there exists a natural homomorphic ring homomorphism $\mathbb{Z}_2 \hookrightarrow \mathbb{F}_N$.

In algebra, it is well known that finite fields are *perfect*, i.e. each extension of finite fields is *separable*. The primitive element theorem postulates that each finite separable field extension $k \subset \bar{k}$ is generated by one element. I.e., there is one $\zeta \in \bar{k}$ s.t. \bar{k} has the basis $1, \zeta, \dots, \zeta^{r-1}$ as k -vector space (where $r = \dim_k \bar{k}$ is the dimension of \bar{k} as k -vector space). In particular, we have

$$\bar{k} = k[\zeta] \cong k \oplus \zeta \cdot k \oplus \zeta^2 \cdot k \oplus \dots \oplus \zeta^{r-1} \cdot k \cong k^r \quad (15)$$

where the second and third equalities are isomorphisms of vector spaces (\oplus denotes the direct sum of vector fields).

Since the extension $\mathbb{Z}_2 \subset \mathbb{F}_N$ is separable, it is also generated by one element. Let's call this element ζ , too. Eq. (15) implies that the linear map

$$\psi : \mathbb{Z}_2^{\lceil \log n \rceil} \longrightarrow \mathbb{F}_N \quad (16)$$

$$(b_1, \dots, b_{\lceil \log n \rceil}) \longmapsto b_1 + b_2 \cdot \zeta + \dots + b_{\lceil \log n \rceil} \cdot \zeta^{\lceil \log n \rceil - 1} \quad (17)$$

is an isomorphism of vector spaces. This leads to the following two observations:

1. Since ψ is bijective, i.e. one-to-one and onto, it maps uniformly random vectors in $\mathbb{Z}_2^{\lceil \log n \rceil}$ to uniformly random elements in \mathbb{F}_N . I.e., if we sample $b \leftarrow \mathbb{Z}_2^{\lceil \log n \rceil}$ then $\psi(b)$ is identically distributed as $c \leftarrow \mathbb{F}_N$.
2. The natural inclusion $\mathbb{Z}_2 \hookrightarrow \mathbb{F}_N$ extends to a natural inclusion of polynomial rings

$$\mathbb{Z}_2[X_1, \dots, X_n] \hookrightarrow \mathbb{F}_N[X_1, \dots, X_n]. \quad (18)$$

In fact, we can consider $\mathbb{Z}_2[X_1, \dots, X_n]$ to be a subring of $\mathbb{F}_N[X_1, \dots, X_n]$. This inclusion preserves functionality and degree: if we have a polynomial $f \in \mathbb{Z}_2[X_1, \dots, X_n]$ we can – via the inclusion $\mathbb{Z}_2[X_1, \dots, X_n] \subset \mathbb{F}_N[X_1, \dots, X_n]$ – consider f to be a polynomial in $\mathbb{F}_N[X_1, \dots, X_n]$. Interpreted as a polynomial in $\mathbb{F}_N[X_1, \dots, X_n]$, f will evaluate on $\{0, 1\}^n$ to the same values as before. Further, f – interpreted as element in $\mathbb{F}_N[X_1, \dots, X_n]$ – has the same degree as $f \in \mathbb{Z}_2[X_1, \dots, X_n]$. This comes from the fact that – when we go from $\mathbb{Z}_2[X_1, \dots, X_n]$ to $\mathbb{F}_N[X_1, \dots, X_n]$ – we only “change” the coefficients of f .

ψ now extends to an isomorphism of vector spaces

$$\psi : (\mathbb{Z}_2[X_1, \dots, X_n])^{\lceil \log n \rceil} \longrightarrow \mathbb{F}_N[X_1, \dots, X_n] \quad (19)$$

$$(f_1, \dots, f_{\lceil \log n \rceil}) \longmapsto f_1 + \zeta \cdot f_2 + \dots + \zeta^{\lceil \log n \rceil - 1} \cdot f_{\lceil \log n \rceil}. \quad (20)$$

This isomorphism is degree-preserving: if $f_1, f_2, \dots, f_{\lceil \log n \rceil}$ are of degree $\leq d$, then so are their scaled versions $f_1, \zeta \cdot f_2, \dots, \zeta^{\lceil \log n \rceil - 1} \cdot f_{\lceil \log n \rceil}$ and the linear combination $\psi(f_1, \dots, f_{\lceil \log n \rceil}) = f_1 + \zeta \cdot f_2 + \dots + \zeta^{\lceil \log n \rceil - 1} \cdot f_{\lceil \log n \rceil}$.

Now, given a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of *degree* d and a point $y \in \{0, 1\}^m$, we want to apply ψ on blocks of length $\lceil \log n \rceil$ of F and y . Assume for the simplicity of this exposition that m is a multiple of $\lceil \log n \rceil$ i.e. there is an $m' \in \mathbb{N}$ s.t. $m = \lceil \log n \rceil \cdot m'$. Further, consider the following matrix ¹ of shape $m' \times m$ that applies the linear map ψ block-wise on its input:

$$A := I_{m'} \otimes (1 \zeta \dots \zeta^{\lceil \log n \rceil - 1}) = \begin{pmatrix} 1 \zeta \dots \zeta^{\lceil \log n \rceil - 1} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & 1 \zeta \dots \zeta^{\lceil \log n \rceil - 1} \end{pmatrix} \quad (21)$$

where $I_{m'}$ denotes the identity matrix of shape $m' \times m'$ and \otimes the Kronecker product. Just as before, we can apply A via matrix-vector multiplication on F and y . Concretely, we compute

$$y' := A \cdot y \in \mathbb{F}_N^{m'}, \quad (22)$$

$$G(X) := A \cdot F(X). \quad (23)$$

Then, y lies in the image of F iff y' lies in the image of G . Further, y is uniformly random from $\{0, 1\}^m$ iff y' is uniformly distributed in $\mathbb{F}_N^{m'}$. Note that for this property we do not need to sample A uniformly at random as before, and we do not need to invoke the leftover hash lemma. However, more importantly, the degree of the PRG $G : \{0, 1\}^n \rightarrow \mathbb{F}_N^{m'}$, which computes $F(x)$ and then applies A , equals the *degree* of F over \mathbb{Z}_2 . This comes from the fact that ψ preserves the algebraic degree of polynomials and A applies ψ block-wise to the polynomials computing F .

It follows that we can invoke the distinguishing algorithm from Theorem 1 on G and y' to decide if y lies in the image of F . If $m \geq n^{1+e}$, then $m' \geq \frac{n^{1+e}}{\lceil \log n \rceil}$ and the attack from Theorem 1 has time and space complexities in

$$2^{O(\log(n)^{\deg G / (\deg G - 1)} \cdot n^{1-e / (\deg G - 1)})} = 2^{O(\log(n)^{\deg F / (\deg F - 1)} \cdot n^{1-e / (\deg F - 1)})} \quad (24)$$

and an advantage of

$$\geq 1 - O\left(\frac{\log(n)^{1/(\deg G - 1)} \cdot n^{1-e/(\deg G - 1)}}{N}\right) \quad (25)$$

$$\geq 1 - O\left(\frac{\log(n)^{1/(\deg G - 1)} \cdot n^{1-e/(\deg G - 1)}}{n}\right) \quad (26)$$

$$= 1 - O\left(\frac{\log(n)^{1/(\deg G - 1)}}{n^{e/(\deg G - 1)}}\right) \quad (27)$$

$$= 1 - O\left(\frac{\log(n)^{1/(\deg F - 1)}}{n^{e/(\deg F - 1)}}\right) \geq 1 - o(1). \quad (28)$$

¹ It is easy to notice the resemblance of A and the gadget matrix of Micciancio & Peikert [MP12] in the setting of lattice-based cryptography.

This gives us a subexponential distinguishing attack on poly-stretch PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ whose runtime is independent of the locality of F . I.e., this attack is even applicable on PRGs computed by dense polynomials of constant degree.

While we only handled the case of binary PRGs here, the extension technique can naturally be used for PRGs $F : k^n \rightarrow k^m$ that are polynomial over any small field k .

We can summarize our result as follows:

Theorem 2 (Main Result). *Let $F : k^n \rightarrow k^m$ be a PRG of constant degree $\deg F$ over some field k . Set $r := \max(\log(n)/\log(\#k), 1)$ and assume $m \in \omega(r \cdot n)$.*

Then, there is an algebraic attack on F whose time and space complexities are bounded from above by $n^{O((\frac{r \cdot n \deg F}{m})^{\frac{1}{\deg F - 1}})}$. Further, the advantage of this attack in the pseudorandomness game of F (Definition 4) is lower bounded by

$$\text{adv}_F(\mathcal{A}) \geq 1 - O\left(\frac{r \cdot n}{m}\right)^{1/(\deg F - 1)} \geq 1 - o(1). \quad (29)$$

For $m \geq n^{1+e}$, the attack is a subexponential-time algorithm with time and space complexities in $2^{O(\log(n) \cdot r^{1/(\deg F - 1)} \cdot n^{1 - \epsilon/(\deg F - 1)})}$ and noticeable advantage $\geq 1 - O((r/n^e)^{1/(\deg F - 1)})$.

Remark 1 (Avoiding Field Extensions). The hashing technique in [Üna23b] did indeed hash from $\{0, 1\}^m$ to $\mathbb{Z}_p^{\lfloor m/(3 \lceil \log p \rceil) \rfloor}$ via a random matrix which induced inevitable information loss. However, the new “extension technique” works via a bijection that maps $\{0, 1\}^{\lceil \log n \rceil \cdot m'}$ to $\mathbb{F}_N^{m'}$ without any information loss. In fact, this technique is actually just a change in how we view the PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and is reversible: a PRG $G : \mathbb{F}_N^n \rightarrow \mathbb{F}_N^m$ of degree d over \mathbb{F} can be seen as a PRG $G' : \{0, 1\}^{\lceil \log(n) \rceil \cdot n} \rightarrow \{0, 1\}^{\lceil \log(n) \rceil \cdot m}$ of degree d over \mathbb{Z}_2 and a PRG $H : \{0, 1\}^n \rightarrow \mathbb{F}_N^m$ of degree d over \mathbb{F} can be seen as a PRG $H' : \{0, 1\}^n \rightarrow \{0, 1\}^{\lceil \log(n) \rceil \cdot m}$ of degree d over \mathbb{Z}_2 . We explain this further in Remark 3.

A Search Algorithm for the Average Case. For learning with errors [Reg05] it is known that there exists a simple search-to-decision reduction technique where one solves a search problem with a secret vector $s \in R^n$ over a polynomially large ring R by guessing for each index $i \in [n]$ a potential coordinate $z \in R$ and asking a decision oracle if the resulting problem instance with hard-wired $s_i = z$ is solvable. If $z \neq s_i$, the resulting decision problem is with high probability not solvable and the decision oracle will refute it. This gives a search algorithm that uses a decision algorithm and has a computational overhead of $n \cdot \#R$.

We can do the same thing for a polynomial inversion problem $F(X) = y$ where $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is of constant degree. This gives a search algorithm whose time complexity equals the runtime of the distinguishing attack in Theorem 2 times a polynomial overhead of $2n$. Unfortunately, we can only prove a high advantage in the average case. Formally, we have:

Theorem 3 (Inversion Attack). *Set $r := \max(\log(n)/\log(\#k), 1)$ and assume $m \in \omega(r \cdot n)$.*

There is an algorithm \mathcal{B} that on input a polynomial map $F : k^n \rightarrow k^m$ and a target value y tries to output a preimage $x \in F^{-1}(y)$. The time complexity of \mathcal{B} is bounded by $\#k \cdot n^{O\left(\left(\frac{r \cdot n \cdot \deg F}{m}\right)^{1/(\deg F - 1)}\right)}$. For each $x \in k^n$, the success probability of \mathcal{B} can be bounded by

$$\Pr_F[\mathcal{B}(F, F(x)) = x] \geq 1 - o(1) \quad (30)$$

if the coefficients of F are drawn from k uniformly and independently at random.

The requirement that F is a uniformly random polynomial map of some constant degree can be relaxed, see Definition 6 and Theorem 7.

Analogous to the decisional version, our findings for our search algorithm shine back to Macaulay matrix-based solving algorithms and yield guarantees for their success probability in the average case. We detail this in Remark 4.

Applicability to Learning Parity with Noise. Learning parity with noise (LPN) [Pie12, BFKL94] consists of the problem of distinguishing a uniformly random binary matrix from an LPN sample

$$(B, Bs + e) \quad (31)$$

where $B \leftarrow \mathbb{Z}_2^{m \times n}$ and $s \leftarrow \mathbb{Z}_2^n$ are uniformly random and $e \in \mathbb{Z}_2^m$ is a vector of Bernoulli-distributed entries. I.e., each entry e_i is zero with large probability $1 - \tau$ and one with small probability τ . Usual parameters are $m \in \Omega(n), \tau \in (0, 0.5)$.

A lot of advanced primitives ([BCGI18, BCG+22, BCG+19a, BCG+19b, BCG+20] e.g.) use variations of LPN where one considers different noise-rates, noise-distributions or higher values of m . A modification we will study here is to cast LPN samples over larger fields (or rings) than \mathbb{Z}_2 . In fact, for any field (or ring) k , one can draw $B \leftarrow k^{m \times n}$ and $s \leftarrow k^n$, and consider the sample $(B, Bs + e)$ where each entry of e is zero with probability $1 - \tau$ and uniformly sampled from k^\times with probability τ . The obfuscation schemes of [JLS21, JLS22], for example, require the LPN assumption to hold over large prime fields \mathbb{Z}_p .

Now, one would like to understand how the hardness of LPN over \mathbb{Z}_2 relates to the hardness of LPN over different fields and rings. Liu, Wang, Yang & Yu gave reductions from LPN over the small field \mathbb{Z}_2 to the large integer ring \mathbb{Z}_{2^λ} , and vice-versa. While they also studied LPN over the extension field \mathbb{F}_{128} of \mathbb{Z}_2 , they did not give any reductions between LPN over \mathbb{Z}_2 and over \mathbb{F}_{128} .

We will fill this gap here. Concretely, we will show that LPN over \mathbb{Z}_2 and an extension field \mathbb{F}_{2^r} are – up to factors of r in the parameters – equivalent. In fact, when given the LPN sample $(B, Bs + e) \in \mathbb{Z}_2^{rm \times (n+1)}$ we can use the matrix A from Eq. (21) to compute

$$(AB, ABs + Ae) \in \mathbb{F}_{2^r}^{m \times (n+1)}. \quad (32)$$

Now, AB is uniformly distributed in $\mathbb{F}_{2^r}^{m \times n}$. If $e \in \mathbb{Z}_2^m$ has a noise-rate of τ , then the noise-rate of Ae is close to $r \cdot \tau$. However, the coordinates of Ae are usually not uniformly random from $\mathbb{F}_{2^r}^\times$ when they are not zero. We can come up for this by replacing $1, \zeta, \dots, \zeta^{r-1}$ in the rows of A by random bases of \mathbb{F}_{2^r} as \mathbb{Z}_2 -vector space.

Vice-versa, if we are given $(B, Bs + e) \in \mathbb{F}_{2^r}^{m \times (n+1)}$ with $s \leftarrow \mathbb{F}_{2^r}^n$, then we can extract from each row of $(B, Bs + e)$ a linear equation $b_i'^T \cdot s' + e_i'$ with $b_i' \in \mathbb{Z}_2^{rn}$, $e_i' \in \mathbb{Z}_2$ and $s' \in \mathbb{Z}_2^{rn}$ unique with $As' = s$. Each new noise term e_i' is zero with probability approximately $1 - \tau$.

In conclusion, we have the following reductions:

- Theorem 4.** 1. Solving LPN over \mathbb{F}_{2^r} with m samples, n secrets and noise-rate of τ is at least as hard as solving LPN over \mathbb{Z}_2 with $r \cdot m$ samples, n secrets and noise-rate of $\tau' \approx \tau/r$.
2. Solving LPN over \mathbb{Z}_2 with m samples, rn secrets and noise-rate of τ is at least as hard as solving LPN over \mathbb{F}_{2^r} with m samples, n secrets and noise-rate of $\tau' \approx \tau$.

We detail both reductions in Section 6. Both reductions can be generalized to arbitrary field extensions $k \subset \bar{k}$. This shows that for each prime p , the difficulties of LPN over different fields of characteristic p are comparable. What we still do not know is how to compare the difficulties of LPN over fields of different characteristics. E.g., we do not know if LPN over \mathbb{Z}_2 is substantially simpler or more difficult than LPN over \mathbb{Z}_p , for $p > 2^\lambda$ prime.

1.3 Related Work

To better understand the new attacks of this work we will present here existing attacks on PRGs of constant degree and constant locality, and finally compare their performances.

Attacks on PRGs of Constant Degree.

Relinearization Attacks. Each known attack on PRGs of constant degree over arbitrarily large fields is of algebraic nature. A first approach is to understand the equation $F(X) = y$ as a polynomial equation system with n variables X_1, \dots, X_n and m polynomial equations $f_1(X) = y_1, \dots, f_m(X) = y_m$. By relinearizing this equation system one can generate a linear equation system, on which one can apply Gaussian elimination. If we have enough equations, i.e. $m \geq \binom{n+\deg F}{\deg F}$, then with high probability [AG11] this linear equation system can be solved for a possible seed x , or at least the satisfiability of the linear equation system can be checked. This leads to a basic attack on polynomial PRGs that is efficient and very reliable (its advantage is provably noticeable). This attack can already be improved: we don't need that m is greater than $\binom{n+\deg F}{\deg F}$, in fact, it suffices that $m \in \Omega(n^{\deg F})$. If m is smaller than $\binom{n+\deg F}{\deg F}$, but has the same asymptotic complexity then it suffices to populate the linear equation system with more

polynomial equations that can be generated from $F(X) = y$ up to some constant degree.

Groebner Bases. Extending the idea of the relinearization-and-elimination algorithm above leads to Groebner basis-based attacks. Groebner bases together with a first algorithm for computing them have been introduced by Buchberger [Buc76]. Faster algorithms have been given by Faugère [Fau99, Fau02], these algorithms are based on Macaulay matrices [Mac16, Laz83]. Additionally, the XL-algorithm with a lot of variations [CKPS00, CCNY12, DBM⁺08, MMDB08, YC05] have been introduced, which also aim to compute something that is similar to a Groebner basis. The core idea of those algorithms is to solve the polynomial equation system $F(X) = y$ by computing a Groebner basis for the ideal $(f_1(X) - y_1, \dots, f_m(X) - y_m) \subset k[X]$ for some monomial ordering. Most algorithms do this by computing a Macaulay matrix up to an increasing degree and applying Gaussian elimination on it: the Macaulay matrix for degree D is the matrix where each row represents a polynomial $X^I \cdot (f_i(X) - y_i)$, for a multi-index I with $|I| \leq D - \deg f_i$, and where each column represents a monomial of $k[X]$ up to degree D . I.e., the rows of the Macaulay matrix are the coefficient vectors of polynomials $X^I \cdot (f_i(X) - y_i)$. The columns are ordered according to the monomial ordering. By applying Gaussian elimination to the Macaulay matrix of degree D one can extract a Groebner basis from it, if D is large enough. In most cases, the Groebner basis will be of the shape $\{X_1 - x_1, \dots, X_n - x_n\}$, which allows to directly read off the solution $X = x \in k^n$ of the polynomial equation system $F(X) = y$. Hence, Groebner basis-based attacks are usually inversion attacks that try to extract the seed $X = x$ from the PRG problem $F(X) = y$.

While Groebner basis-based algorithms perform well in practice, it is hard to give formal guarantees for them. In the worst case, the highest degree of polynomials of a reduced Groebner basis for an equation system $F(X) = y$ is doubly exponential [Dub90]. However, the doubly exponential degree only occurs in extreme cases. On average, the maximum degree for which a Macaulay matrix must be computed is suspected to be upper-bounded by the degree of regularity (in the case of graded anti-lexicographic monomial orders [CG21, CG23]). The degree of regularity is a popular heuristic for Groebner basis-based algorithms, it has been shown to be smaller than $O(n^{1-e/(d-1)})$ for a system of $m \geq n^{1+e}$ equations of degree d [Üna23a]. This would yield an inversion attack of suspected time complexity $n^{O(n^{1-e/(deg F-1)})}$.

In the case of refutation, better bounds can be given: a Groebner basis-based algorithm for refutation problems only checks if the equation $1 = 0$ can be deduced from a Macaulay matrix of sufficiently large degree (so the monomial ordering does not matter, those algorithms are actually just Macaulay matrix-based). If up to some degree the span of the rows of the Macaulay matrix does not contain a vector that corresponds to a constant non-zero polynomial, then the algorithm assumes that the system $F(X) = y$ is solvable and decides that y lies in the image of the PRG F . Otherwise, the algorithm could prove that $F(X) = y$ is unsatisfiable and refutes y . It has been shown [Üna23a] that such algorithms only

need to compute the Macaulay matrix up to some degree in $O(n^{1-e/(\deg F-1)})$. If the base field k is large enough ($\#k \geq n$ e.g.), then this approach will provably have an advantage of $1 - o(1)$. Otherwise, for small fields, the advantage can still be lower bounded by an inverse of a subexponential function. Our results here extend to this algebraic refutation approach: according to Remark 3, it suffices to compute the Macaulay matrix up to some degree in $O(\log(n)^{1/(\deg F-1)} \cdot n^{1-e/(\deg F-1)})$ for this algorithm to have a provable advantage of $1 - o(1)$, even if k is small.

However, we want to point out that extending a PRG $F : k^n \rightarrow k^m$ to $G : k^n \rightarrow \bar{k}^{m'}$, $m' \approx m \cdot \log(\#k) / \log(n)$, will not improve the performance of a Groebner basis resp. Macaulay matrix-based algorithm (i.e. solving $G(x) = y'$ will not be simpler than solving $F(x) = y$ for such algorithms). The reason is that – simply put – the field extension technique here will already be implicitly used when computing Macaulay matrices, even when all computations happen over the small base field k (as Remark 3 explains one polynomial over \bar{k} in m' variables corresponds to many polynomials over k in m variables of the same degree). Therefore, to use the results here on Groebner basis resp. Macaulay matrix-based algorithms, it suffices to increase the degree up to which the Macaulay matrix is computed by a small factor of $\log(n)^{1/(\deg F-1)}$.

New Algebraic Attacks. In his master thesis, Zichron [Zic17], gave a new algebraic attack on polynomial PRGs of constant degree over any field. The new idea is to find an algebraic relationship² among the polynomials computing the output values of the PRG. The author proved that the time complexity of this algorithm is subexponential and gave lower bounds for the advantage of this algorithm [Üna23b]. In the case of small base fields (like $k = \mathbb{Z}_2$), the author gave a hashing technique that improves the advantage substantially by hashing from \mathbb{Z}_2 values to \mathbb{Z}_p values.

On a PRG $F : k^n \rightarrow k^{n^{1+e}}$, the algorithm of Zichron and the author has a time complexity of $n^{O(n^{1-e/(\deg F-1)})}$ and a noticeable advantage if k is large. If k is small, the advantage of this algorithm can only be lower-bounded by a subexponentially small function.

Using the hashing technique of [Üna23b] on PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$ yields an algorithm of time complexity $n^{O(\log(n)^{1/(\log F-1)} \cdot n^{1-e/(\log F-1)})}$ and advantage $1 - o(1)$. However, note that in this case the complexity depends on the locality of the PRG, which makes this algorithm with hashing technique only applicable to binary PRGs of constant *locality*.

Barrier of Applebaum and Lovett. Unfortunately, the time complexity of algebraic algorithms must be subexponential in general. In fact, Applebaum &

² Note that such a relationship $h(Y)$ corresponds to a non-trivial element of the elimination ideal $(f_1(X) - Y_1, \dots, f_m(X) - Y_m) \cap k[Y]$. Since a generating set of this elimination ideal can be computed by a Groebner basis (via an elimination order) we see that algebraic relationships and Macaulay matrices are related. See [Üna23a] for a detailed discussion of parallels.

Lovett proved that the time complexity of an algebraic algorithm deciding if $y \in \{0, 1\}^{n^{1+e}}$ lies in the image of a random local function $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$ is lower-bounded by $n^{O(n^{1-16\cdot e/(d-1)})}$ where d is the *rational degree*³ of the predicate of F [AL18, Theorem 5.4].

It is an interesting open problem to construct new algebraic algorithms that perform provably faster than the barrier of Applebaum and Lovett and have non-negligible advantage. Note that such new algorithms must avoid computing Macaulay matrices.

The Polynomial Method. In a recent line of work [LPT⁺, BKW19, Din, Din21], worst-case algorithms for finding solutions of a polynomial equation system $F(X) = y$ have been given. While these algorithms all have an exponential runtime of $2^{O(n)}$, they significantly beat the typical brute-force search for a solution x by an exponential speed-up.

For solving quadratic systems over \mathbb{Z}_2 with n variables, the first algorithm has been given by Lokshtanov, Paturi, Tamaki, Williams & Yu [LPT⁺] with a worst-case runtime of $O(2^{0.8765n})$. Their algorithm has been improved to have an asymptotic runtime of $\text{poly}(n) \cdot 2^{0.804n}$ by Björklund, Kaski & Williams [BKW19]. Dinur [Din] further improved this asymptotically to $O(2^{0.6943n})$ and gave a better algorithm for concrete parameters [Din21].

At the heart of all of those algorithms is the *polynomial method*, which is usually used to prove lower bounds in circuit complexity theory [Wil14]. The idea is that – instead of checking if $f_i(x) = y_i$ for each $i \in [m]$ separately on input $x \in \{0, 1\}^n$ – we consider the polynomial

$$p(X) := \prod_{i=1}^m (1 + f_i(X) - y_i) - 1 \in \mathbb{Z}_2[X]. \quad (33)$$

$p(X)$ does vanish on input $x \in \{0, 1\}^n$ iff each equation $f_i(x) = y_i$ is fulfilled. However, the degree of p is very large, so instead of looking for roots of p , one uses a randomized simplification \tilde{p} of p with significantly lower degree. By various algebraic tricks, we can now iterate fast over $x \in \{0, 1\}^n$ and check if \tilde{p} vanishes on x .

Attacks on PRGs of Constant Locality. Local PRGs [Gol11] have been subject to cryptanalysis for a long time. The reader can find well-written surveys on their security in [App13, CDM⁺18].

We will try to give here an overview on the most important attacks on local PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of polynomial stretch $m \geq n^{1+e}$ over \mathbb{Z}_2 . For simplicity, we will assume that F is a random local function, i.e. for computing

³ The *rational degree* of a predicate $P : \{0, 1\}^{\text{loc } F} \rightarrow \{0, 1\}$ is defined as the smallest number e s.t. there exist polynomials $Q, R \in \mathbb{Z}_2[X_1, \dots, X_{\text{loc } F}]$ of degree e that fulfil $P(X) \cdot Q(X) = R(X) \pmod{(X_1^2 - X_1, \dots, X_{\text{loc } F}^2 - X_{\text{loc } F})}$ and $Q \neq 0$. In other words, P can be written as the rational function $P(X) = \frac{R(X)}{Q(X)}$ of degree e whenever Q does not evaluate to zero.

each output bit it evaluates the same predicate $P : \{0, 1\}^{\text{loc } F} \rightarrow \{0, 1\}$ on a constant set of input bits. The set of input bits depends on the corresponding output bit. The security of random local functions has been extensively discussed by Applebaum [App13].

Linear Tests as Sanity Checks. A linear test is a degree-1 polynomial $L \in \mathbb{Z}_2[Y_1, \dots, Y_m]$ that gets evaluated on the potential output y of a PRG. If $y \leftarrow \{0, 1\}^m$ is uniformly random, then the output $L(y)$ is balanced. However, for an image $y = F(x)$, the output $L(y)$ may be biased towards zero or one. In fact, the bias of a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is defined to be the maximum of distinguishing advantages of linear tests of F , and one strives to construct PRGs of negligible bias.

It has been shown that the bias of F depends mainly on the predicate P that is used by F [ABR12]. Further, for random local functions linear tests are a good first check to probe their security: if a random local function has a low bias, it is also secure against a large corpus of other attacks [ABR12, App13]. Additionally, we know of random local functions that have a provably negligible bias: in fact, [MST03] constructed a predicate of locality 5 and a local PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^{cn}$ that uses this predicate s.t. the bias of F is provably $2^{-n/O(c^4)}$ (for small choices of $c > 1$).

Also, Viola [Vio08] showed that one can generically create PRGs that are secure against low-degree tests by using PRGs of low bias.

Approximation-Based Attacks. Bogdanov and Qiao [BQ09] showed that sufficiently close approximations x' of the seed x can help to invert the function $F(x)$ efficiently. This leads to a subexponential inversion attack on local PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$. This attack iterates over a set of $2^{O(n - \frac{n}{2 \log c^F})}$ bitstrings x' of length n . With good probability one of those bitstrings x' will be very close to the correct solution x . By using this approximate solution x' one can then invert $F(x) = y$ efficiently and check if the yielded solution is correct.

This leads to an inversion attack with time complexity $2^{O(n^{1 - \frac{e}{2 \log c^F}})}$.

Correlation-Based Attacks. A predicate $P : \{0, 1\}^{\text{loc } F} \rightarrow \{0, 1\}$ is called c -correlated if there are c different input variables X_{i_1}, \dots, X_{i_c} s.t. $P(X) \oplus X_{i_1} \oplus \dots \oplus X_{i_c}$ is unbalanced (i.e. its probability to evaluate to zero on a random input $x \leftarrow \{0, 1\}^{\text{loc } F}$ is not $1/2$). A predicate that has a high correlation is also called *resilient* (when we say “high correlation” we mean it is only c -correlated for large values of c).

Local random functions of low correlation can be efficiently inverted if they have sufficient stretch [App13, ABW10, MST03, BQ09]. In fact, if $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ stems from a predicate of correlation c , and we have $m \in \Omega(n^{c/2}) + \omega(n)$, then we can deduce from $\Omega(n^{c/2})$ equations of the system $F(X) = y$ a new system of $\Omega(n)$ noisy equations of the shape $X_i \oplus X_j = y'_i$. By using an SDP algorithm [CW04, GW95] on this noisy system of locality 2, we can extract an approximation of x . By using this approximation and $\omega(n)$ fresh equations of $F(X) = y$ we can efficiently deduce the correct solution x [BQ09].

Siegenthaler showed that a predicate can either have a high algebraic degree or a high correlation, but not both [Sie84]. In fact, a balanced predicate $P : \{0, 1\}^{\text{loc } F} \rightarrow \{0, 1\}$ must be c -correlated for $c \leq \text{loc } F - \text{deg } F - 1$. This leads to the following important attack: Given a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of locality $d = \text{loc } F$ and stretch $m \in \omega(n^{\lfloor 2d/3 \rfloor / 2})$ that stems from a predicate $P : \{0, 1\}^d \rightarrow \{0, 1\}$, we can distinguish two cases: first case, we have $\text{deg } F \leq \lfloor d/3 \rfloor \leq \lfloor 2d/3 \rfloor / 2$. In this case, we can invert F by using relinearization (with high probability), since we have $m \in \omega(n^{\lfloor 2d/3 \rfloor / 2})$ many equations. Second case, we have $\text{deg } F > \lfloor d/3 \rfloor$. In this case, we have for the correlation $c \leq d - \lfloor d/3 \rfloor - 1$ and, hence, $c \leq \lfloor 2d/3 \rfloor$. In this case, the correlation-based attack can be applied to efficiently invert $F(x)$. It follows that a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of stretch $m \in \omega(n^{\lfloor 2d/3 \rfloor / 2})$ and locality d cannot be secure, at all.

O'Donnell and Witmer [OW14] showed that we really need a stretch of $m \in \omega(n^{\lfloor 2d/3 \rfloor / 2})$ and that the correlation-based attack does not degrade gracefully for smaller stretches m . I.e., if the stretch m lies in $o(n^{\lfloor 2d/3 \rfloor / 2})$ the above attack does not yield a subexponential attack, since the SDP algorithms have a minimum number of equations they need to use.

Guess-and-Determine Attacks. Couteau, Dupin, Méaux, Rossi & Rotella [CDM⁺18] gave multiple new attacks on PRGs $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+\epsilon}}$ that they call *guess-and-determine* attacks. Their attacks are subexponential inversion attacks.

The PRG they attack does not need to be local, but it needs to use a predicate $P : \{0, 1\}^{d+\ell} \rightarrow \{0, 1\}$ of the following form

$$P(X) = M(X_1, \dots, X_d) \oplus X_{d+1} \oplus \dots \oplus X_{d+\ell} \quad (34)$$

for some unbalanced predicate $M : \{0, 1\}^d \rightarrow \{0, 1\}$. Note that we have $\text{deg } F \leq d < \text{loc } F$ if F is not to be trivially broken. This is important, since d determines the time complexity of their attack.

Their first attack proceeds as follows: In a so-called *selection phase*, their attack chooses (intelligent and greedily) $n^{1-\epsilon/(d-1)}$ input bits s.t. sufficiently many equations of $F(X) = y$ will become linear once these input bits are fixed to constant values. In a second phase, the algorithm iterates over all possible values for these input bits. For each assignment, the algorithm gets a linear equation system in the remaining variables with at least n equations. At this point, the algorithm distinguishes two cases: if the resulting matrix has full rank, the algorithm solves the linear equation system and receives a possible assignment for the remaining inputs. Given such a candidate seed x' , the algorithm can check if x' is a correct solution for $F(X) = y$. Otherwise, if the yield matrix does not have full rank, the algorithm can deduce multiple linear equations that does have to hold over $y = F(x)$. If all of these equations do hold over y , then the algorithm knows that, with noticeable probability, y must lie in the image of F .

The runtime of this algorithm is $\text{poly}(n) \cdot 2^{n^{1-\epsilon/(d-1)}}$, and in the special case of the tri-sum-and predicate, the runtime of their algorithm is $\text{poly}(n) \cdot 2^{n^{1-\epsilon/2}}$.

Couteau, Dupin, Méaux, Rossi & Rotella could show that this attack has a noticeable advantage at distinguishing the output of F from uniform randomness for predicates $P = M \oplus \bigoplus_{i=1}^{\ell} X_{2+i}$ with $M = X_1 X_2$. For other choices of $M : \{0, 1\}^d \rightarrow \{0, 1\}$, they can only prove a high advantage of their algorithm assuming a specially tailored assumptions which depends on $\text{loc } F$.

Further, they use the notion of *bit fixing algebraic immunity* [MJSC16] and show that the selection phase of their algorithm can be used to exploit low bit fixing algebraic immunity of predicates. In fact, if the algebraic immunity of a predicate M deteriorates fast, then instead of trying to extract a linear system the algorithm tries to fix less input bits to extract a polynomial equation system of low degree. This polynomial equation system can then be subexponentially solved by a Groebner basis-based algorithm.

Yang, Guo, Johansson & Lentmaier [YGJL22] improved the guess-and-determine attack in the case of the famous tri-sum-and predicate of locality 5. They could show that their attack achieves a better runtime of $\text{poly}(n) \cdot 2^{n^{1-e}/4}$. Further, they gave a new attack, which they dubbed *guess-and-decode* attack, that works similarly but, instead of inverting a linear subsystem, tries to decode a quadratic equation system by using belief propagation. While their guess-and-decode has the same asymptotic time complexity $\text{poly}(n) \cdot 2^{n^{1-e}/4}$, and performs better in practice, they could only verify experimentally its success probability.

Shrinking-Set Attacks. A *shrinking set* for a local PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$ is a subset $T_x \subset [n]$ of input bit positions s.t. the set

$$T_y = \{j \in [m] \mid f_j \text{ only depends on input bits in } T_x\} \quad (35)$$

of output bits that only depend on the input bits specified by T_x is truly larger than T_x , i.e. $\#T_y > \#T_x$. Zichron [Zic17] showed that F must have a shrinking set T_x of size $\#T_x \in \Theta(n^{1-e/(\text{loc } F-1)})$. It can be easily shown that, when we sample $T_x \subset [n]$ uniformly at random of size $\Theta(n^{1-e/(\text{loc } F-1)})$, T_y will have more elements than T_x with non-negligible probability.

This leads to the following distinguishing attack [AIK08] on F : sample $T_x \subset [n]$ uniformly at random of size $\Theta(n^{1-e/(\text{loc } F-1)})$. If T_y is larger than T_x , we essentially have a smaller PRG

$$F' : \{0, 1\}^{\#T_x} \longrightarrow \{0, 1\}^{\#T_y} \quad (36)$$

that maps the input bits chosen by T_x to the output bits specified by T_y . Given a string $y \in \{0, 1\}^{n^{1+e}}$ and the task to decide if y lies in the image of F , we can instead consider the substring $y' := (y_j)_{j \in T_y}$ and check if y' lies in the image of F' . Indeed, if $y = F(x)$ then $y' = F'(x')$ (for a substring x' of x). Otherwise, if y is uniformly random, then y' will not lie in image of F' with probability at least $1/2$ (since $\#T_y > \#T_x$). Therefore, it suffices to check if y' lies in the image of F' . This can be done by brute-force: we iterate over all possible assignments for the input bits chosen by T_x and see if y' is a possible output of F' .

Since the size of T_x lies in $O(n^{1-e/(\text{loc } F-1)})$, this yields a distinguishing attack of time complexity $O(2^{n^{1-e/(\text{loc } F-1)}})$. Note that this attack can be combined with

inversion attacks of [LPT⁺, BKW19, Din, Din21] to achieve better constants in the exponent.

Comparing Runtimes. We will now consider the problem of distinguishing the outputs of a local PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of poly-stretch $m \geq n^{1+e}$ from true randomness. Remember that the attack that we present in this work has a time complexity upper-bounded by

$$n^{O(\log(n)^{1/(\deg F-1)} \cdot n^{1-e/(\deg F-1)})} = 2^{O(\log(n)^{\deg F/(\deg F-1)} \cdot n^{1-e/(\deg F-1)})} \quad (37)$$

and an advantage lower-bounded by

$$1 - O(\log(n)^{1/\deg F} / n^{e/(\deg F-1)}) \geq 1 - o(1). \quad (38)$$

From all attacks listed in this overview, only three have a (potentially) faster runtime for small values of e than this attack, since the runtimes of all other attacks depend on the locality of d or are directly exponential. Those attacks are the basic algebraic attack that is based on algebraic relationships [Zic17, Üna23a, Üna23b], the Groebner basis-based resp. Macaulay matrix-based attacks [Fau99, Fau02, CKPS00, CCNY12, DBM⁺08, MMDB08, YC05] and the guess-and-determine attacks of Couteau, Dupin, Méaux, Rossi & Rotella [CDM⁺18] and Yang, Guo, Johansson & Lentmaier [YGJL22]. We will compare each of those attacks individually with our attack:

1. The basic algebraic attack [Zic17, Üna23a, Üna23b] from Theorem 1 works by finding an algebraic relationship among the polynomials $f_1, \dots, f_m \in \mathbb{Z}_2[X]$ (without going from \mathbb{Z}_2 to a larger field). The runtime of this algorithm is upper-bounded by $n^{O(n^{1-e/(\deg F-1)})} = 2^{O(\log(n) \cdot n^{1-e/(\deg F-1)})}$ which is by the factor $\log(n)^{1/(\deg F-1)}$ in the exponent smaller than the runtime of the algorithm of Theorem 2. However, in the case of binary PRGs, the advantage of this algorithm is only known to be lower-bounded by a subexponentially small function (note, that for this case the algorithm must be slightly adapted s.t. it retrieves an algebraic relationship that is reduced modulo the field equations of \mathbb{Z}_2 , the details are described in [Üna23a]). Hence, while the basic algebraic algorithm is faster than the algorithm of Theorem 2, it only gives an unsatisfactory advantage for breaking the pseudorandomness of binary PRGs.
2. When trying to solve the equation system $F(X) = y$, the algorithms of Faugère [Fau99, Fau02] and the XL-algorithms [CKPS00, CCNY12, DBM⁺08, MMDB08, YC05] will aim to compute a Groebner basis for the ideal $(f_1(X) - y_1, \dots, f_m(X) - y_m)$. While this may work well in praxis, in the theoretical very worst case, computing a Groebner basis may have a doubly exponential time complexity.

In the simpler case of only checking the system of $F(X) = y$ for satisfiability, it suffices to consider algorithms that compute the Macaulay matrix up to

some degree D and inspect if they can deduce a contradiction from it. Our results here show that there is some $D \in O(\log(n)^{1/(\deg F-1)} \cdot n^{1-e/(\deg F-1)})$ s.t. it suffices to check the Macaulay matrix up to degree D to have a high distinguishing advantage. It follows that for the Macaulay matrix-based algorithm and the algorithm from Theorem 2, we can give the same bound $2^{O(\log(n)^{\deg F/(\deg F-1)} \cdot n^{1-e/(\deg F-1)})}$ on the time complexity and the same bound $1 - o(1)$ on their advantages. However, following the discussion in [Üna23a], while both algorithms have the same asymptotic bounds, it is known that the hidden constants of the algorithm from Theorem 2 are better in the non-uniform model: in fact, if the PRG F is known ahead, the algebraic relationship h can be computed in a *preprocessing phase*. In the actual security game, the non-uniform adversary then gets the relationship h as a hint and can solve the fresh challenge y by simply evaluating h on it. This is faster than computing the Macaulay matrix of $F(X) = y$ up to degree $D = \deg F \cdot \deg h$.

3. The guess-and-determine attacks of Couteau, Dupin, Méaux, Rossi & Rotella [CDM⁺18] and Yang, Guo, Johansson & Lentmaier [YGJL22] have a time complexity of $2^{O(n^{1-e/(d-1)})}$ for an integer $d \in [\deg F, \text{loc } F - 1]$ that depends on the predicate used by F . In the best cases, we have $d = \deg F$ and the algorithm performs by a factor of $\log(n)^{\deg F/(\deg F-1)}$ in the exponent faster than the algorithm from Theorem 2. However, in the general case, d will be larger than the degree of the predicate of F (for example the predicate $P(X_1, \dots, X_\ell)$ could be replaced by $P(X_1, \dots, X_\ell) \oplus P(X_{\ell+1}, \dots, X_{2\ell})$). Hence, in most cases the guess-and-determine attacks will perform worse than the algorithm from Theorem 2.

Acknowledgements. Before writing this paper I presented its results in a talk at the crypto seminar at NYU. I want to thank the audience for interesting discussions we had while I gave my talk.

2 Preliminaries

2.1 Notation

Denote by $\mathbb{N} = \{1, 2, 3, \dots\}$ the set of natural numbers.

For a field k , we denote by $k[X_1, \dots, X_n]$ the ring of polynomials in n variables over k . Sometimes, by abuse of notation, we will just write $k[X] := k[X_1, \dots, X_n]$. For $f \in k[X]$, we denote by $\deg f$ its degree, that is its *total degree*.

For a prime $p \in \mathbb{N}$, we denote by $\mathbb{Z}_p := \mathbb{Z}/p\mathbb{Z}$ the field with p elements. Further, for $e \in \mathbb{N}$, we denote by \mathbb{F}_{p^e} the Galois field with p^e elements. Up to isomorphism, \mathbb{F}_{p^e} is uniquely determined by its size.

In this text, n will always denote the *security parameter*. An additional parameter is given by the stretch $m = m(n)$ that depends on n . We will always tacitly assume that m is time-constructible.

Let V be some vector space. Given subspaces $V_1, \dots, V_r \subset V$ that fulfil

$$V_i \cap \left(\sum_{j \neq i} V_j \right) = 0 \quad (39)$$

for each $i \in [r]$, we will denote by $V_1 \oplus \dots \oplus V_r$ the smallest subspace of V that contains each V_i . I.e., $V_1 \oplus \dots \oplus V_r$ equals $V_1 + \dots + V_r$, but by using the \oplus symbol we emphasize that the enumeration of basis vectors of V_1, \dots, V_r stays linearly independent.

2.2 Mathematical Preliminaries

Definition 1. Let k, \bar{k} be fields. If there exists a homomorphism of rings $\iota : k \rightarrow \bar{k}$, we will call the pair k, \bar{k} a **field extension**. Note that each ring homomorphism must send 1 to 1, therefore each ring homomorphism must be injective on fields. In particular, $\iota : k \rightarrow \bar{k}$ is one-to-one and – without loss of generality – we can assume that k is a subset of \bar{k} . By abuse of notation, we will denote field extensions always as subset-relationships $k \subset \bar{k}$.

We will repeat here a well-known fact from algebra and the theory of field extensions.

Lemma 2. Let $k \subset \bar{k}$ be an extension of finite fields. Then, $k \subset \bar{k}$ is simple, i.e., there exists an element $\zeta \in \bar{k}$ s.t. $\bar{k} = k[\zeta]$. I.e., each element of \bar{k} can be written as $f(\zeta)$ where $f \in k[Z]$ is a univariate polynomial.

Proof (Sketch). First note that the unit group $\bar{k}^\times = \bar{k} \setminus \{0\}$ must be cyclic. Otherwise, there would be a proper divisor $d | (\#\bar{k} - 1)$ s.t. we have $x^d = 1$ for each $x \in \bar{k}^\times$. However, the polynomial $X^d - 1$ can have at most $d < \#\bar{k}^\times$ roots in \bar{k} . Ergo, there must exist at least one element in \bar{k}^\times of proper order $\#\bar{k} - 1$.

Let $\zeta \in \bar{k}$ be a generator of \bar{k}^\times . Then, we have $k[\zeta] = \bar{k}$. In fact, besides zero each element of \bar{k} can be written as a power of ζ .

Note that each simple and finite field extension $k \subset \bar{k}$ that is generated by one element ζ can be written as

$$k[\zeta] = \bar{k} \cong k \oplus \zeta \cdot k \oplus \dots \oplus \zeta^{r-1} \cdot k \cong k^r \quad (40)$$

where $r = [\bar{k} : k] := \dim_k \bar{k}$ is the *degree* of the extension $k \subset \bar{k}$. I.e., as a k -vector space \bar{k} has the basis $1, \dots, \zeta^{r-1}$ and each element $c \in \bar{k}$ has a unique representation

$$c = b_1 + b_2 \cdot \zeta + \dots + b_r \cdot \zeta^{r-1} \quad (41)$$

for $b_1, \dots, b_r \in k$.

We cite here the general Schwartz-Zippel Lemma, which will be implicitly used by Theorem 5.

Lemma 3 (Schwartz-Zippel [DL78, Zip79, Sch80]). *Let k be any field and let $S \subset k$ be a finite set. Let $h \in k[Y_1, \dots, Y_m]$ be a polynomial in m variables. We have*

$$\Pr_{y \leftarrow S^m} [h(y) = 0] \leq \frac{\deg h}{\#S}. \quad (42)$$

2.3 Cryptographic Preliminaries

In the following, we will revisit the definitions and results from [Üna23a, Üna23b] for pseudorandom number generators. For compactness, the definitions here will be less detailed than in [Üna23a, Üna23b].

Definition 2 (Pseudorandom Number Generators). *Let k be a field. A **pseudorandom number generator (PRG)** is a deterministic algorithm*

$$F : k^n \longrightarrow k^m \quad (43)$$

that is parametrized by n .

We call $m = m(n)$ the **stretch** of F . If $m \geq n^{1+e}$ for some constant $e > 0$, we call F a **poly-stretch PRG**.

Definition 3 (Locality and Degree of PRGs). *Let $F : k^n \rightarrow k^m$ be a PRG over a finite field k . Let $f_1, \dots, f_m : k^n \rightarrow k$ be the functions that compute the corresponding output values of F . Note, that each f_i can be computed by a polynomial in $k[X_1, \dots, X_n]$. In this text, we will always assume that each f_i is in fact a polynomial in $k[X_1, \dots, X_n]$ that is reduced modulo the field equations of k .*

1. We define the (algebraic) **degree** of F as the maximum of all degrees of the polynomials computing its output values. I.e.

$$\deg F := \max_{i \in [m]} (\deg f_i) \quad (44)$$

2. We define the **locality** of a polynomial $f \in k[X_1, \dots, X_n]$ as the number of variables that occur non-trivially in f . I.e.

$$\text{loc } f := \min \{ \#S \mid S \subseteq [n], f \in k[X_i \mid i \in S] \}. \quad (45)$$

We define the **locality** of F as the maximum of all localities of the polynomials computing its output values. I.e.

$$\text{loc } F := \max_{i \in [m]} (\text{loc } f_i) \quad (46)$$

Definition 4 (Security Game for PRGs). *Let $F : k^n \rightarrow k^m$ be a PRG over a finite field k . The **security-game** for F with an adversary \mathcal{A} is given by:*

1. A challenger draws a bit $b \leftarrow \{0, 1\}$. If $b = 0$, it samples a preimage $x \leftarrow k^n$ uniformly at random, computes $F(x)$ and sends $(F, F(x))$ to \mathcal{A} . If $b = 1$, it samples $y \leftarrow k^m$ and sends (F, y) to \mathcal{A} .

2. \mathcal{A} receives (F, y^*) for some $y^* \in k^m$ and must decide which bit b has been drawn by the challenger. \mathcal{A} makes some computations on its own and finally sends a bit b' to the challenger.

\mathcal{A} wins an instance of this game iff $b = b'$ holds at the end. We define \mathcal{A} 's advantage against F by

$$\text{adv}_F(\mathcal{A}) := 2 \cdot \Pr[\mathcal{A} \text{ wins}] - 1 \quad (47)$$

$$= \Pr_{x \leftarrow k^n}[\mathcal{A}(F, F(x)) = 0] + \Pr_{y \leftarrow k^m}[\mathcal{A}(F, y) = 1] - 1 \quad (48)$$

where we take the probability over the randomness of \mathcal{A} and the challenger.

Definition 5. We say that an algorithm is **subexponential** if there is a constant $e \in [0, 1)$ s.t. its time complexity lies in $2^{O(n^e)}$.

Theorem 5 ([Üna23a] Theorem 2). Let $m \in \omega(n)$. Let $F : k^n \rightarrow k^m$ be a PRG of constant degree $\deg F$ over k .

Then, there is an algebraic attack on F whose time and space complexities are bounded from above by $n^{O((n^{\deg F}/m)^{\frac{1}{\deg F-1}})}$. Further, the advantage of this attack in the pseudorandomness game of F (Definition 4) is lower bounded by

$$\geq 1 - O\left(\left(n^{\deg F}/m\right)^{\frac{1}{\deg F-1}} \cdot \frac{1}{\#k}\right). \quad (49)$$

Remark 2. In [Üna23a], Theorem 5 has only been stated for fields $k = \mathbb{Z}_p$, however, it is clear that the attack and the proof of Theorem 5 work for any field k (in fact, the algorithm \mathcal{B} that is used as a subroutine of this attack is stated for any field and the Schwartz-Zippel lemma works for any field, too).

While in Definition 4 we only define the security of PRGs with regard to the uniform distribution over k^n , the attack of Theorem 5 works for any distribution of seeds in k^n . The reason is that the attack will always correctly recognize image points $F(x)$ independent of the seed x . However, a random point $y \leftarrow k^m$ will only be refuted by the attack with probability $\geq 1 - O\left(\left(n^{\deg F}/m\right)^{\frac{1}{\deg F-1}} \cdot \frac{1}{\#k}\right)$.

3 The New Distinguishing Attack

We will prove here our main result:

Theorem 6. Let $F : k^n \rightarrow k^m$ be a PRG of constant degree $\deg F$ over some field k . Set $r := \max(\log n / \log \#k, 1)$ and assume $m \in \omega(r \cdot n)$.

Then, there is an algebraic attack on F whose time and space complexities are bounded from above by $n^{O\left(\left(\frac{r \cdot n}{m}\right)^{1/(\deg F-1)}\right)}$. Further, the advantage of this attack in the game of Definition 4 against F is lower bounded by

$$\geq 1 - O\left(\left(\frac{r \cdot n}{m}\right)^{1/(\deg F-1)}\right) \geq 1 - o(1). \quad (50)$$

We will need the following lemma:

Lemma 4. *Let $k \subset \bar{k}$ be a field extension generated by an element $\zeta \in \bar{k}$. Let $r = [\bar{k} : k] = \dim_k \bar{k}$ be the degree of this extension.*

Then, the $m' \times (rm')$ matrix

$$A_{m'} := I_{m'} \otimes (1 \ \zeta \ \dots \ \zeta^{r-1}) = \begin{pmatrix} 1 \ \zeta \ \dots \ \zeta^{r-1} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & 1 \ \zeta \ \dots \ \zeta^{r-1} \end{pmatrix} \quad (51)$$

gives an isomorphism of k -vector spaces

$$A_{m'} : k^{rm'} \longrightarrow \bar{k}^{m'} \quad (52)$$

and a degree-preserving isomorphism of k -vector spaces

$$A_{m'} : (k[X_1, \dots, X_n])^{rm'} \longrightarrow (\bar{k}[X_1, \dots, X_n])^{m'} \quad (53)$$

by left-multiplication.

Proof. Note, that $1, \zeta, \dots, \zeta^{r-1}$ is a basis of \bar{k} as a k -vector space. Indeed, since ζ generates \bar{k} over k , each element of \bar{k} can be written as a polynomial $g(\zeta)$ evaluated on ζ . Since the algebraic degree of ζ over k is r , there must exist a non-zero polynomial $h \in k[Z]$ of degree exactly r that has ζ as root. Therefore, the polynomial representation $g \in k[Z]$ of any element of $\bar{k} = k[\zeta]$ can be reduced modulo h to a polynomial of degree $< r$.

Since each element of $c \in \bar{k}$ has a unique representation as a linear combination

$$c = b_1 + b_2 \cdot \zeta + \dots + b_r \cdot \zeta^{r-1} \quad (54)$$

for b_1, \dots, b_r , $A_{m'}$ gives us a linear bijective map from $k^{rm'}$ to $\bar{k}^{m'}$.

Now, if we are given polynomials $f_1, \dots, f_r \in k[X_1, \dots, X_n]$ of degree $\leq d$, then $f_1 + f_2 \cdot \zeta + \dots + f_r \cdot \zeta^{r-1}$ is an element of $\bar{k}[X_1, \dots, X_n]$. Since scaling and adding polynomials does not increase their degree, the degree of $f_1 + f_2 \cdot \zeta + \dots + f_r \cdot \zeta^{r-1}$ is at most d (in fact, it equals the maximum of degrees of f_1, \dots, f_r). It follows, that for each $d \in \mathbb{N}$, $A_{m'}$ gives us a bijective linear map

$$A_{m'} : (k[X_1, \dots, X_n]^{\leq d})^{rm'} \longrightarrow (\bar{k}[X_1, \dots, X_n]^{\leq d})^{m'}. \quad (55)$$

Proof (Theorem 6). Denote by \mathcal{A} the algorithm from Theorem 5. \mathcal{A} is an algorithm that – when given a PRG $G : \bar{k}^n \rightarrow \bar{k}^{m'}$ of degree d and a point $y' \in \bar{k}^{m'}$ – will always output 0 if $y' \in G(\bar{k}^n)$ and will output 1 with probability $\geq 1 - O\left(\left(n^d/m'\right)^{\frac{1}{d-1}} \cdot \frac{1}{\#\bar{k}}\right)$ for $y' \leftarrow \bar{k}^{m'}$. The runtime of \mathcal{A} is bounded by $\leq n^{O((n^d/m')^{1/(d-1)})}$.

Let k be a field and $F : k^n \rightarrow k^m$ a PRG of degree $d := \deg F$ over k . We can assume that $\#k < n$, since otherwise the claim of the theorem follows directly from Theorem 5. Set $r := \left\lceil \frac{\log n}{\log \#k} \right\rceil = \lceil \log_{\#k} n \rceil$ and $m' := \lfloor \frac{m}{r} \rfloor$.

We will attack the pseudorandomness of F by giving a reduction \mathcal{R} that transforms F to a PRG whose pseudorandomness can be broken by \mathcal{A} with noticeable advantage. On input $F : k^n \rightarrow k^m$ and a point $y \in k^m$, \mathcal{R} has to decide if y lies in the image of F or has been sampled uniformly at random. We assumed that the size of k is less than n , so k is finite. \mathcal{R} now constructs an extension field \bar{k} of k s.t. \bar{k} has $(\#k)^r \geq n$ elements. In particular, the algebraic degree of the extension $k \subset \bar{k}$ will be r . According to Lemma 2, there is an element $\zeta \in \bar{k}$ that generates \bar{k} over k , i.e.

$$\bar{k} = k[\zeta]. \quad (56)$$

\mathcal{R} can find an extension field \bar{k} together with a generator ζ by searching for an irreducible polynomial $g \in k[Z]$ of degree r , which can be done in poly-time. When given g , \bar{k} is isomorphic to $k[Z]/(g(Z))$ and the residue class of Z corresponds to the generator ζ . Alternatively, if \mathcal{R} already knows a suitable extension field \bar{k} of k , \mathcal{R} can find a generator for this extension by searching for an element $\zeta \in \bar{k}^\times$ whose multiplicative order is exactly $\#\bar{k} - 1$. The time complexity of searching for ζ is at most polynomial in the size of \bar{k} .

Given $k \subset \bar{k}$ and ζ , \mathcal{R} computes the matrix $A_{m'} \in \bar{k}^{m' \times (r \cdot m')}$ from Lemma 4. Note that $r \cdot m' \leq m < (r+1) \cdot m'$. Since rm' may be smaller than m , \mathcal{R} constructs a new matrix $B \in \bar{k}^{m' \times m}$

$$B = (A_{m'} \ 0_{m' \times (m - rm')}) = \begin{pmatrix} 1 \ \zeta \ \dots \ \zeta^{\lceil \log n \rceil - 1} & & & 0 \ \dots \ 0 \\ & \ddots & & \vdots \\ & & 1 \ \zeta \ \dots \ \zeta^{\lceil \log n \rceil - 1} & 0 \ \dots \ 0 \end{pmatrix} \quad (57)$$

which consists of $A_{m'}$ and $m - rm' < r$ columns of zero. Multiplying $A_{m'}$ with F and y is equivalent to truncating F and y to their first rm' functions resp. outputs and multiplying those with $A_{m'}$. Lemma 4 therefore yields that the map

$$G := B \cdot F : k^n \longrightarrow \bar{k}^{m'} \quad (58)$$

$$x \longmapsto B \cdot F(x) \quad (59)$$

is polynomial of degree $d = \deg F$ over \bar{k} , and that $y' := B \cdot y$ is uniformly distributed in $\bar{k}^{m'}$ if $y \leftarrow k^m$.

\mathcal{R} therefore computes $y' = B \cdot y$ and polynomials representing $G = B \cdot F$ and submits both to \mathcal{A} . We can now consider two cases:

1. If y equals $F(x)$ for some $x \in \{0, 1\}^n$, then $y' = B \cdot y$ equals $G(x) = B \cdot F(x)$. In this case, \mathcal{A} will always output 0.

2. If $y \leftarrow k^m$ has been sampled uniformly at random, then $y' = B \cdot y$ is distributed uniformly in $\overline{k}^{m'}$. In this case, \mathcal{A} will output 1 with probability

$$\geq 1 - O\left(\left(\frac{n^d}{m'}\right)^{\frac{1}{d-1}} \cdot \frac{1}{\#\overline{k}}\right) \geq 1 - O\left(\left(\frac{r \cdot n^d}{m}\right)^{\frac{1}{d-1}} \cdot \frac{1}{n}\right) \quad (60)$$

$$\geq 1 - O\left(\left(\frac{r \cdot n}{m}\right)^{\frac{1}{d-1}}\right) \quad (61)$$

It follows that the advantage of \mathcal{R} is lower-bounded by

$$\text{adv}_F(\mathcal{R}) \geq 1 - O\left(\left(\frac{r \cdot n}{m}\right)^{\frac{1}{d-1}}\right) \quad (62)$$

and hence high.

Further, the time complexity of \mathcal{A} is upper-bounded by

$$\leq n^{O\left(\left(\frac{n^d}{m'}\right)^{1/(d-1)}\right)} \leq n^{O\left(\left(\frac{r \cdot n^d}{m}\right)^{1/(d-1)}\right)}. \quad (63)$$

Hence, \mathcal{R} 's time complexity is bounded by $\leq n^{O\left(\left(\frac{r \cdot n^d}{m}\right)^{1/(d-1)}\right)}$, too.

Remark 3. If one does not want to use extension fields \overline{k} , one can modify the algorithm from Theorem 5 s.t. one receives an algebraic algorithm over the base field k that has the same runtime and advantage bounds as in Theorem 2.

Given a binary PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $m \geq n^{1+e}$, the idea is – when running the algorithm from Theorem 5 – to compute a basis for $\ker \phi_D$ where

$$\phi_D : \mathbb{Z}_2[Y_1, \dots, Y_m] \longrightarrow \mathbb{Z}_2[X_1, \dots, X_n] \quad (64)$$

$$Y_i \longmapsto f_i \quad (65)$$

for $D \in O\left(\left(\log n\right)^{1/(\deg F-1)} \cdot n^{1-e/(\deg F-1)}\right)$ large enough.

When one needs to decide if a point $y \in \{0, 1\}^m$ lies in the image of F , one evaluates *all* basis polynomials of $\ker \phi_D$ on y : if one polynomial in $\ker \phi_D$ does not vanish on y , then we know that y cannot lie in the image of F . Otherwise, if the whole space $\ker \phi_D$ vanishes on y , then we know that with high probability y must lie in the image of F .

The reason for this is that the algebraic relationship $h \in \mathbb{F}_N[Y'_1, \dots, Y'_{m'}]$, $N = 2^{\lceil \log n \rceil}$, that is yield in the algorithm of Theorem 5 after we went from $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ to $G : \{0, 1\}^n \rightarrow \mathbb{F}_N^{m'}$ can be interpreted as $\lceil \log n \rceil$ polynomials in $\mathbb{Z}_2[Y_1, \dots, Y_m]$ of degree $O\left(\left(\log n\right)^{1/(\deg F-1)} \cdot n^{1-e/(\deg F-1)}\right)$.

In fact, the matrix $A_{m'}$ from Lemma 4 maps the variables Y_1, \dots, Y_m (which represents binary values) to the variables $Y'_1, \dots, Y'_{m'}$, $m' = \lfloor m / \lceil \log n \rceil \rfloor$, (which represent extension field values) by

$$Y'_i = Y_{(i-1) \cdot \lceil \log n \rceil + 1} + \zeta \cdot Y_{(i-1) \cdot \lceil \log n \rceil + 2} + \dots + \zeta^{\lceil \log n \rceil - 1} \cdot Y_{i \cdot \lceil \log n \rceil}. \quad (66)$$

Further, $h \in \mathbb{F}_N[Y'_1, \dots, Y'_{m'}]$ can be written as

$$h(Y') = \sum_{I \in \mathbb{N}_0^{m'}, |I| \leq D} (c_{I,1} + \zeta \cdot c_{I,2} + \dots + \zeta^{\lceil \log n \rceil - 1} \cdot c_{I, \lceil \log n \rceil}) \cdot Y'^I \quad (67)$$

for coefficients $c_{I,i} \in \mathbb{Z}_2$. By substituting the Y' -variables in h by Y -variables according to Eq. (66) and by sorting the terms in h by powers of ζ , we see that h can be written as

$$h(Y') = u_1(Y) + \zeta \cdot u_2(Y) + \dots + \zeta^{\lceil \log n \rceil - 1} \cdot u_{\lceil \log n \rceil}(Y) \quad (68)$$

for polynomials $u_1, \dots, u_{\lceil \log n \rceil} \in \mathbb{Z}_2[Y_1, \dots, Y_m]$. The degree of the polynomials $u_1, \dots, u_{\lceil \log n \rceil}$ is bounded by the degree of h .

It follows that checking one polynomial over the extension field \mathbb{F}_N is equivalent to checking $\lceil \log n \rceil$ polynomials over the base field \mathbb{Z}_2 . Hence, if $y \leftarrow \{0, 1\}^m$ is truly random, with probability $1 - o(1)$ one polynomial in $\ker \phi_D$ will not vanish on y for D large enough.

4 Solving the Polynomial Search Problem in the Average Case

We will describe here a subexponential algorithm that – given a polynomial map $F : k^n \rightarrow k^{n^{1+e}}$ and a function value $y = F(x)$ for any input x – will extract a pre-image x' s.t. $F(x') = y$ with high probability in the average case. Unfortunately, this algorithm will not be successful in the worst case. Its success probability will depend on the randomness of F (however it is independent of the distribution of x).

We will first define the set of suitable distributions of polynomial maps F , for which our algorithm does have a high advantage:

Definition 6. Let \mathcal{D} be a distribution with support in $(k[X_1, \dots, X_n]^{\leq d})^m$, this means \mathcal{D} outputs m polynomials over n variables of degree $\leq d$.

We say that \mathcal{D} has a **linear core**, if there is a distribution \mathcal{D}' s.t. \mathcal{D} is identically distributed as the output of the following procedure:

$$\text{Draw } F'(X) \leftarrow \mathcal{D}' \text{ and } R \leftarrow k^{m \times n}. \quad (69)$$

$$\text{Set } F(X) := F'(X) + R \cdot X. \quad (70)$$

$$\text{Output } F(X). \quad (71)$$

Here, $R \cdot X$ denotes the linear function that maps $x \in k^n$ to Rx .

Now, we can state our formal result:

Theorem 7. Let k be a field with q elements. Let \mathcal{D} be a distribution over $(k[X_1, \dots, X_n]^{\leq d})^m$ with linear core, where $m \in \omega(r \cdot n)$ for $r = \max\{1, \frac{\log n}{\log q}\}$.

There is an algorithm \mathcal{B} s.t. for each $x \in k^n$ and $F \leftarrow \mathcal{D}$, $\mathcal{B}(F, F(x))$ will output x with high probability. Formally,

$$\forall x \in k^n : \Pr_{F \leftarrow \mathcal{D}} [x \leftarrow \mathcal{B}(F, F(x))] \geq 1 - o(1). \quad (72)$$

The time-complexity of \mathcal{B} is always bounded by

$$\leq q \cdot n^{O(n \cdot (r \cdot n/m)^{1/(d-1)})}. \quad (73)$$

For $m \geq n^{1+e}$ and $q = 2$, we have a runtime of

$$\leq 2^{O(\log(n)^{1/(d-1)} \cdot n^{1-e/(d-1)})}. \quad (74)$$

The idea for our algorithm will essentially follow the typical search-to-decision reduction paradigm that is already known from learning with errors [Reg05]: that is, for each input value x_i , $i \in [n]$, guess a potential candidate $z \in k$. Then, ask the algorithm \mathcal{A} from Theorem 6 if the polynomial equation system

$$F(X_1, \dots, X_{i-1}, z, X_{i+1}, \dots, X_n) = F(x) \quad (75)$$

is solvable. Since the distribution of F has a linear core, it seems reasonable to assume that $F(x)$ will have enough entropy for a high distinguishing advantage of \mathcal{A} .

Unfortunately, this direct approach makes it hard to prove that we indeed have a high distinguishing advantage. Hence, in the following, we will not give $F(X_1, \dots, X_{i-1}, z, X_{i+1}, \dots, X_n)$ and $F(x)$ to \mathcal{A} . Instead, we will give the “sanitized” polynomial

$$G_{i,z}(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n) \quad (76)$$

$$:= F(X_1, \dots, X_{i-1}, z, X_{i+1}, \dots, X_n) - F(0, \dots, 0, z, 0, \dots, 0) \quad (77)$$

and the target value

$$y_{i,z} := F(x) - F(0, \dots, 0, z, 0, \dots, 0) \quad (78)$$

to \mathcal{A} . In this case, we can prove that $y_{i,z}$ is uniformly random and independent of $G_{i,z}$.

Proof (Theorem 7). Let k be a field of size q and let $m \in \omega(\frac{\log n}{\log q} \cdot n)$. Let \bar{k} be an extension field \bar{k} of k of degree $r := 1 + \lceil 2 \frac{\log n}{\log q} \rceil$. It follows

$$\#\bar{k} = q^r = q^{1 + \lceil 2 \frac{\log n}{\log q} \rceil} \geq q \cdot (q^{\frac{\log n}{\log q}})^2 = q \cdot n^2. \quad (79)$$

Without loss of generality, we can assume that r divides m and set $m' := m/r$.

Now, let $A_{m'} \in \bar{k}^{m' \times m}$ be the matrix from Lemma 4 and let \mathcal{A} be the algorithm from Theorem 5 for polynomial maps $G : k^{n-1} \rightarrow \bar{k}^{m'}$. Remember that \mathcal{A} works by computing an algebraic relationship h of degree $D \in O(n \cdot (n/m')^{1/(d-1)}) \in o(n)$ for G and checking if h vanishes on a potential image point of G .

Finally, fix any $x \in k^n$ and let \mathcal{D} be a distribution over $(k[X_1, \dots, X_n]^{\leq d})^m$ with linear core.

On input $F \leftarrow \mathcal{D}$ and $y := F(x)$, our algorithm \mathcal{B} proceeds as follows:

1. \mathcal{B} computes $r = 1 + \lceil 2 \frac{\log n}{\log q} \rceil$, $m' = m/r$ and the extension field \bar{k} of k of degree r .

2. \mathcal{B} applies $A_{m'}$ to F and y . I.e., it computes

$$G(X) := A_{m'} \cdot F(X) \in (\bar{k}[X_1, \dots, X_n]^{\leq d})^{m'} \quad (80)$$

$$y' := A_{m'} \cdot y = G(x) \in \bar{k}^m. \quad (81)$$

Note, that G is distributed with linear core, since F is distributed with linear core and $A_{m'}$ gives an isomorphism $k^{m \times n} \rightarrow \bar{k}^{m' \times n}$.

3. For each $i \in [n]$ and $z \in k$, \mathcal{B} computes

$$G_{i,z}(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n) \quad (82)$$

$$:= G(X_1, \dots, X_{i-1}, z, X_{i+1}, \dots, X_n) - G(0, \dots, 0, z, 0, \dots, 0) \quad (83)$$

and

$$y_{i,z} := y' - G(0, \dots, 0, z, 0, \dots, 0). \quad (84)$$

4. For each $i \in [n]$ and $z \in k$, \mathcal{B} uses algorithm \mathcal{A} from Theorem 5 to compute an algebraic relationship $h_{i,z} \in \bar{k}[Y_1, \dots, Y_{m'}]$ for $G_{i,z}$. We have

$$h_{i,z}(Y) \neq 0, \quad (85)$$

$$h_{i,z} \circ G_{i,z}(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n) = 0, \quad (86)$$

$$\deg h_{i,z} \leq D \in o(n). \quad (87)$$

5. For each $i \in [n]$, \mathcal{B} does the following:

\mathcal{B} enumerates all $z \in k$ and checks if $h_{i,z}(y_{i,z}) = 0$. If there is exactly one $z \in k$ s.t. $h_{i,z}(y_{i,z}) = 0$, then \mathcal{B} sets $x'_i := z$. Otherwise, \mathcal{B} aborts.

6. \mathcal{B} outputs $x' := (x'_1, \dots, x'_n) \in k^n$ as solution.

\mathcal{B} 's runtime is dominated by step 4 and 5. We can upper-bound it by

$$q \cdot n \cdot n^{O(D)} \leq q \cdot n^{O(n \cdot (n/m')^{1/(d-1)})}. \quad (88)$$

To bound \mathcal{B} 's advantage, we fix one $i \in [n]$ and first bound the probability that $x'_i = x_i$. Without loss of generality, we can assume $i = 1$.

First note that if $z = x_1$ in step 5, then $G_{1,z}(X) = y_{1,z}$ is always solvable, and $h_{1,z}$ must vanish at $y_{1,z}$. Hence, we only need to show that $h_{1,z}(y_{1,z})$ is non-zero with high probability whenever $z \neq x_1$. Therefore, let $z \in k \setminus \{x_1\}$ and let $G(X) = G'(X) + R \cdot X$, where $R \leftarrow \bar{k}^{m' \times n}$ is distributed uniformly and independently of G' at random. Denote by $R_1, \dots, R_n \in \bar{k}^m$ the columns of R . We have

$$G_{1,z}(X_2, \dots, X_n) \quad (89)$$

$$= G(z, X_2, \dots, X_n) - G(z, 0, \dots, 0) \quad (90)$$

$$= G'(z, X_2, \dots, X_n) - G'(z, 0, \dots, 0) \quad (91)$$

$$+ R \cdot (z, X_2, \dots, X_n) - R \cdot (z, 0, \dots, 0) \quad (92)$$

$$= G'(z, X_2, \dots, X_n) - G'(z, 0, \dots, 0) + R_2 \cdot X_2 + \dots + R_n \cdot X_n \quad (93)$$

and

$$y_{1,z} = y' - G(z, 0, \dots, 0) \quad (94)$$

$$= G(x_1, \dots, x_n) - G(z, 0, \dots, 0) \quad (95)$$

$$= G'(x_1, \dots, x_n) - G'(z, 0, \dots, 0) + R \cdot x - R \cdot (z, 0, \dots, 0) \quad (96)$$

$$= G'(x_1, \dots, x_n) - G'(z, 0, \dots, 0) \quad (97)$$

$$+ R_2 \cdot x_2 + \dots + R_n \cdot x_n + R_1 \cdot (x_1 - z). \quad (98)$$

Now, $h_{i,z}$ depends only on the map $G_{1,z}$, whose distribution is independent of the random vector R_1 . On the other hand, $y_{1,z}$ gets perturbed by the uniformly random vector $R_1 \cdot (x_1 - z)$. Therefore, $y_{1,z}$ remains uniformly random in \bar{k}^m even if we condition its distribution on $G_{1,z}$. Because of the Schwartz-Zippel Lemma 3, we have for each $x \in k^n$ and $z \in k \setminus \{x_1\}$

$$\Pr[h_{1,z}(y_{1,z}) = 0] \leq \frac{\deg h}{\#\bar{k}} \leq \frac{D}{q \cdot n^2} \in o\left(\frac{1}{q \cdot n}\right) \quad (99)$$

where the probability is taken over the randomness of F . Taking a union bound over all $z \in k \setminus \{x_1\}$, we get

$$\Pr[\exists z \in k \setminus \{x_1\} : h_{1,z}(y_{1,z}) = 0] \leq q \cdot \frac{D}{q \cdot n^2} \in o\left(\frac{1}{n}\right). \quad (100)$$

It follows that for each $i \in [n]$, the probability that \mathcal{B} does not recover x_i is lower-bounded by $o(1/n)$. Again, taking a union bound over all $i = 1, \dots, n$, the probability that there is at least one i s.t. \mathcal{B} does not successfully recover x_i is bounded by $n \cdot o(1/n) = o(1)$. It follows that \mathcal{B} 's success probability is at least

$$\geq 1 - o(1). \quad (101)$$

Remark 4. In [Üna23b], we discussed implications for Macaulay matrix-based algorithms that arise from the algorithm in Theorem 5. Analogously, we can apply the results for the algorithm from Theorem 7 to Macaulay matrices and distil new insights.

For this end, let k be a field with q elements, let \bar{k} be a field of extension degree r , draw $F \leftarrow \mathcal{D}$ and set $G := A_{m'} \cdot F$ for $m' = m/r$. Fix some $x \in k^n$.

For $\ell \in \mathbb{N}$, set

$$M_\ell(G) := M_\ell(g_1, \dots, g_{m'}) \quad (102)$$

$$:= \text{span}_{\bar{k}} \{g_i(X) \cdot u(X) \mid i \in [m], u \in \bar{k}[X], \deg g_i + \deg u \leq \ell\}. \quad (103)$$

Note, that we have for ℓ large enough

$$M_\ell(G) = M_{\ell-1}(G) + \sum_{i=1}^n X_i \cdot M_{\ell-1}(G) \quad (104)$$

and that $M_\ell(G)$ is the image of the Macaulay matrix up to degree ℓ for the equation system $G(X) = 0$.

Assume that the algorithm \mathcal{B} from Theorem 7 can extract x from F and $F(x)$. This implies that the system $G(X) - G(x) = 0$ with hardwired $X_i = z$ can be refuted by an algebraic relationship of degree $D \in O(n \cdot (n/m')^{1/(d-1)})$ for each $i \in [n]$ and $z \neq x_i$. From our reasoning in [Üna23b], it follows that $M_{dD}(G(X) - G(x), X_i - z)$ must contain 1. Hence, for each $i \in [n]$ and $z \neq x_i$, the space $M_{dD}(G(X) - G(x))$ contains the element

$$1 - (X_i - z)t_{i,z}(X) \quad (105)$$

for some $t_{i,z} \in \bar{k}[X]^{dD-1}$. Let $b(X_i)$ be the field equation of k i.e.

$$b(X_i) := \prod_{a \in k} (X_i - a). \quad (106)$$

Now, the space

$$M_{dD+q-1}(G(X) - G(x), b(X_i)) \quad (107)$$

contains the element

$$b(X_i) \cdot t_{i,z}(X) + (1 - (X_i - z)t_{i,z}(X)) \cdot \frac{b(X_i)}{X_i - z} = \frac{b(X_i)}{X_i - z} \quad (108)$$

for $z \neq x_i$. Subsequently, we can remove each factor $X_i - z$ with $z \neq x_i$ from $b(X_i)$ and are left with $X_i - x_i$. It follows for each $i \in [n]$

$$X_i - x_i \in M_{dD+q-1}(G(X) - G(x), b(X_i)). \quad (109)$$

I.e., with high probability over the randomness of $F \leftarrow \mathcal{D}$, the row space

$$M_{dD+q-1}(G(X) - G(x), b_1, \dots, b_n) \quad (110)$$

of the Macaulay matrix of degree $dD + q - 1$ will contain the Groebner basis $X_1 - x_1, \dots, X_n - x_n$ of the ideal $(G(X) - G(x), b_1, \dots, b_n)$.

Now, the space $M_{dD+q-1}(G(X) - G(x), b_1, \dots, b_n)$ is contained in the tensor product of \bar{k} and the row space of the k -valued Macaulay matrix of $(F(X) - F(x), b_1, \dots, b_n)$ up to degree $dD + q - 1$. Hence, the Macaulay matrix of $(F(X) - F(x), b_1, \dots, b_n)$ of degree $dD + q - 1$ over k will contain the k -valued polynomials $X_1 - x_1, \dots, X_n - x_n$, too.

5 On the Security of Polynomial and Local PRGs

We try to derive here some insights on parameter choices for PRGs of constant locality and constant degree. By the attacks of this work, we know that the number of *security bits*⁴ of a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$ provably lies in

$$O(\log(n)^{\deg F / (\deg F - 1)} \cdot n^{1-e / (\deg F - 1)}). \quad (111)$$

⁴ By “number of security bits” we mean the logarithm of the average of the concrete number of bit operations a “strong” computational machine needs to run the fastest attacks on PRGs of constant degree.

For simplicity, we will pretend that the number of security bits of F is upper-bounded by

$$n^{1-e/(\deg F-1)}. \quad (112)$$

From a theoretical point of view, this is a very pessimistic view, since we are ignoring the factor $\log(n)^{\deg F/(\deg F-1)}$ and all potentially hidden constants in Eq. (111). However, if we compare the simple theoretical security estimations here with the experimental estimations of Couteau, Dupin, Méaux, Rossi & Rotella [CDM⁺18] and Yang, Guo, Johansson & Lentmaier [YGJL22], we see that their results are far more pessimistic. Hence, we think that the term in Eq. (112) is a justified upper-bound for the number of security bits of F .

While n denotes the number of bits of the seed of F , we will use λ to denote the desired number of security bits. In this view, n gives the number of bits we are ready to invest, while λ denotes the number of bits we are going to get in the end. Optimally, n and λ would be close, however, in our case, it is more suitable to relate them polynomially: by $\delta > 1$ we will denote the *security leverage* and relate λ and n by

$$\lambda^\delta = n. \quad (113)$$

I.e., to achieve a security of λ we will need to invest polynomially more bits. Now, we upper-bound λ by the term in Eq. (112). This leads to the formula

$$1 \leq \delta \cdot \left(1 - \frac{e}{\deg F - 1}\right). \quad (114)$$

By rearranging everything, we get the following lower bound for the degree of F

$$\deg F \geq \frac{\delta}{\delta - 1} \cdot e + 1. \quad (115)$$

Eq. (115) tells us how high we need to set the algebraic degree of F if we want to maintain a security leverage of δ and a stretch of e . Note that for $\delta \rightarrow \infty$, the RHS of Eq. (115) converges against $e + 1$. In fact, $\deg F > e + 1$ is the trivial bound that we always have for $\deg F$ and e , since for $\deg F \leq e + 1$ the relinearization attack on F becomes efficient. In other words, the higher our security leverage δ becomes, the less important subexponential attacks on F get. Efficient attacks like relinearization and correlation attacks, however, will always stay relevant and give hard bounds for degree and locality of F .

If F is additionally supposed to have constant locality, we need to take the correlation-based attack from Section 1.3 into account. If c denotes the correlation of F , then it must hold $e < c/2$, since otherwise the correlation-based attack becomes applicable. Since we have $c \leq \text{loc } F - \deg F - 1$, we get $e < (\text{loc } F - \deg F - 1)/2$ and have therefore the minimum bound

$$\text{loc } F > \deg F + 2e + 1 \geq \left(\frac{\delta}{\delta - 1} + 2\right) \cdot e + 2 \quad (116)$$

for the locality of F .

Eq. (115) and Eq. (116) now relate the degree, the locality, the stretch and the security leverage of F . In particular, they tell us lower bounds for the degree and locality of F for given δ and e . For given values of δ , we can now list the concrete bounds on $\deg F$ and $\text{loc } F$ in Table 1⁵.

| δ | $\deg F$ | $\text{loc } F$ |
|----------|--------------------------|-----------------------|
| 1.1 | $\geq 11 \cdot e + 1$ | $> 13 \cdot e + 2$ |
| 1.2 | $\geq 6 \cdot e + 1$ | $> 8 \cdot e + 2$ |
| 1.3 | $\geq 13/3 \cdot e + 1$ | $> 19/3 \cdot e + 2$ |
| 1.4 | $\geq 7/2 \cdot e + 1$ | $> 11/2 \cdot e + 2$ |
| 1.5 | $\geq 3 \cdot e + 1$ | $> 5 \cdot e + 2$ |
| 1.6 | $\geq 8/3 \cdot e + 1$ | $> 14/3 \cdot e + 2$ |
| 1.7 | $\geq 17/7 \cdot e + 1$ | $> 31/7 \cdot e + 2$ |
| 1.8 | $\geq 9/4 \cdot e + 1$ | $> 17/4 \cdot e + 2$ |
| 1.9 | $\geq 19/9 \cdot e + 1$ | $> 37/9 \cdot e + 2$ |
| 2.0 | $\geq 2 \cdot e + 1$ | $> 4 \cdot e + 2$ |
| 2.1 | $\geq 21/11 \cdot e + 1$ | $> 43/11 \cdot e + 2$ |
| 2.2 | $\geq 11/6 \cdot e + 1$ | $> 23/6 \cdot e + 2$ |
| 2.3 | $\geq 23/13 \cdot e + 1$ | $> 49/13 \cdot e + 2$ |
| 2.4 | $\geq 12/7 \cdot e + 1$ | $> 26/7 \cdot e + 2$ |
| 2.5 | $\geq 5/3 \cdot e + 1$ | $> 11/3 \cdot e + 2$ |
| 2.6 | $\geq 13/8 \cdot e + 1$ | $> 29/8 \cdot e + 2$ |
| 2.7 | $\geq 27/17 \cdot e + 1$ | $> 61/17 \cdot e + 2$ |
| 2.8 | $\geq 14/9 \cdot e + 1$ | $> 32/9 \cdot e + 2$ |
| 2.9 | $\geq 29/19 \cdot e + 1$ | $> 67/19 \cdot e + 2$ |
| 3.0 | $\geq 3/2 \cdot e + 1$ | $> 7/2 \cdot e + 2$ |
| 3.5 | $\geq 7/5 \cdot e + 1$ | $> 17/5 \cdot e + 2$ |
| 4.0 | $\geq 4/3 \cdot e + 1$ | $> 10/3 \cdot e + 2$ |
| 4.5 | $\geq 9/7 \cdot e + 1$ | $> 23/7 \cdot e + 2$ |
| 5.0 | $\geq 5/4 \cdot e + 1$ | $> 13/4 \cdot e + 2$ |
| ∞ | $> e + 1$ | $> 3 \cdot e + 2$ |

Table 1. This table shows for some values of the security leverage $\delta = \log(n)/\log(\lambda)$ corresponding lower bounds for the degree and locality of the PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+e}}$. Each row can be read as: “if we want a security of λ bits while our seeds contain $n = \lambda^\delta$ bits, then the degree of the PRG needs to be at least $\frac{\delta}{\delta-1} \cdot e + 1$ and the locality of F needs to be larger than $\left(\frac{\delta}{\delta-1} + 2\right) \cdot e + 2$.”

If we want, for example, a PRG $F : \{0, 1\}^n \rightarrow \{0, 1\}^{n^2}$ of quadratic stretch, i.e. $e = 1$, then it makes sense to pick a security leverage that is a by an epsilon larger than 2. If $\delta > 2$, the degree bound $\deg F \geq \frac{\delta}{\delta-1} \cdot e + 1$ is already implied

⁵ In the initial version of this text, Table 1 contained wrong values (the correct values are a bit more pessimistic). The reason is that I initially computed the values in Table 1 by hand. This time, I used a python script to generate the table. You can find it here: <https://github.com/Semigroup/algebraic-prg-attack-scripts>

by the trivial bound $\deg F > e + 1 = 2$ that stems from relinearization, since $\frac{\delta}{\delta-1} \cdot e + 1 < 2 \cdot e + 1 = 3$ and $\deg F$ needs to be an integer. On the other hand, for the locality bound $\text{loc } F > \left(\frac{\delta}{\delta-1} + 2\right) \cdot e + 2$ we have $\left(\frac{\delta}{\delta-1} + 2\right) \cdot e + 2 < 4 \cdot e + 2 = 6$. I.e., we can choose $\text{loc } F = 6$, and indeed this bound is also already implied by the trivial bound $\text{loc } F > 3 \cdot e + 2 = 5$.

However, note that a typical security of $\lambda = 128$ security bits would – under a security leverage of $\delta > 2$ – imply a seed of more than 128^2 bits, i.e. 2 kilobytes, and an output size of F of > 4 megabytes.

6 Reductions of Learning Parity with Noise

In the following, we will give some efficient reductions for the learning parity with noise (LPN) problem over fields of different sizes. Our reductions will show that LPN with low-rate noise over fields of constant size is (up to logarithmic factors) equivalent to LPN with similar noise rate over fields of polynomial size.

First, we will define the distribution of LPN samples:

Definition 7. *Let k be a finite field and let $\tau \in [0, 1]$ be a real number. The distribution $\text{Ber}_k(\tau)$ is defined to output zero with probability $1 - \tau$ and a uniformly random element of k^\times with probability τ . I.e.,*

$$\Pr_{e \leftarrow \text{Ber}_k(\tau)}[e = 0] = 1 - \tau, \quad (117)$$

$$\forall x \in k^\times : \Pr_{e \leftarrow \text{Ber}_k(\tau)}[e = x | e \neq 0] = \frac{1}{\#k - 1}. \quad (118)$$

$$(119)$$

Let $n \in \mathbb{N}$ and $s \in k^n$. The distribution $\text{LPN}_k(s, \tau)$ is defined by the following sample procedure:

$$\text{Draw } b \leftarrow k^n. \quad (120)$$

$$\text{Draw } e \leftarrow \text{Ber}_k(\tau). \quad (121)$$

$$\text{Set } c := b^T \cdot s + e. \quad (122)$$

$$\text{Output } (c, b). \quad (123)$$

Where $(b_1, \dots, b_n)^T \cdot (s_1, \dots, s_n) = b_1 s_1 + \dots + b_n s_n$ denotes the scalar product of k^n .

For $m \in \mathbb{N}$, we will denote by $\text{LPN}_k(s, \tau)^m$ a tuple of m independent samples of $\text{LPN}_k(s, \tau)$.

For a field k and a field extension \bar{k} of extension degree $r = [\bar{k} : k]$, we will give two reductions of LPN here: the first reduction converts samples of $\text{LPN}_k(s, \tau)^r$ to $\text{LPN}_{\bar{k}}(s, \tau')$ for some $\tau' \leq r\tau$ and $s \in k^n$. The second reduction converts samples of $\text{LPN}_{\bar{k}}(\bar{s}, \tau)$ to $\text{LPN}_k(s, \tau')$ for $\tau' \leq \tau$ and $\bar{s} \in \bar{k}^n$ and $s \in k^{rn}$ s.t. \bar{s}, s are in an efficiently computable one-to-one correspondence.

Remark 5. We will assume that the following algorithms know the field extension $k \subset \bar{k}$ and can compute arithmetic operations efficiently in both fields. Further, we will implicitly assume that they can parse \bar{k} as k -vector space (according to some unspecified basis). This condition is naturally satisfied if elements of \bar{k} are stored as tuples of elements of k .

Proposition 1. *Let $k \subset \bar{k}$ be an extension of finite fields of degree $r = [\bar{k} : k]$. Let $\tau = \tau(n) \in [0, 1]$.*

There is an efficient reduction \mathcal{R} s.t. we have the following:

1. *For each $s \in k^n$, \mathcal{R} will, on input $((c_1, b_1), \dots, (c_r, b_r)) \leftarrow \text{LPN}_k(s, \tau)$, output an element (\bar{c}, \bar{b}) distributed according to $\text{LPN}_{\bar{k}}(s, \tau')$, where*

$$\tau' = 1 - (1 - \tau)^r \in [\tau, r\tau]. \quad (124)$$

2. *If $((c_1, b_1), \dots, (c_r, b_r))$ are distributed uniformly at random in $k^{r(n+1)}$, then the output of \mathcal{R} will be distributed uniformly at random in \bar{k}^{n+1} .*

Proof. Let $(c_1, b_1), \dots, (c_r, b_r)$ be the input of \mathcal{R} . We set

$$c := \begin{pmatrix} c_1 \\ \vdots \\ c_r \end{pmatrix} \in k^r \text{ and } B := \begin{pmatrix} b_1^T \\ \vdots \\ b_r^T \end{pmatrix} \in k^{r \times n}. \quad (125)$$

\mathcal{R} samples a random basis z_1, \dots, z_r of \bar{k} as k vector space and sets

$$z := \begin{pmatrix} z_1 \\ \vdots \\ z_r \end{pmatrix} \in \bar{k}^r. \quad (126)$$

\mathcal{R} computes and outputs

$$\bar{c} := z^T \cdot c \text{ and } \bar{b} := z^T \cdot B. \quad (127)$$

Since z is a basis of \bar{k} as k , it gives a bijection

$$k^{r \times (n+1)} \longrightarrow \bar{k}^{n+1} \quad (128)$$

$$(c, B) \longmapsto (z^T c, z^T B). \quad (129)$$

Hence, (\bar{c}, \bar{b}) are uniformly distributed iff $(c_1, b_1), \dots, (c_r, b_r)$ are.

Now, sample $(c_1, b_1), \dots, (c_r, b_r) \leftarrow \text{LPN}_k(s, \tau)$ for some $s \in k^n$ and $\tau \in [0, 1]$. We can rewrite c as

$$c = Bs + e \quad (130)$$

for $e \leftarrow \text{Ber}_k(\tau)^r$. $\bar{c} = z^T c$ is then equal to $z^T Bs + z^T e$. It is clear that $z^T B$ is distributed uniformly, we need to show that $\bar{e} := z^T e = z_1 e_1 + \dots + z_r e_r$ is distributed according to $\text{Ber}_{\bar{k}}(\tau')$.

We can distinguish two cases:

1. Each noise term e_1, \dots, e_r is zero. This will happen exactly with probability $(1 - \tau)^r = 1 - \tau'$. In this case, $\bar{e} = z^T e$ is zero, too.
2. Otherwise, the vector $e = (e_1, \dots, e_r)$ is non-zero. Since z is a random basis of \bar{k} as k vector space, $\bar{e} = z^T e$ will be a uniformly random element of \bar{k}^\times . This element will be independent of $z^T B$, since B is distributed uniformly at random in $k^{r \times n}$.

It follows that \bar{e} is distributed according to $\text{Ber}_{\bar{k}}(\tau')$ for $\tau' = 1 - (1 - \tau)^r$. Hence, (\bar{c}, \bar{b}) is distributed according to $\text{LPN}_{\bar{k}}(s, \tau')$.

To prove the correctness of the second reduction, we need the following two lemmata:

Lemma 5. *Let $n \in \mathbb{N}$ and let $k \subset \bar{k}$ be an extension of degree $r = [\bar{k} : k]$. The matrix A_n from Lemma 4 has an inverse mapping $A_n^{-1} : \bar{k}^n \rightarrow k^{rn}$ that is k -linear s.t. we have for each $\bar{x} \in \bar{k}^n$ and $x \in k^{rn}$*

$$A_n^{-1}(A_n \cdot x) = x \text{ and } A_n \cdot A_n^{-1}(\bar{x}) = \bar{x}. \quad (131)$$

A_n^{-1} can be efficiently evaluated.

Proof. Let $\zeta \in \bar{k}$ be the generator that is used by A_n . Since A_n implies an isomorphism $k^{rn} \rightarrow \bar{k}^n$, the map $A_n^{-1} : \bar{k}^n \rightarrow k^{rn}$ must be bijective. Since A_n is k -linear and its domain⁶ and codomain are k -vector spaces, A_n^{-1} is k -linear, too.

Computing $A_n^{-1}(\bar{x})$ is equivalent to solving a k -linear equation system of shape $rn \times rn$. Hence, $A_n^{-1}(\bar{x})$ can be efficiently evaluated.

For $m \in \mathbb{N}$, we extend A_n^{-1} to a map $\bar{k}^{n \times m} \rightarrow k^{rn \times m}$ by applying it column-wise, i.e.

$$A_n^{-1}(u_1 \dots u_m) := (A_n^{-1}(u_1) \dots A_n^{-1}(u_m)) \quad (132)$$

for column vectors $u_1, \dots, u_m \in \bar{k}^n$.

Lemma 6. *Let $v = (1, 0, \dots, 0) \in k^r$ be the first unit vector. Define*

$$\psi : \bar{k}^n \longrightarrow k^{rn} \quad (133)$$

$$\bar{x} \longmapsto v^T \cdot A_1^{-1}(\bar{x}^T \cdot A_n). \quad (134)$$

Then, ψ is a k -linear isomorphism.

Proof. Since ψ is a linear map of k -vector spaces of identical dimensions, it suffices to prove that $\ker \psi$ is trivial. For $\bar{x} = (\bar{x}_1, \dots, \bar{x}_n) \in \bar{k}^n$, $\psi(\bar{x})$ is given by

$$\psi(\bar{x}) = v^T \cdot A_1^{-1}(\bar{x}^T \cdot A_n) \quad (135)$$

⁶ This is the main difference to the typical gadget-matrix $G : \{0, 1\}^{\lceil \log p \rceil n} \rightarrow \mathbb{Z}_p^n$ from [MP12]. The subset $\{0, 1\}^{\lceil \log p \rceil n} \subset \mathbb{Z}_p^n$ is not closed under addition, hence $G^{-1} : \mathbb{Z}_p^n \rightarrow \{0, 1\}^{\lceil \log p \rceil n}$ cannot be linear.

$$= (v^T \cdot A_1^{-1}(\bar{x}_1 \cdot A_1) \dots v^T \cdot A_1^{-1}(\bar{x}_n \cdot A_1)). \quad (136)$$

Hence, it suffices to prove

$$v^T \cdot A_1^{-1}(\bar{x} \cdot A_1) = 0 \implies \bar{x} = 0 \quad (137)$$

for each $\bar{x} \in \bar{k}$. Now, $A_1^{-1}(\bar{x} \cdot A_1)$ is given by an $r \times r$ -matrix over k

$$Z := A_1^{-1}(\bar{x} \cdot A_1) \in k^{r \times r}. \quad (138)$$

Z is a matrix representation of the multiplication with \bar{x} on \bar{k} . In fact, we have for each $y \in \bar{k}$

$$Z \cdot A_1^{-1}(y) = A_1^{-1}(\bar{x} \cdot A_1) \cdot A_1^{-1}(y) = A_1^{-1}(\bar{x} \cdot A_1 \cdot A_1^{-1}(y)) = A_1^{-1}(\bar{x} \cdot y). \quad (139)$$

If $v^T Z$ is zero, then Z does not have full rank. This implies that the multiplication with \bar{x} is not invertible in \bar{k} . Since \bar{k} is a field, \bar{x} must be zero. This proves the implication in Eq. (137), hence ψ is injective.

Proposition 2. *Let $k \subset \bar{k}$ be an extension of finite fields of degree $r = [\bar{k} : k]$. Denote by $q = \#k$ the size of k and let $\tau = \tau(n) \in [0, 1]$.*

There is an efficient reduction \mathcal{R} s.t. we have the following:

1. *For each $\bar{s} \in \bar{k}^n$, \mathcal{R} will, on input $(\bar{c}, \bar{b}) \leftarrow \text{LPN}_{\bar{k}}(\bar{s}, \tau)$, output an element (c, b) distributed according to $\text{LPN}_k(A_n^{-1}(\bar{s}), \tau')$ for*

$$\tau' = \tau \cdot \frac{q^r - q^{r-1}}{q^r - 1}. \quad (140)$$

2. *If (\bar{c}, \bar{b}) is distributed uniformly at random in \bar{k}^{n+1} , then the output of \mathcal{R} will be distributed uniformly at random in k^{rn+1} .*

Proof. Let $(\bar{c}, \bar{b}) \in \bar{k}^{n+1}$. On input (\bar{c}, \bar{b}) , \mathcal{R} proceeds as follows:

1. \mathcal{R} computes the vector

$$b_1 := \psi(\bar{b}) := v^T \cdot A_1^{-1}(\bar{b}^T \cdot A_n) \in k^{rn}, \quad (141)$$

where v is the first unit vector, and

$$c := (c_1, \dots, c_r) := A_1^{-1}(\bar{c}) \in k^r. \quad (142)$$

2. \mathcal{R} outputs (c_1, b_1) .

Because of Lemma 6, $\psi(\bar{b})$ is uniformly distributed iff \bar{b} is uniformly distributed. If \bar{c} is uniformly distributed, then c_1 is, too, hence \mathcal{R} maps uniformly random samples to uniformly random samples.

Now, fix $\bar{s} \in \bar{k}^n$, draw $\bar{b} \leftarrow \bar{k}^n$ and $\bar{e} \leftarrow \text{Ber}_{\bar{k}}(\tau)$ and assume that we have

$$\bar{c} = \bar{b}^T \cdot \bar{s} + \bar{e}. \quad (143)$$

Now, we have

$$A_1^{-1}(\bar{c}) = A_1^{-1}(\bar{b}^T \cdot \bar{s} + \bar{e}) \quad (144)$$

$$= A_1^{-1}(\bar{b}^T \cdot \bar{s}) + A_1^{-1}(\bar{e}) \quad (145)$$

$$= A_1^{-1}(\bar{b}^T \cdot A_n \cdot A_n^{-1}(\bar{s})) + A_1^{-1}(\bar{e}) \quad (146)$$

$$= A_1^{-1}(\bar{b}^T \cdot A_n) \cdot A_n^{-1}(\bar{s}) + A_1^{-1}(\bar{e}). \quad (147)$$

Set $e := (e_1, \dots, e_r) := A_1^{-1}(\bar{e})$, then we have

$$c_1 = b_1 \cdot A_n^{-1}(s) + e_1. \quad (148)$$

We know that b_1 is uniformly distributed iff \bar{b} is uniformly distributed. It remains to prove that e_1 is Bernoulli-distributed if \bar{e} is. One can compute

$$\Pr[e_1 = 0 | e \neq 0] = \frac{\Pr[e_1 = 0, e \neq 0]}{\Pr[e \neq 0]} = \frac{\Pr[e_1 = 0, (e_2, \dots, e_r) \neq 0]}{\Pr[e \neq 0]} \quad (149)$$

$$= \frac{\frac{1}{q} \cdot \frac{q^{r-1} - 1}{q^{r-1}}}{\frac{q^r - 1}{q^r}} = \frac{q^{r-1} - 1}{q^r - 1}. \quad (150)$$

It follows

$$\Pr_{\bar{e} \leftarrow \text{Ber}_{\bar{k}}(\tau)}[e_1 = 0] = 1 - \tau + \tau \cdot \frac{q^{r-1} - 1}{q^r - 1} = 1 - \tau'. \quad (151)$$

Ergo, e_1 is distributed according to $\text{Ber}_k(\tau')$ if \bar{e} is distributed according to $\text{Ber}_{\bar{k}}(\tau)$.

Remark 6. Note that the reduction in Proposition 2 loses information. In fact, if (\bar{c}, \bar{b}) is uniformly random, then $c := A_1^{-1}(\bar{c})$ is uniformly random, however, the reduction only outputs the first coordinate of c and throws the remaining coordinates away.

For most applications, it makes more sense that the reduction outputs the whole vector c and the whole matrix $B := A_1^{-1}(\bar{b} \cdot A_n)$, even if B is not uniformly distributed.

Further, if (\bar{c}, \bar{b}) is an LPN sample, then (c, B) is of the form

$$c = B \cdot A_n^{-1}(\bar{s}) + A_1^{-1}(\bar{e}). \quad (152)$$

Now, $e = A_1^{-1}(\bar{e})$ is zero with high probability, but the entries of e are all correlated with each other. However, if the reduction converts a lot of independent LPN samples, then this correlation can be compensated.

References

- ABR12. Benny Applebaum, Andrej Bogdanov, and Alon Rosen. A dichotomy for local small-bias generators. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 600–617. Springer, Heidelberg, March 2012.

- ABW10. Benny Applebaum, Boaz Barak, and Avi Wigderson. Public-key cryptography from different assumptions. In Leonard J. Schulman, editor, *42nd ACM STOC*, pages 171–180. ACM Press, June 2010.
- ADI⁺17. Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. Secure arithmetic computation with constant computational overhead. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 223–254. Springer, Heidelberg, August 2017.
- AG11. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiří Sgall, editors, *Automata, Languages and Programming*, pages 403–415, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- AIK08. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. On pseudorandom generators with linear stretch in nc^0 . *Comput. Complex.*, 17(1):38–69, apr 2008.
- AK23. Benny Applebaum and Niv Konstantini. Actively secure arithmetic computation and vole with constant computational overhead. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 190–219, Cham, 2023. Springer Nature Switzerland.
- AL16. Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 1087–1100. ACM Press, June 2016.
- AL18. Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. *SIAM Journal on Computing*, 47(1):52–79, 2018.
- App12. Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 805–816. ACM Press, May 2012.
- App13. Benny Applebaum. Cryptographic hardness of random local functions-survey. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, page 599. Springer, Heidelberg, March 2013.
- BCG⁺19a. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 291–308. ACM Press, November 2019.
- BCG⁺19b. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 489–518. Springer, Heidelberg, August 2019.
- BCG⁺20. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators from ring-LPN. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 387–416. Springer, Heidelberg, August 2020.
- BCG⁺22. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Nicolas Resch, and Peter Scholl. Correlated pseudorandomness from expand-accumulate codes. In Yevgeniy Dodis and Thomas Shrimpton, editors,

- CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 603–633. Springer, Heidelberg, August 2022.
- BCGI18. Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 896–912. ACM Press, October 2018.
- BFKL94. Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 278–291. Springer, Heidelberg, August 1994.
- BKW19. Andreas Björklund, Petteri Kaski, and Ryan Williams. Solving Systems of Polynomial Equations over $GF(2)$ by a Parity-Counting Self-Reduction. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:13, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- BQ09. Andrej Bogdanov and Youming Qiao. On the security of goldreich’s one-way function. In Irit Dinur, Klaus Jansen, Joseph Naor, and José Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 392–405, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- Buc76. B. Buchberger. A theoretical basis for the reduction of polynomials to canonical forms. *SIGSAM Bull.*, 10(3):19–29, aug 1976.
- CCNY12. Chen-Mou Cheng, Tung Chou, Ruben Niederhagen, and Bo-Yin Yang. Solving quadratic equations with XL on parallel architectures. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 356–373. Springer, Heidelberg, September 2012.
- CDM⁺18. Geoffroy Couteau, Aurélien Dupin, Pierrick Méaux, Mélissa Rossi, and Yann Rotella. On the concrete security of Goldreich’s pseudorandom generator. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 96–124. Springer, Heidelberg, December 2018.
- CG21. Alessio Caminata and Elisa Gorla. Solving multivariate polynomial systems and an invariant from commutative algebra. In Jean Claude Bajard and Alev Topuzoğlu, editors, *Arithmetic of Finite Fields*, pages 3–36, Cham, 2021. Springer International Publishing.
- CG23. Alessio Caminata and Elisa Gorla. Solving degree, last fall degree, and related invariants. *J. Symb. Comput.*, 114(C):322–335, jan 2023.
- CKPS00. Nicolas Courtois, Alexander Klimov, Jacques Patarin, and Adi Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 392–407. Springer, Heidelberg, May 2000.
- CW04. Moses Charikar and Anthony Wirth. Maximizing quadratic programs: Extending Grothendieck’s inequality. In *45th FOCS*, pages 54–60. IEEE Computer Society Press, October 2004.
- DBM⁺08. Jintai Ding, Johannes Buchmann, Mohamed Saied Emam Mohamed, Wael Said Abd Elmageed Mohamed, and Ralf-Philipp Weinmann. Mutantxl. *SCC*, (TUD-CS-2009-0142):16–22, 01 2008.
- Din. Itai Dinur. *Improved Algorithms for Solving Polynomial Systems over $GF(2)$ by Multiple Parity-Counting*, pages 2550–2564.

- Din21. Itai Dinur. Cryptanalytic applications of the polynomial method for solving multivariate equation systems over $\text{gf}(2)$. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, pages 374–403, Cham, 2021. Springer International Publishing.
- DL78. Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978.
- Dub90. Thomas W. Dubé. The structure of polynomial ideals and gröbner bases. *SIAM Journal on Computing*, 19(4):750–773, 1990.
- Fau99. Jean-Charles Faugère. A new efficient algorithm for computing gröbner bases (f4). *Journal of Pure and Applied Algebra*, 139(1):61–88, 1999.
- Fau02. Jean Charles Faugère. A new efficient algorithm for computing gröbner bases without reduction to zero (f5). In *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation*, ISSAC '02, page 75–83, New York, NY, USA, 2002. Association for Computing Machinery.
- Gol11. Oded Goldreich. *Candidate One-Way Functions Based on Expander Graphs*, pages 76–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- GW95. Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, nov 1995.
- HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- IKOS08. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 433–442. ACM Press, May 2008.
- JLS21. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 60–73, New York, NY, USA, 2021. Association for Computing Machinery.
- JLS22. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from lpn over fp , dlin , and prgs in $\text{nc}0$. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 670–699, Cham, 2022. Springer International Publishing.
- Laz83. D. Lazard. Gröbner bases, gaussian elimination and resolution of systems of algebraic equations. In J. A. van Hulzen, editor, *Computer Algebra*, pages 146–156, Berlin, Heidelberg, 1983. Springer Berlin Heidelberg.
- LPT⁺. Daniel Lokshtanov, Ramamohan Paturi, Suguru Tamaki, Ryan Williams, and Huacheng Yu. *Beating Brute Force for Systems of Polynomial Equations over Finite Fields*, pages 2190–2202.
- Mac16. F.S. Macaulay. The algebraic theory of modular systems. *Cambridge Mathematical Library*, xxxi, 1916.
- MJSC16. Pierrick Méaux, Anthony Journault, François-Xavier Standaert, and Claude Carlet. Towards stream ciphers for efficient FHE with low-noise ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 311–343. Springer, Heidelberg, May 2016.

- MMDB08. Mohamed Saied Emam Mohamed, Wael Said Abd Elmageed Mohamed, Jintai Ding, and Johannes A. Buchmann. MXL2: Solving polynomial equations over $\text{GF}(2)$ using an improved mutant strategy. In Johannes Buchmann and Jintai Ding, editors, *Post-quantum cryptography, second international workshop, PQCRYPTO 2008*, pages 203–215. Springer, Heidelberg, October 2008.
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.
- MST03. Elchanan Mossel, Amir Shpilka, and Luca Trevisan. On e-biased generators in NC_0 . In *44th FOCS*, pages 136–145. IEEE Computer Society Press, October 2003.
- OW14. Ryan O’Donnell and David Witmer. Goldreich’s prg: Evidence for near-optimal polynomial stretch. pages 1–12, 06 2014.
- Pie12. Krzysztof Pietrzak. Cryptography from learning parity with noise. In Mária Bieliková, Gerhard Friedrich, Georg Gottlob, Stefan Katzenbeisser, and György Turán, editors, *SOFSEM 2012: Theory and Practice of Computer Science*, pages 99–114, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- Sch80. J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, oct 1980.
- Sie84. T. Siegenthaler. Correlation-immunity of nonlinear combining functions for cryptographic applications (corresp.). *IEEE Transactions on Information Theory*, 30(5):776–780, 1984.
- Üna23a. Akin Ünal. Worst-case subexponential attacks on prgs of constant degree or constant locality. Cryptology ePrint Archive, Paper 2023/119, 2023. <https://eprint.iacr.org/2023/119>.
- Üna23b. Akin Ünal. Worst-case subexponential attacks on prgs of constant degree or constant locality. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 25–54, Cham, 2023. Springer Nature Switzerland.
- Vio08. Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d . In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 124–127, 2008.
- Wil14. Richard Ryan Williams. The Polynomial Method in Circuit Complexity Applied to Algorithm Design (Invited Talk). In Venkatesh Raman and S. P. Suresh, editors, *34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014)*, volume 29 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47–60, Dagstuhl, Germany, 2014. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- YC05. Bo-Yin Yang and Jiun-Ming Chen. All in the x_1 family: Theory and practice. In Choon-sik Park and Seongtaek Chee, editors, *Information Security and Cryptology – ICISC 2004*, pages 67–86, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

- YGJL22. Jing Yang, Qian Guo, Thomas Johansson, and Michael Lentmaier. Revisiting the concrete security of goldreich’s pseudorandom generator. *IEEE Transactions on Information Theory*, 68(2):1329–1354, 2022.
- Zic17. Lior Zichron. Locally computable arithmetic pseudorandom generators, 2017.
- Zip79. Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, pages 216–226, Berlin, Heidelberg, 1979. Springer Berlin Heidelberg.