



Low-Complexity Weak Pseudorandom Functions in $\text{AC0}[\text{MOD2}]$

Elette Boyle^{1(✉)}, Geoffroy Couteau², Niv Gilboa³, Yuval Ishai⁴, Lisa Kohl⁵,
and Peter Scholl⁶

¹ IDC Herzliya, Herzliya, Israel
eboyle@alum.mit.edu

² IRIF, Paris, France
couteau@irif.fr

³ Ben-Gurion University, Beersheba, Israel
gilboan@bgu.ac.il

⁴ Technion, Haifa, Israel
yuvali@cs.technion.ac.il

⁵ Cryptology Group, CWI Amsterdam, Amsterdam, The Netherlands
lisa.kohl@cwi.nl

⁶ Aarhus University, Aarhus, Denmark
peter.scholl@cs.au.dk

Abstract. A *weak pseudorandom function* (WPRF) is a keyed function $f_k : \{0, 1\}^n \rightarrow \{0, 1\}$ such that, for a random key k , a collection of samples $(x, f_k(x))$, for *uniformly random* inputs x , cannot be efficiently distinguished from totally random input-output pairs (x, y) . We study WPRFs in $\text{AC0}[\text{MOD2}]$, the class of functions computable by AC0 circuits with parity gates, making the following contributions.

- **WPRF by sparse polynomials.** We propose the first WPRF candidate that can be computed by sparse multivariate polynomials over \mathbb{F}_2 . We prove that it has subexponential security against linear and algebraic attacks.
- **WPRF in $\text{AC0} \circ \text{MOD2}$.** We study the existence of WPRFs computed by AC0 circuits *over* parity gates. We propose a modified version of a previous WPRF candidate of Akavia et al. (ITCS 2014), and prove that it resists the algebraic attacks that were used by Bogdanov and Rosen (ECCC 2017) to break the original candidate in quasipolynomial time. We give evidence against the possibility of using *public* parity gates and relate this question to other conjectures.
- **Between Lapland and Cryptomania.** We show that WPRFs in $\text{AC0}[\text{MOD2}]$ imply a variant of the Learning Parity with Noise (LPN) assumption. We further show that WPRFs in a subclass of $\text{AC0}[\text{MOD2}]$ that includes a recent candidate by Boyle et al. (FOCS 2020) imply, under a seemingly weak additional conjecture, public-key encryption.

1 Introduction

This work explores the minimal achievable complexity of *weak pseudorandom functions*. Roughly speaking, a pseudorandom function (PRF) family [31] is a collection of efficiently computable functions $f_k(x)$, such that a random function from

the collection induced by a uniform choice of the key k cannot be efficiently distinguished from a truly random function. The existence (or nonexistence) of PRFs in low complexity classes is closely related to questions in computational learning theory [38, 58]: Indeed, any complexity class rich enough to contain PRFs is inherently unlearnable, even when membership queries are allowed. In this light, understanding the feasibility of low-complexity PRFs corresponds to exploring the border between the learnable and the unlearnable. More broadly, the study of low-complexity PRFs has proven to be a rich and fruitful research direction, motivated by many connections with circuit lower bounds [43, 54, 56], derandomization [49, 61], and “high-end” cryptographic applications [2, 8, 14, 15, 17, 42].

We focus on the existence of *weak* pseudorandom functions (WPRFs) in $\text{AC0}[\text{MOD2}]$, the class of polynomial-size, constant-depth circuits over AND, OR, XOR gates and negations.¹

Informally, a WPRF relaxes a PRF by restricting the distinguisher to only get input-output pairs for *uniformly random* inputs x , as opposed to chosen inputs x . WPRFs imply hardness results for learning (without membership queries) under the uniform distribution, and can serve as useful building blocks for most “symmetric” cryptographic primitives, such as private-key encryption and message authentication [46]. As a result, minimizing their complexity can lead to improving the complexity of these primitives.

Levels of security. We say that a WPRF has quasipolynomial, subexponential, or exponential security when the distinguisher’s circuit size is bounded by a corresponding function of the key length. Concretely, there exists $c > 0$ such that every circuit of size $T = n^{\log^c n}$, $T = 2^{n^c}$, or $T = 2^{cn}$ (respectively) has at most $1/T$ distinguishing advantage between f_k and a random function, for all sufficiently large key lengths n , given unlimited access to examples on uniformly random inputs. In the case of quasipolynomial and subexponential security, we can equivalently let n be the input length, since the key length and input length are polynomially related. In this work we consider subexponential security by default. This is typically the best level of security achieved by constructions from standard cryptographic assumptions.

WPRFs in low complexity classes. We return to the question of WPRFs in $\text{AC0}[\text{MOD2}]$. At the lower end, much is known about the power and limitations of AC0 . This includes unconditional circuit lower bounds (e.g. AC0 cannot compute parity [28, 32]), derandomization (e.g. AC0 cannot distinguish any polylog-wise independent distribution from the uniform distribution [16]), and learning algorithms (e.g. AC0 can be learned from quasipolynomially many samples under the uniform distribution [41]). The latter imply, in particular, that AC0 cannot contain a WPRF with better than quasipolynomial security. Slightly above

¹ More precisely, AND/OR/XOR gates can have an unbounded fan-in, and depth is defined to be the length of the longest path from an input to the output, not counting negations. As is common in the study of constant-depth PRFs, we consider the complexity of mapping the input to the output when the key is fixed.

$\text{AC0}[\text{MOD2}]$, the picture is also relatively clear: strong PRFs with subexponential security exist in the class TC0 (of polynomial-size constant-depth circuits with threshold gates) under standard cryptographic assumptions [9, 47, 48]. In contrast, despite some partial results [1, 6, 14, 15, 60], the space in between AC0 and TC0 remains a relatively uncharted territory.

Sparse \mathbb{F}_2 -polynomials. Sparse polynomials are a natural object of study in several areas, including computational learning theory. We will be interested in sparse n -variate polynomials over \mathbb{F}_2 , namely sums of $\text{poly}(n)$ monomials. Sparse \mathbb{F}_2 -polynomials can be viewed as the subclass of $\text{AC0}[\text{MOD2}]$ corresponding to depth-2 circuits that take the XOR of ANDs of inputs. A WPRF in this class would show the hardness of learning sparse \mathbb{F}_2 -polynomials under the uniform distribution. We briefly survey some relevant known results.

A result of Hellerstein and Servedio [34] implies an $2^{\tilde{O}(\sqrt{n})}$ -time PAC learning algorithm (applying to any input distribution) for learning sparse \mathbb{F}_2 -polynomials. In the converse direction, a recent work of Daniely and Vardi [24] shows that sparse \mathbb{F}_2 -polynomials are hard to learn in better than *quasipolynomial time*, albeit only under a specific non-uniform input distribution (a highly biased Bernoulli distribution), under the conjectured existence of polynomial-stretch local pseudorandom generators [4, 30, 35]. Finally, Boneh et al. [14] put forward a WPRF candidate in ACC0 that implies $2^{\Omega(n)}$ -hardness of learning sparse \mathbb{F}_3 -polynomials, again under a special input distribution (uniform over $\{-1, 1\}^n$). To our knowledge, no result is currently known that supports the hardness of learning sparse \mathbb{F}_2 -polynomials in *any* hardness regime under the uniform distribution, or in the subexponential hardness regime under *any* distribution.

The class $\text{AC0} \circ \text{MOD2}$. The class $\text{AC0} \circ \text{MOD2}$ of AC0 on *top* of parities can be seen as a minimal extension of AC0 . Despite its apparent simplicity, it is quite poorly understood. In particular, it is open whether the mod-2 inner-product function is in this class [56]. Akavia et al. [1] put forward the question of WPRFs in $\text{AC0} \circ \text{MOD2}$ as a second-best alternative to WPRFs in AC0 . They presented a candidate construction where $f_k(x)$ applies a specific DNF formula (the “TRIBES” function) to a secret linear mapping $A_k \cdot x$ of the input x , and proved resistance against several classes of attacks. However, this candidate was later broken by a quasipolynomial-time algebraic attack [13] exploiting the low *rational degree* of functions f_k in the family. Namely, there exists a low-degree g for which $f_k \cdot g = 0$ or $(f_k \oplus 1) \cdot g = 0$. This kind of attacks further rules out the possibility of any WPRF with better than quasipolynomial security that can be computed by depth-2 AC0 circuits over XOR.

1.1 Our Contribution

Candidate WPRF by sparse \mathbb{F}_2 -polynomials. We present a candidate WPRF in the class of sparse \mathbb{F}_2 -polynomials that can be conjectured to have subexponential security. More concretely, we conjecture our candidate to be secure against distinguishers of size $T = 2^{n^\epsilon}$ for a constant $\epsilon \geq 1/8$, where

n is the input size. To our knowledge, this is the first proposal for a candidate WPRF in this class. We give several kinds of evidence for the security of our candidate. First, building on previous works, we show that it has high rational degree. This implies that it cannot be broken by a subexponential algebraic attack (the same attack that breaks the candidate of [1] in quasipolynomial time). Second, we reduce its security to a *variable-density* variant of the Learning Parity with Noise (LPN) [12] assumption. This assumption is similar to (but essentially incomparable) to the variable-density LPN assumption used in the recent work of Boyle et al. [15] to build a WPRF in the class of XNF formulas (sparse \mathbb{F}_2 -polynomials in the inputs *and their negations*). Finally, we prove that it cannot be broken by any attack that fits into the framework of *linear attacks*, a general framework that captures in particular all known attacks against the LPN assumption and its variants. Our analysis builds upon the analysis of [15]; however, our setting involves additional challenges that require to significantly refine their proof techniques.

Our candidate WPRF provides an explicit distribution \mathcal{D} over sparse n -variate \mathbb{F}_2 -polynomials such that the following plausibly holds: no circuit of size $2^{n^{1/s}}$, given the values of a secret polynomial $p \in_R \mathcal{D}$ on uniformly random inputs, can predict the value of p on a fresh random input with better than $2^{-n^{1/s}}$ advantage.

As noted above, the recent work of Daniely and Vardi [24] shows hardness of learning sparse \mathbb{F}_2 -polynomials, assuming the existence of local pseudorandom generators. Our results are incomparable (and complementary) to their result:

- The result of [24] only shows the hardness of learning sparse \mathbb{F}_2 -polynomials for inputs sampled from a very specific distribution \mathcal{D} over strings $\{0, 1\}^n$, which outputs n independent samples from a highly biased Bernoulli distribution. In contrast, our results hold with respect to the *uniform* distribution.
- The result of [24] fundamentally cannot apply to the subexponential regime. The core reason is the following: from the existence of a learner for s -sparse polynomials given N examples, [24] only derives a contradiction to the existence of $(\log s)$ -local PRGs which stretch N bits from their input. However, it is known [44] that logarithmic-locality pseudorandom generators cannot possibly achieve stretch beyond quasipolynomial. Therefore, their result does not apply to the setting where s is polynomial and N is subexponential. In contrast, our result applies even to subexponential-time learning algorithms, in the setting where s is polynomial.
- On the other hand, the result of [24] relies on the existence of local PRGs, which is a relatively well-established assumption. In contrast, our result relies on a new variant of LPN, which we support by proving that it resists a large class of attacks (including in particular all standard attacks against LPN).

Candidate weak PRF in $\text{AC0} \circ \text{MOD2}$. We revisit the question of Akavia et al. [1]:

Can weak pseudorandom functions exist in the class $\text{AC0} \circ \text{MOD2}$?

We present a new candidate WPRF in $\text{AC0} \circ \text{MOD2}$ which follows the high-level template of Akavia et al. [1], but with an alternative choice of AC0 circuit structure. The WPRF candidate of Akavia et al. [1] (hereafter referred to as the “ABGKR” candidate) is of the form

$$f_{s,K}(x) = \langle x, s \rangle \oplus g(K \cdot x \pmod 2)$$

for $s \in \{0,1\}^n$, $K \in \{0,1\}^{(n-1) \times n}$, where $g(x) = \bigvee_{i=1}^\lambda \bigwedge_{j=1}^{\log \lambda} x_{ij}$ is a DNF (the so-called TRIBES function). Since $f_{s,K}(x)$ can be written as $(\neg \langle x, s \rangle \wedge g(K \cdot x)) \vee (\langle x, s \rangle \wedge \neg g(K \cdot x))$, it indeed belongs to $\text{AC0} \circ \text{MOD2}$. Notice that this candidate is an instance of the learning parity with simple deterministic noise framework, where $g(\cdot)$ is the noise function. Since the noise function is biased, XORing it with $\langle x, s \rangle$ makes the final function balanced.

Unfortunately, this candidate was broken in [13] by an algebraic attack. In our candidate, we address this issue by simply adding a layer of OR gates after the parity layer, replacing the noise function with:

$$g(x) = \bigvee_{i=1}^\lambda \bigwedge_{j=1}^\lambda \bigvee_{k=1}^w x_{ijk}.$$

We conjecture that our candidate is a subexponentially secure WPRF. We observe that our candidate resists the same classes of attacks as addressed for the ABGKR candidate. However, we are further able to prove that our candidate construction has *high rational degree*, thus circumventing the algebraic attacks under which the ABGKR candidate was insecure.²

We also study the resistance of our candidate against *linear attacks*, a large class of attacks that includes most state-of-the-art attacks on learning parity problems (such as the learning parity with noise assumption), whose structure bears connections to our candidate. We put forth a conjecture which, if true, implies that our candidate (as well as the WPRF candidates of [1, 14]) cannot be broken by linear attacks.

We view our results as providing a strong indication that $\text{AC0} \circ \text{MOD2}$ may not be learnable under the uniform distribution. We compare our results to known results regarding low-complexity PRFs on Table 1. As shown in the Table, our work fills gaps in our understanding of the complexity of weak PRFs.

On WPRFs in AC0 on top of *public* parities. The conjectured security of our candidate above relies on the MOD2 portion of the $\text{AC0} \circ \text{MOD2}$ circuit remaining *secret*, dictated by the secret WPRF key. We further revisit the question:

Can WPRF exist in the class formed by AC0 atop public parities?

² Formally, high rational degree does not prove resistance to the attack from [13], which only requires *proximity* to low rational degree. However, we view this as strong evidence that the attack does not apply to our candidate.

Table 1. Comparison of positive and negative results for low-depth PRFs. We consider the complexity of computing the output for any fixed key, where security level is with respect to the key length (see Definition 4). We write $\text{AC0}[\text{MOD}2]$ to denote the class AC0 with XOR gates at all levels, and ACC0 to denote the class AC0 with MOD_m gates for a fixed integer m ($m = 6$ suffices). RLF refers to the conjectured one-wayness of random local functions [30] and Factor to the intractability of factoring.

Circuit Class	Reference	Flavor	Security	Assumption
AC0	[12]	Weak PRF	Quasipolynomial	Heuristic
	[6, 39]	Weak PRF	Quasipolynomial	Factor, RLF
	[41]	No WPRF with better than quasipolynomial sec.		
AC0 + O(1) XOR, MAJ	[60]	No Strong PRF		
Sparse \mathbb{F}_2 -polynomials	This work	Weak PRF	Subexponential	Heuristic
	[34]	No WPRF with input length n and better than $2^{\tilde{O}(\sqrt{n})}$ sec.		
XNF formulas	[15]	Weak PRF	Subexponential	Heuristic
AC0 \circ MOD2	[1, 13]	Weak PRF	Quasipolynomial	Heuristic
	This work	Weak PRF	Subexponential	Heuristic
AC0[MOD2]	[39, 47, 60]	Strong PRF	Quasipolynomial	DDH, Factor
	[15]	Weak PRF	Subexponential	Heuristic
	[18, 40, 54]	No strong PRF with better than quasipolynomial sec.		
ACC0	[14]	Weak & Strong	Exponential	Heuristic
almost-AC0[MOD2]	[62]	Strong PRF	Subexponential	Low-noise LPN
quasilinear-TC0	[43]	Strong PRF	Exponential	Heuristic
TC0	[9, 47, 48]	Strong PRF	Subexponential	LWE, DDH, Fact

That is, we study the (in)existence of WPRFs of the form $f_k(x) = g_k(G \cdot x)$, where $g_k \in \text{AC0}$ and G is a public matrix. The existence of such a candidate would imply AC0 is not weakly learnable on all linear distributions (i.e. uniform distributions over linear subspaces of \mathbb{F}_2^n). Note, however, that it does not directly imply strong learnability, as boosting techniques would require the learner to modify the input distribution, an option that is not available for WPRFs.

We put forth a conjecture regarding the heavy Fourier coefficients of functions of this form, which implies that no WPRF can exist in AC0 on top of public parities. This is a direct strengthening of a conjecture of [1], which asserts the *existence* of a heavy Fourier coefficient for any function in this class. We conjecture further about the *form* of a heavy Fourier coefficient: namely, its expressibility as $G^T \cdot b$ for a low-weight vector $b \in \{0, 1\}^n$. This conjectured form implies that a heavy Fourier coefficient can be found within quasipolynomial time, and leveraged to obtain nontrivial advantage in distinguishing the function from random.

We demonstrate that both pieces of evidence supporting the (more conservative) conjecture of Akavia et al. [1] apply as well to our strengthened variant. Namely, the conjecture provably holds for the case of:

- Arbitrary $g_k \in \text{AC0}$ and “typical” public matrices G , including random matrices with high probability. More concretely, any G for which $G \cdot x$ for uniform inputs x fools AC0 .
- Arbitrary public G , and g_k of polynomial size and depth 2 (i.e., CNF/DNF).

We observe that Akavia et al.’s proof for the former immediately applies to our setting as well; namely, the heavy Fourier coefficient they demonstrate already is of the desired form. The latter claim holds via a more subtle extension of the argument of Jackson [36], beyond the treatment within [1].

Relation between conjectures. We map the relation between the various conjectures posed within this work and beyond (depicted in Sect. 4.3, Fig. 1).

In particular, we draw a connection between our results and the linear IPPP conjecture of Servedio and Viola [56]: we observe that the nonexistence of WPRF in AC0 over public parities (which follows from our conjecture above), together with the *existence* of a WPRF in $\text{AC0} \circ \text{MOD2}$ (for which we provide a candidate) implies the Linear IPPP conjecture.

A related but technically incomparable observation was recently made in [27], which proves under a standard cryptographic assumption (namely, the learning with rounding assumption [9]) that either (1) the known quasipolynomial time learning algorithm for AC0 under the uniform distribution [41] cannot be extended to all \mathbb{F}_2 -linear distributions, even with subexponential time, or (2) an IPPP-style hardness conjecture is true, in the sense that $\text{AC0} \circ \text{MOD2}$ cannot compute inner-products *over the integers* (as opposed to inner product modulo 2). The paper also achieves related results under the assumption underlying the WPRF candidate of [14]. Our result is incomparable: it relies on new assumptions regarding the security of WPRF candidates in $\text{AC0} \circ \text{MOD2}$ instead of standard cryptographic assumptions, but applies to the “true” Linear IPPP conjecture instead of a variant over the integers.

Between Lapland and Cryptomania. Finally, we put forth the study of a new family of LPN-style assumptions, called *LPN with simple deterministic noise*. Roughly, these assumptions assert that one cannot distinguish pairs $(x, \langle x, s \rangle \oplus g_k(x))$ with random x from random pairs (x, y) , where s is a secret vector, and g_k is a *simple* secret function sampled at random from a family. By simple, we mean that g_k should belong to a low complexity class (such as $\text{AC0}[\text{MOD2}]$).

To our knowledge, this flavor of the learning parity with noise problem has never been studied; it bears some resemblance but is incomparable to the *learning parity with structured noise* framework of Arora and Ge [7], which consider noise patterns which are not deterministic, but satisfy some structure (typically, being roots of a low degree polynomial). This LPN with simple noise formulation captures the candidate weak PRF of [1], our candidate weak PRF in $\text{AC0} \circ \text{MOD2}$, and a recent candidate WPRF from [14] that can be viewed as being based on Learning with Rounding (LWR) [9] modulo 6. In the full version of this paper we

formulate a list of simple combinatorial properties of the noise function that we conjecture to be sufficient for the resulting candidate to defeat all *linear attacks*.

Further, we show that any candidate weak PRF in $\text{AC0}[\text{MOD2}]$ *implies* the existence of a hard instance of learning parity with simple noise. Weak PRFs in $\text{AC0}[\text{MOD2}]$ therefore necessarily live in “**Lapland**”, where there exist some codes (and deterministic noise distributions) for which the learning parity with noise assumption is hard. [1] describe a natural conjecture which implies that LPN is necessary for WPRFs in $\text{AC0} \circ \text{MOD2}$. Our result strengthens this, since it is unconditional and applies to the whole of $\text{AC0}[\text{MOD2}]$; on the other hand, we only show hardness of a specific instance of learning parity with simple noise, rather than standard LPN. It is an interesting open question to obtain a more natural LPN implication from candidate weak PRFs in $\text{AC0}[\text{MOD2}]$.

Our approach uses a result of Razborov and Smolensky [53, 57], who show that any $\text{AC0}[\text{MOD2}]$ function can be approximated by a low-degree polynomial; we show that the approximation noise itself can be used to define a learning parity with noise instance that fits our framework. On the other hand, we observe that the Razborov-Smolensky approximation could also be leveraged in a positive sense, for improving efficiency when evaluating the PRF homomorphically on ciphertexts or as part of a secure computation. For more details we refer to [10].

Note that the seeming contradiction of the Razborov-Smolensky approximation being sufficiently noisy to avoid decoding attacks (as far as we know), but precise enough to be useful for replacing the weak PRF by its approximation in applications, can be explained by the different number of input-output pairs considered in both contexts. An attacker attempting to break the security requires at least a quasipolynomial number of samples (because the low-degree multivariate polynomial potentially consists of a quasipolynomial number of terms), thus noise will occur *almost certainly*, whereas in an honest setting, when only computing a polynomial number of samples, likely the approximation will be perfect on all samples considered.

Since Lapland only has partial overlap with Cryptomania (that is, presently only LPN with low noise rate is known to imply public-key encryption with more than quasi-polynomial security³), one can further ask where in the regime between Lapland and Cryptomania weak PRFs in $\text{AC0}[\text{MOD2}]$ fall.

We put forward a second framework for “variable-density learning parity with noise (VDLPN) assumptions” into which the recent candidate weak PRF of [15] (who coined the term variable-density learning parity with noise for a specific instance of this broader framework) and our weak PRF candidate computed by sparse polynomials fall into. We further observe that any weak PRF candidate within this framework implies an instance of learning parity with simple deterministic noise with noise rate below the bound of [3]. This still does not imply public-key encryption, because the framework of [3] requires the code distribution to be dense, which is not the case for variable density learning parity with

³ More precisely, by a result of [3], random LPN implies public-key encryption if the noise rate is in $o(\sqrt{M})$, where M is the length of the secret.

noise. We still view this as an indication that weak PRFs within this framework “morally” live in Cryptomania.

To formalize this intuition, we put forward a conjecture, stating that with respect to some fixed noise rate, either *all codes* are efficiently decodable, or *almost all codes* are hard to decode. This is backed-up by the common understanding that LPN is in fact hard for random codes when choosing reasonably dense noise. Based on this conjecture we can indeed prove that candidate weak PRFs within the VDLPN framework imply public-key encryption following the strategy of [3]. We are not aware of any such implication for general functions in $\text{AC0}[\text{MOD2}]$ such as our candidate weak PRF in $\text{AC0} \circ \text{MOD2}$, even if willing to assume this conjecture, because the flavor of learning parity with noise implied by Razborov-Smolensky does not give low enough noise rate.

This is particularly interesting in light of recent developments on constructing *pseudorandom correlation functions* [15], since candidate constructions of expressive correlations so far all rely on either the VDLPN assumption [15], factoring-based assumptions [51], or extremely low-noise LPN [23], which with our result in mind, all imply public-key encryption.

2 Preliminaries

We start by recalling some basic properties of Boolean functions. We mostly follow standard notations and terminology (see, e.g., [41, 50]), except that we identify parity functions with vectors in $\{0, 1\}^n$ instead of subsets $S \subseteq \{1, \dots, n\}$.

Boolean functions. A Boolean function is a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$. When considering the Fourier coefficients of a function f we will consider it as a function $f: \{0, 1\}^n \rightarrow \{1, -1\}$ by identifying an output $b \in \{0, 1\}$ with $(-1)^b$.

The set of all *real-valued* functions on the cube $\{0, 1\}^n$ is a 2^n -dimensional real vector space with an inner product defined by $\langle g, f \rangle = 2^{-n} \cdot \sum_{x \in \{0, 1\}^n} f(x) \cdot g(x)$. The *norm* of f is defined as $\|f\| = \sqrt{\langle f, f \rangle}$.

Definition 1 (Characters). For $y \in \{0, 1\}^n$, the character χ_y is defined as $\chi_y(x) = (-1)^{\langle x, y \rangle}$.

Note that $\{\chi_y\}_{y \in \{0, 1\}^n}$ forms an orthonormal basis of the space of all real-valued functions on $\{0, 1\}^n$. Further, for all $y, z \in \{0, 1\}^n$ it holds that $\chi_y \chi_z = \chi_{y \oplus z}$.

Definition 2 (Fourier coefficients). As $\{\chi_y\}_{y \in \{0, 1\}^n}$ forms a basis, we can write every real-valued function f on the cube as $f = \sum_{y \in \{0, 1\}^n} \hat{f}(y) \cdot \chi_y$, for real-valued coefficients $\hat{f}(y)$, called Fourier coefficients.

Note that this is well-defined, as $\{\chi_y\}$ forms a basis for all functions $f: \{0, 1\} \rightarrow \mathbb{R}$. Further, as $\{\chi_y\}$ forms an orthonormal basis, the Fourier coefficient corresponding to $y \in \{0, 1\}^n$ can be computed as $\hat{f}(y) = \langle f, \chi_y \rangle$. For every Boolean function f we have $1 = \|f\|^2 = \langle f, f \rangle = \sum_{y \in \{0, 1\}^n} \hat{f}(y)^2$.

Definition 3 (Degree). *The degree $\text{deg}(f)$ of a Boolean function f is defined as the maximal Hamming weight of a vector $y \in \{0, 1\}^n$ for which $\hat{f}(y) \neq 0$.*

It can be shown that the above notion of degree coincides with standard algebraic degree.

Circuit classes. The class AC0 is the class of functions computed by a family of constant-depth, polynomial-size circuits of over AND/OR gates of unbounded fan-in along with negations. The class $\text{AC0} \circ \text{MOD2}$ is defined similarly, except that one also allows parity (XOR) gates *only at the bottom*. This can be viewed as applying an AC0 function to an \mathbb{F}_2 -linear encoding of the input. We define the circuit *depth* to be the length of the longest path from an input to an output, not counting negations. For instance, a DNF formula has depth 2. For $\text{AC0} \circ \text{MOD2}$ circuits we will consider by default only the depth of the AC0 part, namely ignoring parities. See, e.g., [1, 19, 56] for known facts about AC0 and $\text{AC0} \circ \text{MOD2}$.

In the context of cryptographic primitives, we will consider AC0 or $\text{AC0} \circ \text{MOD2}$ circuit families $\{C_\lambda\}$, parameterized by a *security parameter* λ , where the input length $n = n(\lambda)$ is assumed to be a monotonically-increasing, polynomially-bounded function of λ . We assume by default that such a circuit family is *polynomial-time uniform*, namely there is a polynomial-time algorithm whose output on input 1^λ is a description of C_λ ; however, we drop the uniformity requirement in the context of negative results.

2.1 Pseudorandom Functions

We consider here *weak* PRFs, which relax standard PRFs by only considering distinguishers that get the outputs of the function on uniformly random inputs. We require *subexponential* security by default, namely security against distinguishers of size 2^{n^ϵ} for some $\epsilon > 0$. This is the typical level of security achieved by constructions based on the strongest plausible versions of standard cryptographic assumptions. We formally define this notion below.

Definition 4 ((Weak) pseudorandom function [31, 46]). *Let $\lambda \in \mathbb{N}$ denote a security parameter and $n = n(\lambda)$, $\kappa = \kappa(\lambda)$ be monotonically-increasing and polynomially-bounded input length and key length functions, respectively.*

A (weak) pseudorandom function is syntactically defined by a function family $\mathcal{F} = \{f_\lambda: \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}\}$, where the output $f_\lambda(k, x)$ can be computed from (k, x) in polynomial time. Since λ and κ are determined by the input length n , we will sometimes write $f_k(x)$ instead of $f_\lambda(k, x)$.

For $T = T(\kappa)$ and $\epsilon = \epsilon(\kappa)$, we say that \mathcal{F} is a (T, ϵ) -secure strong pseudorandom function (PRF), if for every $\lambda \in \mathbb{N}$ and every oracle circuit \mathcal{A} of size $T(\kappa)$, it holds

$$\Pr_k[\mathcal{A}^{f_k(\cdot)} = 1] - \Pr_R[\mathcal{A}^{R(\cdot)} = 1] \leq \epsilon(\kappa),$$

where $\kappa = \kappa(\lambda)$, $k \xleftarrow{\$} \{0, 1\}^\kappa$ is chosen at random, and $R: \{0, 1\}^n \rightarrow \{0, 1\}$ is a truly random function. A T -secure PRF is a $(T, 1/T)$ -secure PRF.

We say that \mathcal{F} is a (T, ε) -secure weak PRF (WPRF) or T -secure WPRF if the above holds when \mathcal{A} only gets access to samples $(x_i, f_k(x_i))$, where $x_i \stackrel{\$}{\leftarrow} \{0, 1\}^n$ are chosen uniformly and independently. We say that \mathcal{F} is a (Q, T, ε) -secure (strong/ weak) PRF if \mathcal{A} only gets access to at most Q (chosen/ random) samples. Finally, we say that a (W)PRF \mathcal{F} has polynomial security if it is T -secure for every polynomial T , and that it has subexponential (resp., quasipolynomial, exponential) security if there exists $c > 0$ such that it is T -secure for $T = 2^{\kappa^c}$ (resp., $T = \kappa^{\log^c \kappa}$, $T = 2^{\kappa^c}$).

Our choice of subexponential security as the default level of security is motivated both from a cryptographic perspective and from an algorithmic perspective. From a cryptographic perspective, candidate PRFs with quasipolynomial security are relatively easy to obtain even in very low complexity classes and are considered “borderline insecure.” Subexponential (rather than exponential) security is typically the best level of security one can get from standard assumptions. From an algorithmic perspective, quasipolynomial-time algorithms (such as the LMN learning algorithm [41]) are considered “borderline efficient” and hence ruling out such algorithms requires PRFs with better than quasipolynomial security.

Finally, when referring to a (W)PRF \mathcal{F} in a circuit complexity class such as AC0 or $\text{AC0} \circ \text{MOD2}$, the default convention is that for each key sequence $k(\lambda)$, the induced function family f_k is in the class. We note that even when considered as a function of both the input and the key, our candidate constructions remain in $\text{AC0}[\text{MOD2}]$. On the other hand, our negative results and conjectures are stronger in that they apply to the fixed-key case and do not assume polynomial-time uniformity.

2.2 Preliminaries on Probability

Given t distributions $(\mathcal{D}_1, \dots, \mathcal{D}_t)$ over \mathbb{F}_2^n , we denote by $\bigoplus_{i \leq t} \mathcal{D}_i$ the distribution obtained by *independently* sampling $\mathbf{v}_i \stackrel{\$}{\leftarrow} \mathcal{D}_i$ for $i = 1$ to t and outputting $\mathbf{v} \leftarrow \mathbf{v}_1 \oplus \dots \oplus \mathbf{v}_t$.

Definition 5 (Bias of a Distribution). *Given a distribution \mathcal{D} over \mathbb{F}_2^n and a vector $\mathbf{u} \in \mathbb{F}_2^n$, the bias of \mathcal{D} with respect to \mathbf{u} , denoted $\text{bias}_{\mathbf{u}}(\mathcal{D})$, is equal to $\text{bias}_{\mathbf{u}}(\mathcal{D}) = \left| \frac{1}{2} - \Pr_{\mathbf{v} \stackrel{\$}{\leftarrow} \mathcal{D}} [\mathbf{u}^\top \cdot \mathbf{v} = 1] \right|$. Then, the bias of \mathcal{D} , denoted $\text{bias}(\mathcal{D})$, is defined as $\text{bias}(\mathcal{D}) = \max_{\mathbf{u} \neq \mathbf{0}^n} \text{bias}_{\mathbf{u}}(\mathcal{D})$.*

2.3 Algebraic Attacks and Rational Degree

Algebraic attacks have been introduced in [52] and were extended and abstracted in [20–22]. In its most basic form, an algebraic attack proceeds as follows: given a function $F : \{0, 1\}^n \mapsto \{0, 1\}$, it finds low degree multivariate polynomials (g, h) such that $F \cdot g = h$. If polynomials (g, h) of degree at most d are found, then the function F can be inverted given $n^{\tilde{O}(d)}$ random samples $(x, F(x))$. The

hardness of inverting a function with an algebraic attack is measured by its rational degree:

Definition 6 (Rational Degree). *The rational degree of a boolean function F is defined as the following quantity:*

$$\text{RD}(F) = \min_{g \neq 0} \{\text{deg}(g) \mid Fg = 0 \vee (F \oplus 1)g = 0\}.$$

Observe that the smallest d such that there exist polynomials (g, h) of degree at most d satisfying $F \cdot g = h$ necessarily satisfies $d \geq \text{RD}(F)$.

3 WPRFs by Sparse Multivariate Polynomials

In this section, we put forth a new candidate WPRF in a very low subclass of $\text{ACO}[\text{MOD}2]$: the class of *sparse multivariate polynomials* over \mathbb{F}_2 . That is, the key defines a sum of $\text{poly}(n)$ monomials in the inputs x_1, \dots, x_n . We conjecture that our candidate achieves subexponential security. To our knowledge, this is the first proposal for a WPRF in this class with plausible subexponential security.

In more detail, our candidate is inspired by a WPRF candidate from [15], which belongs to the class of *XNF formulas*, i.e., sparse polynomials in the inputs *and their negations*. Multivariate polynomials are an important object of study in learning theory. Our candidate WPRF provides an explicit distribution \mathcal{D} over sparse n -variate \mathbb{F}_2 -polynomials such that the following plausibly holds: there is a constant $\varepsilon > 0$ such that no 2^{n^ε} -time algorithm, given the values of a polynomial p sampled from \mathcal{D} on uniformly random inputs, can predict the value of p on a fresh random input with better than 2^{-n^ε} advantage. In contrast, the candidate of [15] only implies hardness of learning sparse polynomials under a somewhat artificial input distribution: the distribution over vector pairs (\mathbf{x}, \mathbf{y}) where \mathbf{y} is the bitwise negation of \mathbf{x} . To our knowledge, the only previous results in this setting are limited to showing *quasi-polynomial* hardness of learning sparse \mathbb{F}_2 -polynomials under the uniform distribution [24]. Our candidate complements the results of [34], which imply a $2^{\tilde{O}(\sqrt{n})}$ -time learning algorithms for sparse \mathbb{F}_2 -polynomials.

To support the conjectured subexponential security of our new candidate, we first observe that known results imply that it cannot be broken by *algebraic attacks*, as defined in Sect. 2. Furthermore, we show that its security can be formulated as an *LPN-style* assumption, which closely resembles (but is technically incomparable to) the variable-density learning parity with noise assumption of [15]. We provide support for the security of the candidate by proving that it cannot be broken in subexponential time by any *linear attack*, a large class of attacks which captures essentially all known attacks against LPN and its variants. Our analysis builds upon, but does not follow from, the analysis of [15]. In the full version we elaborate on the specific challenges that arise when trying to extend the analysis of [15] to our candidate.

3.1 Our Candidate

Our candidate builds upon the candidate of [15], which was carefully crafted as a XOR of variable-size terms (products of variables and negated variables), where the purpose of terms of size i is to defeat all linear attacks that depend on (approximately) 2^i samples. In [15], the set of input variables in each term is fixed in advance; the WPRF key simply tells, for each variable in each term, whether to use the input or its negation. To confine our candidate to the subclass of sparse \mathbb{F}_2 -polynomials, we must refrain from using negations of inputs. This suggests a very natural variant: instead of selecting between bits x and $1 - x$, the key is used to randomly select one out of b random bits $x_1 \cdots x_b$ for each variable of each monomial. When b is large enough, since the fraction of zeroes and ones in random b -bit strings is tightly concentrated around $1/2$, this intuitively provides security guarantees comparable to that of [15]. We formally introduce the candidate below.

- **Input domain:** $x \in \{0, 1\}^n$ with $n = w \cdot D \cdot (D - 1) \cdot b/2$. We view x as a concatenation of D blocks $(x_i)_{i \leq D}$, where block x_i contains w sub-blocks $x_{i,1}, \dots, x_{i,w}$, and each sub-block $x_{i,j}$ is composed of i b -bit strings $(x_{i,j,\ell})_{\ell \leq i}$. Given a string $x_{i,j,\ell}$, we write $x_{i,j,\ell}[k]$ to denote its k -th bit.
- **Key domain:** $K = (K_{i,j,\ell})_{i \leq D, j \leq w, \ell \leq i} \in [b]^s$ with $s = w \cdot \sum_{i=1}^D i$.
- **Candidate:**

$$F_K(x) = \bigoplus_{i=1}^D \bigoplus_{j=1}^w \bigwedge_{\ell=1}^i x_{i,j,\ell}[K_{i,j,\ell}]$$

Security against algebraic attacks. The security of our candidate against algebraic attacks [22] follows directly from a known bound on the rational degree of triangular functions.

Lemma 7. *For any $K \in [b]^s$, an algebraic attack in the sense of [22] requires (time and) number of samples lower bounded by $n^{\Omega(D)} = 2^{\Omega(D \log(D+w+b))}$.*

Lemma 7 follows readily from the fact that our candidate weak PRF has high *rational degree*: for any $K \in \{0, 1\}^s$, it holds that $\text{RD}(F_K) \geq D$. The proof follows immediately from [42]: for any fixed choice of key K , F_K is a direct sum of w independent *triangular functions of degree D* , each evaluated on distinct portions of the input, where (denoting $D' = D(D - 1)/2$) the triangular function of degree D is the function $T_D(x_1, \dots, x_{D'}) = x_1 \oplus x_2 x_3 \oplus \dots \oplus \bigwedge_{\ell=D'-D}^{D'} x_\ell$. By Lemma 3 of [42], the rational degree of a direct sum of functions is at least the largest rational degree of its components, and by Lemma 6 of [42], the rational degree of T_D is exactly D .

3.2 Variable-Density LPN Formulation

We now show that the security of our weak PRF candidate follows from a VDLPN-style assumption, in the spirit of [15]. We note, however, that the concrete assumption is not directly comparable to that of [15]: while the corresponding noise distributions are similar, the variable-density matrix distribution

in our work is very different. In the following, for each $(i, j) \in [D] \times [w]$, it is convenient to view $K_{i,j} = (K_{i,j,\ell})_{\ell \leq i}$ as a single integer from the set $[b^i]$, via the natural embedding. Then, let $\mathbf{u}(K_{i,j})$ denote the unit length- b^i vector with a 1 at position $K_{i,j}$ and 0's elsewhere. We can rewrite F_K as

$$\begin{aligned}
 F_K(x) &= \bigoplus_{i=1}^D \bigoplus_{j=1}^w \left\langle \bigotimes_{\ell=1}^i x_{i,j,\ell}, \mathbf{u}(K_{i,j}) \right\rangle \\
 &= \left\langle x_{1,1,1} \|\cdots\| \bigotimes_{\ell=1}^D x_{D,w,\ell}, \mathbf{u}(K_{1,1}) \|\cdots\| \mathbf{u}(K_{D,w}) \right\rangle = \langle h(x), e(K) \rangle
 \end{aligned}$$

where $h : x \rightarrow (x_{1,1,1} \|\cdots\| \bigotimes_{\ell=1}^D x_{D,w,\ell})$ and $e : K \rightarrow (u(K_{1,1}) \|\cdots\| \mathbf{u}(K_{D,w}))$. Now, given a bound N on the number of samples, we let $\mathcal{H} = \mathcal{H}(D, w, \mathbf{b}, N)$ denote the distribution over matrices H in $\mathbb{F}_2^{N \times (w \cdot \sum_{i=1}^D b^i)}$ whose N rows are sampled as $h(x)$ for independent samples $x \xleftarrow{\$} \{0, 1\}^n$. Furthermore, we let $\mathcal{N} = \mathcal{N}(D, w, \mathbf{b})$ denote the distribution over vectors \mathbf{e} in $\mathbb{F}_2^{w \cdot \sum_{i=1}^D b^i}$ induced by sampling $K \xleftarrow{\$} [b]^s$ and outputting $e(K)$. Clearly, breaking the security of our candidate given N samples is equivalent to breaking the $(\mathcal{H}, \mathcal{N})$ -dualLPN assumption. This variant of the dual LPN assumption is very close in spirit to the regular VDLPN assumption from [15]: the noise distribution is the same up to setting $\mathbf{b} = 2$. The matrix distribution, on the other hand, is quite different, but satisfies the same sparsity condition: the matrix H is divided into D submatrices H_i , and the average sparsity of the rows of H_i is $(w \cdot (b/2)^i) / (w \cdot b^i) = 1/2^i$. The matrix distribution in [15] satisfies the same variable density structure, which motivated the name ‘‘variable-density LPN’’. Therefore, we view our new candidate as belonging to the same family of LPN variants.

3.3 Security Against Linear Attacks

We turn to consider the class of *linear attacks*, which in the context of pseudorandom *generators* captures the notion of small-bias generators [45]. Linear attacks capture, intuitively, every attack where the distinguisher is restricted to compute a linear function of the LPN samples, the identity of which can be arbitrarily determined from the public LPN matrix and inputs. This captures essentially all known attacks against standard variants of LPN, such as those based on Gaussian elimination, statistical decoding, information set decoding, and BKW-style attacks. The work of [15] provided support for their VDLPN conjecture by proving subexponential security against such linear attacks.

In the context of a WPRF, a linear distinguisher is first given N random inputs x_1, \dots, x_N , and then must choose a subset of indices $S \subset \{1, \dots, N\}$ such that the distribution $\bigoplus_{i \in S} f_k(x_i)$, for a random choice of k , is biased towards 0 or 1. More formally, we use the following notion of an (ϵ, δ, N) -biased WPRF, which naturally extends the standard notion of an ϵ -biased pseudorandom generator.

Definition 8 ((ε, δ, N) -biased weak PRF family, [15]). A function family $\{F_K : \mathbb{F}_2^{n(\lambda)} \mapsto \mathbb{F}_2\}_{K \in \mathbb{F}_2^{s(\lambda)}}$ is (ε, δ, N) -biased if for every large enough $\lambda \in \mathbb{N}$, letting $\mathcal{D}_{\lambda, N}(\mathbf{x})$ (for some $\mathbf{x} \in (\mathbb{F}_2^{n(\lambda)})^N$) biased with inputs of length samples $K \xleftarrow{\$} \mathbb{F}_2^{s(\lambda)}$ and outputs $\mathbf{y} = (F_K(x^{(1)}), \dots, F_K(x^{(N)}))$, it holds that

$$\Pr_{x^{(1)}, \dots, x^{(N(\lambda))} \xleftarrow{\$} \mathbb{F}_2^{n(\lambda)}} [\text{bias}(\mathcal{D}_{\lambda, N}(\mathbf{x})) > \varepsilon(\lambda)] \leq \delta(\lambda).$$

Notation and theorem statement. We first introduce some notation. Recall that a sample H from \mathcal{H} is a concatenation of D matrices H_i , where each matrix H_i is itself a concatenation of w submatrices $H_{i,j} \in \mathbb{F}_2^{N \times b^i}$ whose rows are of the form $\bigotimes_{\ell=1}^i x_{i,j,\ell}$, where the $(x_{i,j,\ell})_{\ell \leq i}$ are i uniformly random independent b -bit strings. For any fixed matrix H in the support of \mathcal{H} , we let $\mathcal{D}_{\text{out}}(H)$ denote the distribution induced by sampling $\mathbf{e} \leftarrow \mathcal{N}$ and outputting $H \cdot \mathbf{e}$.

Theorem 9 (Low bias). Fix a security parameter λ . There exist constants $0 < \beta, \nu, \mu < 1$ such that for any parameters (D, w, \mathbf{b}, N) satisfying $w = \text{poly}(\lambda)$, $\mathbf{b} = \text{poly}(\lambda)$, $D^2 \leq \beta \cdot w$, $D \leq \frac{\sqrt{b}}{2\lambda} + 1$, and $N \leq 2^D$, letting $\mathcal{H} = \mathcal{H}(D, w, \mathbf{b}, N)$, it holds that

$$\Pr_{H \leftarrow \mathcal{H}} [\text{bias}(\mathcal{D}_{\text{out}}(H)) > \mu^w] \leq \nu^D + \nu^{\lambda^2}.$$

For example, using the choice of parameters $(D, w, \mathbf{b}, N) = (\lambda, \lambda^2/\beta, 4\lambda^4, 2^\lambda)$, our candidate is $(2^{-\Omega(\lambda^2)}, 2^{-\Omega(\lambda)}, 2^\lambda)$ -biased with inputs of length $O(\lambda^8)$, and keys of length $\tilde{O}(\lambda^4)$.

To facilitate comparison with the analysis of [15], we let \mathcal{H}' and \mathcal{N}' denote respectively the matrix and noise distributions for the VDLPN variant of [15], where a sample $H \leftarrow \mathcal{H}'$ can also be broken into D matrices $H_i = H_{i,1} \parallel \dots \parallel H_{i,w}$ where the $H_{i,j}$ are independent matrices; we denote by \mathcal{H}'_i the distribution over H_i induced by $H \leftarrow \mathcal{H}'$ for any $i \leq D$.

High level overview. At a high level, the security analysis follows the same approach as the analysis in [15] (which should come as no surprise due to the similarities between the candidates); however, the analysis is significantly more involved due to the more complex structure of the matrix distribution for our candidate. Fix $i \leq D$. The analysis of [15] proceeds roughly as follows.

1. Using a strong concentration bound (McDiarmid’s bounded difference inequality), it shows that for any fixed *attack vector* $\mathbf{v} \in \mathbb{F}_2^N$ whose Hamming weight is between 2^{i-1} and 2^i , except with probability at most $\exp(-\Omega(w \cdot 2^i))$, a random matrix $H_i \leftarrow \mathcal{H}'_i$ satisfies $\mathcal{HW}(\mathbf{v}^\top \cdot H_{i,j}) / |\mathbf{v}^\top \cdot H_{i,j}| \in [\varepsilon, 1 - \varepsilon]$, where ε is some constant (that is, $\mathbf{v}^\top \cdot H_{i,j}$ has a fraction of ones bounded by a constant, and bounded away from 1 by a constant), for a fraction at least $w/2$ of the w submatrices $H_{i,j}$ of H_j . Such a matrix H_i is called *good with respect to* \mathbf{v} .

2. From a union bound over all vectors \mathbf{v} of weight between 2^{i-1} and 2^i , it follows that, except with probability at most $\exp(-\Omega((\log N - w) \cdot 2^i))$, a random matrix $H_i \leftarrow \mathcal{H}'_i$ will be good with respect to all vectors \mathbf{v} in this weight range. When w is sufficiently larger than $\log N$, this probability is bounded by $\exp(-\Omega(w))$ for any $i \leq D$.
3. By a union bound over all $i \leq D$, with probability at least $1 - D \cdot \exp(-\Omega(w)) = 1 - \exp(-\Omega(w))$, a random matrix $H \leftarrow \mathcal{H}'$ satisfies the following: for every nonzero vector \mathbf{v} , there is an $i^* \leq D$ such that H_{i^*} is good with respect to \mathbf{v} . Then, for any such matrix H , $H \cdot \mathbf{e}$ for $\mathbf{e} \leftarrow \mathcal{N}'$ is the vector obtained by sampling a uniformly random column from each $(H_{i,j})_{i \leq D, j \leq w}$ and XORing them all. Since H_{i^*} is good with respect to \mathbf{v} , $H \cdot \mathbf{e}$ will include at least $w/2$ terms sampled randomly and independently from bitstrings $\mathbf{v} \cdot H_{i^*,j}$ with a fraction of ones in $[\varepsilon, 1 - \varepsilon]$. It follows that, with probability at least $1 - \exp(-\Omega(w))$ over the random choice of $H \leftarrow \mathcal{H}'$, the distribution of $H \cdot \mathbf{e}$ for $\mathbf{e} \leftarrow \mathcal{N}'$ has bias with respect to \mathbf{v} at most $(1 - \varepsilon)^{w/2}/2 = 2^{-\Omega(w)}$, for any possible nonzero vector \mathbf{v} .

Looking ahead, our security analysis will follow the same three steps as above, and the steps 2 and 3 will be the same as in [15]. However, while the first step also consists in proving a similar bound, the actual analysis turns out to be much more involved due to the different matrix structure. Due to space limitations, the proof of Theorem 9 is deferred to the full version.

4 WPRFs in $\text{AC0} \circ \text{MOD2}$

In this section we present a candidate construction of a weak PRF in $\text{AC0} \circ \text{MOD2}$ (recall, unlike $\text{AC0}[\text{MOD2}]$, here the parity gates must lie at the input layer of the circuit). We follow the high-level template of Akavia et al. [1]. Their construction, referred to as ABGKR, is of the form

$$f_{s,K}(x) = \langle x, s \rangle \oplus g(K \cdot x \pmod 2)$$

for $s \in \{0, 1\}^n$, $K \in \{0, 1\}^{(n-1) \times n}$, where $g(x) = \bigvee_{i=1}^\lambda \bigwedge_{j=1}^{\log \lambda} x_{ij}$ is a DNF (the so-called TRIBES function). Since $f_{s,K}(x)$ can be written as $(\neg \langle x, s \rangle \wedge g(K \cdot x)) \vee (\langle x, s \rangle \wedge \neg g(K \cdot x))$, it indeed belongs to $\text{AC0} \circ \text{MOD2}$.

The rationale behind the design of Akavia et al. is the following: even when picking a very simple function g (in their case, a DNF), the function $g_K(x) = g(K \cdot x)$ can already not be distinguished from a random c -unbalanced function (i.e. a random function f with $\Pr_x[f(x) = 1] = c$ for some constant $c \neq 1/2$) for various natural attacks (e.g. correlations with small function families and closeness to low-degree polynomial). Then, this function g_K is XORed with $\langle x, s \rangle$ to make the final function balanced.

From unbalanced WPRFs to standard WPRFs. We observe that this transformation does actually provably turn an unbalanced WPRF into a “standard” WPRF, under the LPN assumption. The proof of this observation is

straightforward; for details we refer to the full version. In spite of its simplicity, this observation had to our knowledge never been made.

We further note that there exists an alternative, unconditional transformation from a c -unbalanced WPRF in $\text{AC0} \circ \text{MOD2}$ into a standard WPRF in $\text{AC0} \circ \text{MOD2}$ which relies on the Von Neumann randomness extractor: assume w.l.o.g. that $c < 1/2$. Use (say) $2n$ parallel instances of the c -unbalanced WPRF on independent inputs and keys, grouped into n pairs. Then, take the first pair of distinct output bits (since c is a constant, there is one such pair with overwhelming probability $1 - 2^{-O(n)}$): if it is 01, define the output of the WPRF to be 0; else, define it to be 1. It is relatively straightforward to prove that if g_k is a c -unbalanced WPRF, the resulting function is a WPRF. This process can be executed in AC0 , hence the resulting function is in $\text{AC0} \circ \text{MOD2}$.

Our approach. The above discussion justifies focusing on the task of building *unbalanced* WPRFs in $\text{AC0} \circ \text{MOD2}$, since the latter imply standard WPRFs in the same class through simple transformations. The ABGKR candidate instantiates this unbalanced WPRF with a DNF on top of parities; however, the attack of [13] allows to distinguish *any* depth-2 AC0 circuit on top of parities from unbalanced random functions, since any such function must have low rational degree. Therefore, any unbalanced WPRF in $\text{AC0} \circ \text{MOD2}$ must have at least three layers of AND/OR gates. With the goal of finding the simplest possible modification of the ABGKR candidate which can retain subexponential security, we ask:

Is there a subexponentially secure unbalanced WPRF computable by a depth-3 AC0 circuit on top of parities?

Our candidate. We put forth the following candidate unbalanced WPRF: $g_k(x) = g(K \cdot x)$, with

$$g(x) = \bigvee_{i=1}^{\lambda} \bigwedge_{j=1}^{\lambda} \bigvee_{k=1}^w x_{ijk}, \tag{1}$$

where λ is a security parameter (i.e., we will bound the complexity of various attacks on our candidate as a function of λ) and m, w are chosen such that $w = \lceil \log \lambda - \log \log \lambda \rceil$ and $m = \lambda^2 w$. That is, we simply add a single layer of ORs after the parity layer, with parameters chosen to guarantee that $\Pr_x[g(x) = 1]$ is constant. Note that choosing the fan-in of the gates more carefully, one can actually obtain bias $1 = 2 + o_n(1)$. In this case the function $g(x)$, which replaces the TRIBES function in ABGKR, corresponds to the degree-3 Sipser function. For more details, we refer to [33, 55].

We conjecture that this candidate achieves subexponential security. Observe that since the attack of [13] distinguishes any depth-2 AC0 circuit on top of

parities from unbalanced random functions, our candidate actually enjoys optimal depth.⁴

4.1 Provable Resistance to Algebraic Attacks

Algebraic attacks are a general class of cryptanalytic algorithms that aim to either invert a function or distinguish it from random, by obtaining many samples and using these to derive a system of linear equations over the secret inputs. This class of attack was first developed by the applied cryptographic community and used to break public-key encryption schemes and stream ciphers [20, 22, 52]. It generalizes in particular the correlation attacks [37] that have been developed for attacking LFSRs. Correlation attacks have been considered in the theory community in the context of constructing local pseudorandom generators [44].

The resistance of a WPRF $f_k : \{0, 1\}^n \rightarrow \{0, 1\}$ to algebraic attacks can be measured by its *rational degree*, that is, the smallest d for which there exist non-zero polynomials p and q of algebraic degree at most d , such that

$$f_k(x) \cdot p(x) = q(x), \quad \forall x \in \{0, 1\}^n. \quad (2)$$

Applebaum and Lovett [5] formally studied algebraic attacks of local functions, and showed that if a predicate has large rational degree then it provably resists a natural class of algebraic attacks.

On the other hand, if a WPRF candidate f_k has low rational degree d , then it can be distinguished from random via a simple algebraic attack, which obtains $O(n^d)$ samples and tests whether (2) holds for each of them. This is exactly the type of attack that Bogdanov and Rosen [13] observed breaks the candidate of Akavia et al. [1] in quasipolynomial time, since it has rational degree $O(\log \lambda)$.

We, on the other hand, show that our candidate has rational degree λ . Even though, formally, this does not rule out the attack of [13], which only requires proximity to low rational degree, we view this as strong evidence that the attack does not apply to our candidate.

To analyze the rational degree of our candidate, we first give a general method for determining the exact rational degree of any function in $\text{AC}0$ that can be expressed as alternating layers of AND and OR gates that each depend on disjoint subsets of the input. We then use this to compute the rational degree of our noise function, and finally our candidate unbiased WPRF.

Towards understanding our techniques, we first briefly recall the attack of Bogdanov and Rosen [13]. To that end, note that the rational degree can be characterized as the minimal d such that there exists a polynomial $p \neq 0$ of algebraic degree d such that $f \cdot p = 0$ or $(f \oplus 1) \cdot p = 0$ (also referred to as the *algebraic immunity* in the literature). The attack of Bogdanov and Rosen [13]

⁴ However, transforming our candidate into a standard WPRF, e.g. using the LPN-based transformation, results in a candidate computed by a depth-4 $\text{AC}0$ circuit on top of parities. It is an interesting question whether the optimal depth can be achieved for standard WPRFs, i.e., whether there exists subexponentially-secure standard WPRFs computable by depth-3 $\text{AC}0$ circuit on top of parities.

builds on the observation that $f = \bigvee f_i$ always has rational degree at most $\min_i \deg f_i$, as $f_i(x) = 1$ implies $f(x) = 1$ and thus $(f \oplus 1) \cdot f_i = 0$. Therefore, for any DNF either all inner conjunctions have high algebraic degree (and thus the DNF is highly biased towards 0), or the function is susceptible to rational degree attacks.

We observe that while a disjunction does not increase the rational degree of a function, it does have an effect that can be leveraged. Namely, consider a function $p \neq 0$ of minimal algebraic degree such that $f \cdot p = 0$. We will prove that if p_i are the minimal annihilating functions for f_i (and all functions depend on disjoint parts of the input), p must have algebraic degree at least $\sum_i p_i$.

Now, using that conjunctions behave in a dual way, alternating between conjunctions and disjunctions allows to increase the rational degree while keeping the function's bias constant. In order to prove this, we introduce the notion of *primal* and *dual* rational degree.

Definition 10 (Primal and dual rational degree). For $f: \{0, 1\}^n \rightarrow \{0, 1\}$, we define the primal rational degree ρ as the minimal ρ such that there exists a polynomial $p \neq 0$ with algebraic degree ρ and $f \cdot p = 0$. Further, we define the dual rational degree ρ' of f as the primal rational degree of its negation. Namely, we define the dual rational degree as the minimal ρ' such that there exists a polynomial $p \neq 0$ with algebraic degree ρ' and $(f \oplus 1) \cdot p = 0$. Note that the rational degree of f is $d = \min(\rho, \rho')$.

With the notion of primal and dual rational degree we can distill our main observation in the following lemma, which we prove in the full version.

Lemma 11. Let $f, h: \{0, 1\}^n \rightarrow \{0, 1\}$ be Boolean functions that depend on disjoint parts of the input⁵, where f and h have primal rational degree ρ_f and ρ_h and dual rational degree ρ'_f and ρ'_h , respectively. Then:

- (i) The primal rational degree of $f \vee h$ is lower bounded by $\rho_f + \rho_h$.
- (ii) The dual rational degree ρ' of $f \vee h$ is lower bounded by $\min(\rho'_f, \rho'_h)$.

With this, it is straightforward to compute the exact rational degree of a disjunction, where all terms depend on disjoint parts of the input. Similarly, we can also apply this to compute the rational degree of a conjunction, since $\bigwedge_{i=1}^s f_i = \bigvee_{i=1}^s (f_i \oplus 1) \oplus 1$.

Put together, and applied to our candidate, we obtain the following.

Lemma 12. Let $m = m(\lambda) \in \mathbb{N}$, let $g: \{0, 1\}^m \rightarrow \{0, 1\}$ be as in Eq. 1, let $n = m + 1$, and let $s \in \{0, 1\}^n, K \in \{0, 1\}^{m \times n}$ be such that the map $x \mapsto (\langle x, s \rangle, K \cdot x) \pmod 2$ is invertible. Then, our candidate weak PRF $f_{s,K}: \{0, 1\}^n \rightarrow \{0, 1\}$ defined via

$$f(x) \mapsto \langle x, s \rangle + g(K \cdot x \pmod 2)$$

has rational degree at least λ .

For further details and discussion, we refer the reader to the full version.

⁵ We say that f depends on the i -th index of the input, if x_i appears with a non-zero coefficient in some term in f .




Fig. 1. Relation between different assumptions/ conjectures. $A \rightarrow B$ means that A implies B . By a linear distribution we mean the uniform distribution over a linear subspace $V \subseteq \{0, 1\}^n$, where dotted implications were already observed by [1].

4.2 On Resistance to Linear Attacks

We also consider the resistance of our candidate to linear attacks, as was done for our other candidate in Sect. 3. While we have not been able to prove resistance of linear attacks for this candidate, we formulate a combinatorial conjecture which states, informally, that if the deterministic noise function is c -unbalanced for some constant c and far from all low-degree polynomials, then no attack from the linear attack framework can break the corresponding LPN with simple noise assumption. If true, this conjecture would imply that our candidate, the ABGKR candidate, as well as the “LWR mod 6” candidate from [14], cannot be broken by any of the above attacks. We provide preliminary observations regarding the plausibility of the conjecture; we view proving or disproving this conjecture as an interesting open question. For more details we refer to the full version

4.3 On WPRFs in AC0 with Public Parities

In this work we give a candidate construction of a weak PRF in $AC0 \circ MOD2$, where the parities are *secret*. In particular, we conjecture that such a weak PRF exists (this is in the following referred to as Conjecture I).

We also consider the natural question of the existence of a simpler class of WPRF of the form $f_k(x) = g_k(G \cdot x)$, where G is a *public* matrix. Note that if G is removed (or surjective) then f_k could be learned by the algorithm of Linial, Mansour and Nisan [41] for learning AC0 under the uniform distribution.

While Akavia et al. [1] conjectured that any function in $\text{AC0} \circ \text{MOD2}$ has a large Fourier coefficient, we take this further by suggesting that, in the case of a public matrix G , the heavy Fourier coefficient of f_k stems from a low-order coefficient of g_k (in the following referred to as Conjecture II). This would imply that the high-weight Fourier coefficient can be used to distinguish the function from random in quasipolynomial time even given only access to random samples, and therefore allows to conclude that there cannot exist a weak PRF in AC0 on top of *public* parities (in the following referred to as Conjecture III).

We prove Conjecture II for the case when g_k is a family of DNFs, by extending the work of Jackson [36] to show that the coefficient is of the right form. The idea of Jackson is that any DNF correlates with a parity of its term that is “most likely” to be satisfied, which implies a heavy Fourier coefficient. We further observe that this means the function is either biased, or the term can contain only a few non-correlated variables. Since an AND clause is only satisfied for exactly one setting of inputs, if there are too many independent terms in the DNF then the function is biased. Otherwise, there are many dependencies between the individual terms, which we show implies the heavy Fourier coefficient comes from a vector of the form $a = G^\top v$ for some low-weight v .

We further prove Conjecture II for arbitrary $g_k \in \text{AC0}$ if the matrix G is random (or, more generally, defines a polylog-wise independent map).

We present the formal Conjectures I, II and III as well as the proof of Conjecture II for the above mentioned special cases in the full version.

Linear IPPP and Relations Between Conjectures. Finally, in the full version, we also elaborate on the relations between our conjectures, and previous conjectures in the literature including the “Linear IPPP” conjecture [56], asserting that mod-2 inner product is not in $\text{AC0} \circ \text{MOD2}$. These connections are illustrated in Fig. 1.

5 Between Lapland and Cryptomania

In this section we present two abstract frameworks. We first introduce the notion of learning parity with simple deterministic noise, which captures our candidate weak PRF in $\text{AC0} \circ \text{MOD2}$ from Sect. 4. Further, we show that every weak PRF candidate in $\text{AC0}[\text{MOD2}]$ implies some form of learning parity with simple deterministic noise.

Next, we introduce an abstract framework that captures variable-density learning parity with noise style assumptions such as the candidate weak PRF of [15] and our candidate weak PRF from Sect. 3.

Further, if one believes that either no code is hard to decode or almost all codes are hard to decode with respect to some noise level, then we show that each

candidate that fits into the VDLPN framework lives in Cryptomania. We are not aware of any similar implications for functions that can be cast as learning parity with $\text{AC0}[\text{MOD2}]$ -noise more generally.

5.1 Learning Parity with Simple Deterministic Noise

We observe that the Akavia et al. [1] candidate as well as our own candidate in $\text{AC0} \circ \text{MOD2}$ can be cast as a form of new LPN-style assumption, that we refer to as *LPN with simple deterministic noise*. This can be viewed as a generic method to transform a biased weak PRF into a weak PRF. Formally, we define learning parity with simple noise as follows.

Definition 13 (Learning parity with simple deterministic noise). *Let $n = n(\lambda), \kappa = \kappa(\lambda) \in \mathbb{N}$ and let $\mathcal{G} = \{g_k: \{0, 1\}^n \rightarrow \{0, 1\} \mid k \in \{0, 1\}^\kappa\}$ be a family of keyed functions in a low-complexity class. We say a function family $\mathcal{F} = \{f_{s,k}: \{0, 1\}^n \rightarrow \{0, 1\} \mid s \in \{0, 1\}^n, k \in \{0, 1\}^\kappa\}$ is an instance of learning parity with simple deterministic noise from \mathcal{G} , if $f_{s,k}: \{0, 1\}^n \rightarrow \{0, 1\}$ is of the form $f_{s,k}(x) = \langle x, s \rangle \oplus g_k(x)$.*

In this paper by simple we usually refer to noise functions in $\text{AC0}[\text{MOD2}]$. Note that if \mathcal{G} is in $\text{AC0}[\text{MOD2}]$, then so is \mathcal{F} . Further note that $f_{s,k}$ can be written as

$$f_{s,k}(x) = (\neg \langle x, s \rangle \wedge g_k(x)) \vee (\langle x, s \rangle \wedge \neg g_k(x)).$$

This shows that for $g_k \in \text{AC0} \circ \text{MOD2}$ we also have $f_{s,k} \in \text{AC0} \circ \text{MOD2}$ (where we consider the key as fixed). Note that this transformation from a biased weak PRF g_k to a weak PRF $f_{s,k}$ is not depth-preserving, however.

This framework can be extended to capture more general input distributions as follows.

Definition 14 (Extension to general input distributions). *Let $n = n(\lambda), \kappa = \kappa(\lambda), M = M(\lambda) \in \mathbb{N}$, let $\mathcal{G} = \{g_k: \{0, 1\}^n \rightarrow \{0, 1\} \mid k \in \{0, 1\}^\kappa\}$ be a family of keyed functions in a low-complexity class, and let $h: \{0, 1\}^n \rightarrow \{0, 1\}^M$ a function. We say that a function family $\mathcal{F} = \{f_{s,k}: \{0, 1\}^n \rightarrow \{0, 1\} \mid s \in \{0, 1\}^M, k \in \{0, 1\}^\kappa\}$ is an instance of learning parity with simple deterministic noise from \mathcal{G} with respect to the input distribution generated by h , if $f_{s,k}: \{0, 1\}^n \rightarrow \{0, 1\}$ is of the form $f_{s,k}(x) = \langle h(x), s \rangle \oplus g_k(x)$.*

Of course not every class of noise functions gives rise to a candidate weak PRF. In the full version we make progress on studying learning parity with simple noise by presenting a combinatorial conjecture about properties that the family of noise functions \mathcal{G} has to satisfy (informally speaking, these are the properties of being “balanced” and having “high-degree”), that we believe are sufficient in order to resist all linear attacks. However, note that as the attack by Bogdanov and Rosen [13] showed, satisfying these properties is still not sufficient to be a weak PRF, because other classes of attacks such as algebraic attacks might apply.

5.2 Weak PRFs in $\text{AC0}[\text{MOD2}]$ Live in Lapland

The results on circuit lower bounds by Razborov and Smolensky [53, 57] show that every function in $\text{AC0}[\text{MOD2}]$ can be approximated by a polynomial of poly-logarithmic degree. More formally, their result can be stated as follows.

Theorem 15 (Razborov-Smolensky [53, 57]). *Let $n, d, S \in \mathbb{N}$. If $f: \{0, 1\}^n \rightarrow \{0, 1\}$ can be computed by depth- d , size- S circuit with MOD2 gates, then for any integer $\varepsilon > 0$, there exists a polynomial $p(x) \in \mathbb{F}_2[x_1, \dots, x_n]$ of degree at most $(\log(S/\varepsilon))^d$ such that $\Pr_x[f(x) \neq p(x)] \leq \varepsilon$.*

The theorem implies that if there is a weak PRF in $\text{AC0}[\text{MOD2}]$, then learning parity with noise is hard, where the code is a “punctured” Reed-Muller code of quasipolynomial dimension (i.e. the row corresponding to an input x consists of the of all low-degree monomials evaluated on x), the secret corresponds to the coefficient of the polynomial that approximates the weak PRF, and the noise corresponds to the approximation error. In other words, the existence of a weak PRF in $\text{AC0}[\text{MOD2}]$ says that this kind of punctured Reed-Muller codes are hard to decode for some nontrivial noise rate.

Corollary 16. *Let $n = n(\lambda), \kappa = \kappa(\lambda) \in \mathbb{N}$. If there exists a (Q, T, ε) -weak PRF in $\text{AC0}[\text{MOD2}]$, then there exists $c, C \in \mathbb{N}$ with $c < C$, a family of keyed functions $\mathcal{G} = \{g_k: \{0, 1\}^n \rightarrow \{0, 1\} \mid k \in \{0, 1\}^\kappa\}$ and a function $h: \{0, 1\}^n \rightarrow \{0, 1\}^M$ where $M = 2^{\log^C \kappa}$, such that the learning parity function $f_{s,k} = \langle h(x), s \rangle \oplus g_k(x)$ with deterministic noise $g_k \in \mathcal{G}$ respective to the input distribution generated by h is a $(Q, \mathcal{O}(T), \varepsilon)$ -weak PRF. Further, for the corresponding noise rate we have that $\Pr_{x,k}[g_k(x) = 1] \leq 2^{-\log^c \kappa}$.*

Note that the Razborov-Smolensky result does not make any guarantees as to the *distribution* over the approximating low-degree polynomial for the functions in the PRF family, corresponding to distribution over the secret s in the LPN instance. However, the corresponding LPN instance reduces to the case of average-case s . Namely, samples $\langle x, s^* \rangle \oplus g_k(s)$ for arbitrary s^* can be generically converted to consistent samples for uniform secret $s^* + s'$, by offsetting each sample by $\langle x, s' \rangle$.

Note that having a superpolynomial secret in Corollary 16 only “scales down” the LPN security when expressed as a function of the secret size, and in the subexponential regime the resulting guarantee remains meaningful. More explicitly, the corollary can be understood as follows: If there exists a weak PRF in $\text{AC0}[\text{MOD2}]$ with subexponential security 2^{κ^δ} , then there exists an instance of deterministic LPN that has secret length $M = 2^{\log^C \kappa}$ and security in the order of $2^{\kappa^\delta} = 2^{2^{\delta \cdot \log^{1/C} M}}$. Thus, the existence of weak PRF candidates in $\text{AC0}[\text{MOD2}]$ with subexponential security implies what can be viewed as an instance of deterministic LPN with “subsubexponential hardness” in the secret length (which lies strictly between quasipolynomial and subexponential).

Consider a hardness of decoding interpretation of Corollary 16. Observe that the noise rate ε implied by Razborov-Smolensky is *above* the minimal distance

of the corresponding (punctured) Reed-Muller code of low-degree multivariate polynomials, therefore unique decoding will in general not be possible. That is, we expect many low-degree multivariate polynomials $p(x)$ to agree with a given function f_k in $\text{AC0}[\text{MOD2}]$ up to this noise rate. Identifying *any* such $p(x)$ constitutes an attack on the pseudorandomness of f_k , as it provides a low-error prediction of f_k evaluations. Note that the *number* of (punctured) Reed-Muller codewords within this distance is bounded: in particular, for $Q = 2^{\kappa^\delta}$, the probability that a random word in the space $\{0, 1\}^Q$ will be within Hamming distance $\Delta = 2^{\kappa^\delta - \log^c \kappa}$ of a codeword will be negligible. Thus, we can conclude that the existence of a weak PRF in $\text{AC0}[\text{MOD2}]$ implies that the punctured Reed-Muller code is hard to decode in some non-unique decoding regime. We formalize this in the following corollary.

Corollary 17. *Suppose for every $c, C \in \mathbb{N}$ with $c \leq C$, there is an algorithm \mathcal{A} running in time $2^{n^{o(1)}}$ such that, given a generating matrix G of a punctured RM code over \mathbb{F}_2 with parameters $(\log^C n, n)$ and corrupted codeword y , \mathcal{A} finds a codeword which is within relative distance $2^{-\log^c n}$ from y . Then there are no WPRFs in $\text{AC0}[\text{MOD2}]$.*

While it is known that the decoding of some linear codes and even structured codes such as Reed-Solomon codes for certain noise rates is NP-hard [11, 29], we are not aware of similar result for (punctured) Reed-Muller codes as the one described above. Also, to our knowledge known results on NP-hardness of computing and approximating the minimum distance of codes [25, 59] do not apply to our example. We leave it as an interesting open question to find a more natural implication from weak PRFs in $\text{AC0}[\text{MOD2}]$ to the hardness of decoding linear codes.

5.3 A Framework for VDLPN Assumptions

In [15], a candidate weak PRF in $\text{AC0}[\text{MOD2}]$ was given, with security based on a specific variable-density learning parity with noise assumption. In the following we give a framework of variable-density learning parity with noise that captures the weak PRF candidate of [15] and also our candidate based on sparse polynomials presented in Sect. 3 in $\text{AC0}[\text{MOD2}]$. Note that the VDLPN framework is not restricted to functions in $\text{AC0}[\text{MOD2}]$. And, on the other hand, not all function families in $\text{AC0}[\text{MOD2}]$ fall within this framework. Therefore, the conditional public-key implication that we give in the following only applies to candidates such as the one given in [15] and our candidate based on sparse polynomials, but not our candidate weak PRF in $\text{AC0} \circ \text{MOD2}$.

Definition 18 (A framework for VDLPN). *Let $n = n(\lambda), N = N(\lambda), \kappa = \kappa(\lambda) \in \mathbb{N}$. Let $h: \{0, 1\}^n \rightarrow \{0, 1\}^N$ and $e: \{0, 1\}^\kappa \rightarrow \{0, 1\}^N$. We say that (h, e) defines an instance of variable-density learning party with noise, if $f_k(x) := \langle h(x), e(k) \rangle$ is efficiently computable, and there exist $\zeta_i = \zeta_i(\lambda), \eta_i = \eta_i(\lambda) \in [0, 1]$ for all $i \in [N]$, such that*

1. for all $i \in [N]$ it holds: $\Pr_x[h(x)_i = 1] = \zeta_i$ and $\Pr_k[e(k)_i = 1] = \eta_i$,
2. for all $i \in [N]$ it holds: $\zeta_i \geq \zeta_{i+1}$ and $\eta_i \geq \eta_{i+1}$,
3. there exist polynomials $p = p(\lambda), q = q(\lambda) \in \mathbb{N}$ such that: $\sum_{i=1}^N \zeta_i \leq p$ and $\sum_{i=1}^N \eta_i \leq q$.

We say that the variable-density learning parity with noise (VDLPN) assumption with respect to (h, e) is (Q, T, ϵ) -hard, if $f_k(x) := \langle h(x), e(k) \rangle$ is a (Q, T, ϵ) -weak PRF.

Note that – even though not directly falling into the framework of learning parity with simple noise – VDLPN implies an instance thereof. To see this consider a VDLPN tuple (h, e) . Now, let $h_0: \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $h_1: \{0, 1\}^n \rightarrow \{0, 1\}^{N-n}$ such that $h(x) = (h_0(x), h_1(x))$ for all $x \in \{0, 1\}^n$, and similarly let $e_0: \{0, 1\}^\kappa \rightarrow \{0, 1\}^n, e_1: \{0, 1\}^\kappa \rightarrow \{0, 1\}^{N-n}$, such that $e(k) = (e_0(k), e_1(k))$ for all $k \in \{0, 1\}^\kappa$. Let $\mathcal{G} = \{g_k: \{0, 1\}^n \rightarrow \{0, 1\} \mid k \in \{0, 1\}^\kappa\}$, where $g_k(x) = \langle h_1(x), e_1(k) \rangle$, and let $f_{s,k} = \langle h_0(x), s \rangle \oplus g_k(x)$. Now, if VDLPN with respect to (h, e) is hard, then so is learning parity with simple deterministic noise \mathcal{G} with respect to the input distribution generated by h_0 , due to the same reduction of LPN with arbitrary secret s^* to a uniform secret s' mentioned in a comment following Corollary 16.

5.4 Connections of VDLPN to Cryptomania

In the following we outline why VDLPN “morally” implies LPN with low noise and therefore public-key encryption. We cannot show a direct PKE implication, because the Alekhovich construction [3] does not apply directly if the matrix is also sparse, since the dual LPN assumption (i.e. the assumption that the pair (H, v) for a matrix H that generates the dual code and $v = H \cdot e$ for a sparse noise vector e is indistinguishable from (H, r) for a uniformly random vector r) cannot hold true in this case, as v will be biased towards 0.

What we mean by “morally” is that the noise rate itself is sufficiently low to imply PKE, and because typically LPN is considered to be hard on average for random codes (if the noise is sufficiently dense). In order to formalize this observation we formulate a conjecture stating that if there exists a code that is hard to decode with respect to some noise rate (where the noise itself can depend on the generator matrix of the code), then “almost all” codes are hard to decode with respect to this noise rate. In order to deal with the fact that the noise might depend on the matrix (and therefore replacing the matrix might in fact trivially render LPN insecure), we simultaneously replace the noise by noise that is Bernoulli distributed at the same rate.

Conjecture 19 (Random LPN is the hardest). Let $n = n(\lambda), Q = Q(\lambda), \kappa = \kappa(\lambda), M = M(\lambda) \in \mathbb{N}$, let $\mathcal{G} = \{g_k: \{0, 1\}^n \rightarrow \{0, 1\} \mid k \in \{0, 1\}^\kappa\}$ be a family of keyed functions and let $h: \{0, 1\}^n \rightarrow \{0, 1\}^M$ such that learning parity with simple noise \mathcal{G} is (Q, T, ϵ) -hard for the input distribution generated by h . Then, we conjecture that the standard LPN problem with noise with rate η is (Q, T, ϵ) -hard. More precisely, we conjecture that if $A \stackrel{\$}{\leftarrow} \{0, 1\}^{Q \times M}, s \stackrel{\$}{\leftarrow} \{0, 1\}^M$ are both

sampled uniformly at random, and a noise vector e is sampled according to the Bernoulli distribution over $\{0, 1\}^Q$ with rate $\eta \geq \Pr_{x \xleftarrow{\$} \mathcal{D}, k \xleftarrow{\$} \{0, 1\}^\kappa} [g_k(x) = 1]$, then there exists a constant $c > 0$ such that the distribution of $(A, As + f)$ is $(T, \epsilon + 2^{-\lambda^c})$ -indistinguishable from the uniform distribution.

Note that relaxing the success probability of the adversary to $\epsilon + 2^{-\lambda^c}$ is necessary, because there obviously exist some codes that are easy to distinguish from random for any non-trivial noise rate (e.g. A chosen as the all zero matrix).

In order to further weaken the conjecture, allowing for the possibility that there exist some codes that are significantly harder to decode than random codes, one can require that the input the generated by h (i.e. obtained by sampling $x \xleftarrow{\$} \{0, 1\}^n$ and outputting $h(x)$), have min-entropy at least $\text{polylog}(\lambda)$. This weaker conjecture is still sufficient to prove the PKE implication of VDLPN.

In the full version, we prove the following.

Lemma 20. *Let $n = n(\lambda), N = N(\lambda), \kappa = \kappa(\lambda) \in \mathbb{N}$, $h: \{0, 1\}^n \rightarrow \{0, 1\}^N$ and $e: \{0, 1\}^\kappa \rightarrow \{0, 1\}^N$. Let $T = T(\lambda) \in \mathbb{N}$ and $Q = Q(\lambda)$ such that $Q \in \lambda^{\omega(1)}$. Then, if Conjecture 19 holds and VDLPN is $(Q, 2^{\lambda^c}, 2^{-\lambda^c})$ -hard for (h, e) for some constant $c > 0$, then public-key encryption with quasipolynomial running time and subexponential security exists.*

Remark 21. Note that the noise rate implied by Razborov-Smolensky does not suffice to construct public-key encryption via Alekhovich [3] (even under the “random LPN is the hardest” conjecture), because the noise rate implied by the Razborov-Smolensky approximation is $\omega(1/\sqrt{M})$. In addition, constructions of PKE from LPN with constant noise, e.g., [63], have quasi-polynomial running time and security. We are therefore not aware of any public-key implications for general weak PRFs in $\text{AC0}[\text{MOD}2]$.

Acknowledgements. We thank Andrej Bogdanov, Nicolas Resch, and the anonymous Crypto reviewers for helpful discussions and suggestions.

E. Boyle supported by ISF grant 1861/16, AFOSR Award FA9550-21-1-0046, a Google Research Award, and ERC Project HSS (852952). G. Couteau supported by the ANR SCENE. N. Gilboa supported by ISF grant 2951/20, ERC grant 876110, and a grant by the BGU Cyber Center. Y. Ishai supported by ERC Project NTSC (742754), NSF-BSF grant 2015782, BSF grant 2018393, and ISF grant 2774/20. L. Kohl is funded by NWO Gravitation project QSC. Research of L. Kohl was done in part while at Technion, supported by ERC Project NTSC (742754). P. Scholl supported by the Danish Independent Research Council under Grant-ID DFF-6108-00169 (FoCC) and an Aarhus University Research Foundation starting grant.

References

1. Akavia, A., Bogdanov, A., Guo, S., Kamath, A., Rosen, A.: Candidate weak pseudorandom functions in $AC^0 \circ \text{MOD}_2$. In: ITCS 2014. ACM, January 2014
2. Albrecht, M.R., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 430–454. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_17
3. Alekhovich, M.: More on average case vs approximation complexity. In: 44th FOCS. IEEE Computer Society Press, October 2003
4. Applebaum, B.: Pseudorandom generators with long stretch and low locality from random local one-way functions. *SIAM J. Comput.* **42**(5), 2008–2037 (2013)
5. Applebaum, B., Lovett, S.: Algebraic attacks against random local functions and their countermeasures. *SIAM J. Comput.* **47**(1), 52–79 (2018)
6. Applebaum, B., Raykov, P.: Fast pseudorandom functions based on expander graphs. In: Hirt, M., Smith, A. (eds.) TCC 2016, Part I. LNCS, vol. 9985, pp. 27–56. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_2
7. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011. LNCS, vol. 6755, pp. 403–415. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22006-7_34
8. Ball, M., Holmgren, J., Ishai, Y., Liu, T., Malkin, T.: On the complexity of decomposable randomized encodings, or: how friendly can a garbling-friendly PRF be? In: ITCS 2020. LIPIcs, January 2020
9. Banerjee, A., Peikert, C., Rosen, A.: Pseudorandom functions and lattices. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 719–737. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_42
10. Barkol, O., Ishai, Y.: Secure computation of constant-depth circuits with applications to database search problems. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 395–411. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_24
11. Berlekamp, E., McEliece, R., Van Tilborg, H.: On the inherent intractability of certain coding problems (Corresp.). *IEEE Trans. Inf. Theory* **24**(3), 384–386 (1978)
12. Blum, A., Furst, M., Kearns, M., Lipton, R.J.: Cryptographic primitives based on hard learning problems. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 278–291. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_24
13. Bogdanov, A., Rosen, A.: Pseudorandom functions: three decades later. *Cryptology ePrint Archive*, Report 2017/652 (2017). <http://eprint.iacr.org/2017/652>
14. Boneh, D., Ishai, Y., Passelègue, A., Sahai, A., Wu, D.J.: Exploring crypto dark matter: new simple PRF candidates and their applications. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part II. LNCS, vol. 11240, pp. 699–729. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03810-6_25
15. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Correlated pseudorandom functions via variable-density LPN. In: FOCS (2020)
16. Braverman, M.: Polylogarithmic independence fools AC^0 circuits. *J. ACM (JACM)* **57**(5), 1–10 (2008)
17. Canteaut, A., et al.: Stream ciphers: a practical solution for efficient homomorphic ciphertext compression. In: Peyrin, T. (ed.) FSE 2016. LNCS, vol. 9783, pp. 313–333. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-52993-5_16

18. Carmosino, M.L., Impagliazzo, R., Kabanets, V., Kolokolova, A.: Learning algorithms from natural proofs. In: CCC 2016, pp. 10:1–10:24 (2016)
19. Cheraghchi, M., Grigorescu, E., Juba, B., Wimmer, K., Xie, N.: Ac^0 mod₂ lower bounds for the Boolean inner product. *J. Comput. Syst. Sci.* **97**, 45–59 (2018)
20. Courtois, N.T.: The security of hidden field equations (HFE). In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 266–281. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45353-9_20
21. Courtois, N.T.: Fast algebraic attacks on stream ciphers with linear feedback. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 176–194. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_11
22. Courtois, N.T., Meier, W.: Algebraic attacks on stream ciphers with linear feedback. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 345–359. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-39200-9_21
23. Couteau, G., Meyer, P.: Breaking the circuit size barrier for secure computation under quasi-polynomial LPN. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12697, pp. 842–870. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77886-6_29
24. Daniely, A., Vardi, G.: From local pseudorandom generators to hardness of learning. arXiv preprint [arXiv:2101.08303](https://arxiv.org/abs/2101.08303) (2021)
25. Dumer, I., Micciancio, D., Sudan, M.: Hardness of approximating the minimum distance of a linear code. *IEEE Trans. Inf. Theory* **49**(1), 22–37 (2003)
26. Feldman, V., Gopalan, P., Khot, S., Ponnuswami, A.K.: On agnostic learning of parities, monomials, and halfspaces. *SIAM J. Comput.* **39**(2), 606–645 (2009)
27. Filmus, Y., Ishai, Y., Kaplan, A., Kindler, G.: Limits of preprocessing. In: CCC 2020 (2020)
28. Furst, M., Saxe, J.B., Sipser, M.: Parity, circuits, and the polynomial-time hierarchy. *Math. Syst. Theory* **17**(1), 13–27 (1984)
29. Gandikota, V., Ghazi, B., Grigorescu, E.: On the np-hardness of bounded distance decoding of Reed-Solomon codes. In: 2015 IEEE International Symposium on Information Theory (ISIT), pp. 2904–2908. IEEE (2015)
30. Goldreich, O.: Candidate one-way functions based on expander graphs. *Cryptology ePrint Archive*, Report 2000/063 (2000). <http://eprint.iacr.org/2000/063>
31. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: 25th FOCS. IEEE Computer Society Press, October 1984
32. Håstad, J.T.: *Computational Limitations for Small-depth Circuits*. MIT Press, Cambridge (1987)
33. Håstad, J.: Almost optimal lower bounds for small depth circuits. In: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing, pp. 6–20 (1986)
34. Hellerstein, L., Servedio, R.A.: On PAC learning algorithms for rich Boolean function classes. *Theor. Comput. Sci.* **384**(1), 66–76 (2007). <https://doi.org/10.1016/j.tcs.2007.05.018>
35. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Cryptography with constant computational overhead. In: STOC 2008, pp. 433–442 (2008)
36. Jackson, J.C.: An efficient membership-query algorithm for learning DNF with respect to the uniform distribution. *J. Comput. Syst. Sci.* **55**(3), 414–440 (1997)
37. Johansson, T., Jönsson, F.: Improved fast correlation attacks on stream ciphers via convolutional codes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 347–362. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_24
38. Kearns, M.J., Valiant, L.G.: Cryptographic limitations on learning Boolean formulae and finite automata. *J. ACM* **41**(1), 67–95 (1994)

39. Kharitonov, M.: Cryptographic hardness of distribution-specific learning. In: 25th ACM STOC. ACM Press, May 1993
40. Krause, M., Lucks, S.: On the minimal hardware complexity of pseudorandom function generators. In: Ferreira, A., Reichel, H. (eds.) STACS 2001. LNCS, vol. 2010, pp. 419–430. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44693-1_37
41. Linial, N., Mansour, Y., Nisan, N.: Constant depth circuits, Fourier transform, and learnability. In: 30th FOCS. IEEE Computer Society Press, October/November 1989
42. Méaux, P., Journault, A., Standaert, F.-X., Carlet, C.: Towards stream ciphers for efficient FHE with low-noise ciphertexts. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part I. LNCS, vol. 9665, pp. 311–343. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_13
43. Miles, E., Viola, E.: Substitution-permutation networks, pseudorandom functions, and natural proofs. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 68–85. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_5
44. Mossel, E., Shpilka, A., Trevisan, L.: On e-biased generators in $NC0$. In: 44th FOCS. IEEE Computer Society Press, October 2003
45. Naor, J., Naor, M.: Small-bias probability spaces: efficient constructions and applications. *SIAM J. Comput.* **22**(4), 838–856 (1993)
46. Naor, M., Reingold, O.: Synthesizers and their application to the parallel construction of pseudo-random functions. In: 36th FOCS. IEEE Computer Society Press, October 1995
47. Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. In: 38th FOCS. IEEE Computer Society Press, October 1997
48. Naor, M., Reingold, O., Rosen, A.: Pseudo-random functions and factoring (extended abstract). In: 32nd ACM STOC. ACM Press, May 2000
49. Nisan, N., Wigderson, A.: Hardness vs. randomness (extended abstract). In: 29th FOCS. IEEE Computer Society Press, October 1988
50. O’Donnell, R.: *Analysis of Boolean Functions*. Cambridge University Press, Cambridge (2014)
51. Orlandi, C., Scholl, P., Yakoubov, S.: The rise of Paillier: homomorphic secret sharing and public-key silent OT. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12696, pp. 678–708. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_24
52. Patarin, J.: Cryptanalysis of the Matsumoto and Imai public key scheme of Eurocrypt’88. In: Coppersmith, D. (ed.) CRYPTO 1995. LNCS, vol. 963, pp. 248–261. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-44750-4_20
53. Razborov, A.A.: Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Math. Notes Acad. Sci. USSR* **41**(4), 333–338 (1987)
54. Razborov, A.A., Rudich, S.: Natural proofs. In: 26th ACM STOC. ACM Press, May 1994
55. Rossman, B., Servedio, R.A., Tan, L.Y.: An average-case depth hierarchy theorem for Boolean circuits. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, pp. 1030–1048. IEEE (2015)
56. Servedio, R.A., Viola, E.: On a special case of rigidity. In: *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 19, p. 144. Citeseer (2012)
57. Smolensky, R.: Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In: 19th ACM STOC. ACM Press, May 1987

58. Valiant, L.G.: A theory of the learnable. *Commun. ACM* **27**(11), 1134–1142 (1984)
59. Vardy, A.: The intractability of computing the minimum distance of a code. *IEEE Trans. Inf. Theory* **43**(6), 1757–1766 (1997)
60. Viola, E.: The communication complexity of addition. In: 24th SODA. ACM-SIAM, January 2013
61. Williams, R.: Natural proofs versus derandomization. In: 45th ACM STOC. ACM Press, June 2013
62. Yu, Yu., Steinberger, J.: Pseudorandom functions in almost constant depth from low-noise LPN. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 154–183. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_6
63. Yu, Yu., Zhang, J.: Cryptography with auxiliary input and trapdoor from constant-noise LPN. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 214–243. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_9