Strong Pseudorandom Functions in $\mathsf{AC}^0[2]$ in the Bounded-Query Setting

Marshall Ball * Clément Ducros † Saroja Erabelli ‡ Lisa Kohl §
Nicolas Resch ¶

November 13, 2025

Abstract

Understanding the minimal computational power needed to realize a pseudorandom function (PRF) is a long-standing question in cryptography. By the Razborov–Smolensky polynomial approximation method, it is known that $\mathsf{AC}^0[2]$ cannot support strong pseudorandom functions with subexponential security, since any such function can be distinguished from random with quasipolynomially many samples.

In this work, we initiate the study of low-complexity strong PRFs under a refined framework that separates adversary query complexity from running time, and observe that distinguishing algorithms for $\mathsf{AC}^0[2]$ do not apply if the number of queries is below the threshold implied by the Razborov–Smolensky approximation bound.

We propose the first candidate strong PRF in $\mathsf{AC}^0[2]$, which plausibly offers subexponential security against adversaries limited to a fixed quasipolynomial number of queries. We show that our candidate lacks heavy Fourier coefficients, resists a natural class of adaptive attacks, has high rational degree, is non-sparse over $\mathbb{F}_2$ in expectation, and has low correlation with fixed function families.

Finally, we show that if any strong PRF exists in $\mathsf{AC}^0[2]$ (or a superclass), then we can construct a universal PRF, i.e., a single, fixed function which is guaranteed to be a strong PRF in the same class.

1 Introduction

A pseudorandom function (PRF) is a family of efficiently computable functions that, when keyed with a secret key, is indistinguishable from a truly random function to any efficient adversary. In this work we focus on standard (also referred to as *strong*) PRFs, where the adversary gets chosen-query access to the oracle. The question of how *computationally simple* a PRF can be is a fundamental question at the intersection of cryptography, complexity theory, and learning theory. In modern cryptography, PRFs not only enable direct constructions of encryption and authentication schemes, but they are also central for the efficiency of high-end cryptographic protocols, where a PRF is evaluated within some secure computation protocol [ARS+15, MJSC16, CCF+16, BIP+18, BHI+20, BCG+20]. Furthermore, understanding the minimal computational complexity required to realize PRFs sheds light on fundamental questions in complexity theory. Specifically, if a PRF exists in a complexity class, then any proof technique meeting the Razborov-Rudich "natural proof" criteria cannot be used to establish circuit lower bounds in that class [RR94, MV12, RST15]. Additionally, PRFs help to understand when and how derandomization techniques can be applied [NW88, Wil13]. Moreover, the existence of PRFs is closely tied to questions in computational learning theory. In particular, any complexity class that admits learning with membership queries cannot contain secure PRFs. Constructing PRFs from assumptions that are simple to state provides a concrete target for learning algorithms, offering a more tangible challenge than addressing the full generality of learning an entire complexity class.

*New York University, marshall.ball@cs.nyu.edu
†CISPA, Helmholtz Center for Information Security, clement.ducros@cispa.de
‡New York University, saroja.erabelli@gmail.com
§Cryptology Group, CWI Amsterdam, lisa.kohl@cwi.nl
¶University of Amsterdam, n.a.resch@uva.nl

To be considered secure, a PRF is typically required to admit at least subexponential security, i.e., there exists a constant $c > 0$ such that all adversaries running in time $T = 2^{\kappa^c}$ have distinguishing advantage at most $1/T$, where κ denotes the key length. While results in learning theory usually do not distinguish between the number of samples and the computational time required to *learn* a function, in cryptography these resources are often very different: an adversary may have access to only a small number of queries, while having access to substantial total computational resources to break a pseudorandom function. Motivated by this, we initiate the systematic study of low-complexity PRFs in the bounded-query setting, where the adversary has access to a quasipolynomial number of queries (i.e., $\kappa^{\log^c \kappa}$ for some fixed constant $c > 0$), but may run in subexponential time.

At the lower end of the complexity spectrum, it is known that any function in AC^0 (the class of constant-depth, polynomial-size circuits over AND, OR and NOT gates with unbounded fan-in) can be distinguished from a random function with constant advantage using only two queries [LMN89]. The next class up is $\mathsf{AC}^0[2]$, which augments AC^0 with XOR gates. Here, it is known that any function can be distinguished from random [Raz87, Smo87, KL01] and even learned [CIKK16] in quasipolynomial time. However, this crucially requires access to a sufficiently large quasipolynomial number of queries. In this work, we propose the following question:

Does there exist a PRF in $\mathsf{AC}^0[2]$ that is plausibly secure against any adversary running in time $T = 2^{\kappa^c}$ using $\kappa^{\log^c \kappa}$ chosen queries for some fixed $c > 0$?

We note that this in particular captures any (a-priori unbounded) polynomial number of queries. We decided to explicitly phrase the setting as quasipolynomial to avoid ambiguity with the setting where the queries are limited to a *fixed* polynomial number $q(\kappa)$, as this would be trivially achievable (even information-theoretically) via a family of $(q(\kappa) + 1)$-wise independent hash functions.

1.1 Our Contributions

Strong PRF Candidate in $\mathsf{AC}^0[2]$. We propose the first strong PRF candidate in $\mathsf{AC}^0[2]$ with plausible subexponential security when the adversary is restricted to quasipolynomially many queries. Our construction builds on the weak PRF candidate of Akavia et al. [ABG+14] (in the following referred to as "ABGKR14"), following a similar design rationale. Namely, the ABGKR14 candidate builds on the TRIBES function, which is a DNF of the form

$$\mathsf{TRIBES}_{m,n}(x) = \bigvee_{j=1}^{m} \bigwedge_{k=1}^{n} x_{j,k}$$

where $x \in \{0,1\}^M$ for $M := mn$, and $x_{j,k}$ refers to the $(n \cdot (j-1) + k)$-th index of x. The candidate weak PRF of ABGKR14 uses as its key a secret matrix $\mathbf{A} \in \{0,1\}^{(M+1)\times(M+1)}$, which transforms the input prior to evaluation by the TRIBES function. For suitable choices of m and n,[1] this candidate PRF achieves constant bias. To remove the bias, ABGKR14 augments the construction by adding the XOR-in the last entry of the transformed input to the TRIBES function. More precisely, their weak PRF candidate is of the following form:

$$\begin{aligned} &F_{\mathbf{A}} : \{0,1\}^{M+1} \to \{0,1\} \\ &F_{\mathbf{A}}(x) = g(\mathbf{A}x), \text{ where } g(y,z) = \mathsf{TRIBES}_{m,n}(y) \oplus z, \end{aligned} \tag{1}$$

where $\mathbf{A}x$ is taken modulo 2 and $\mathbf{A}x$ is parsed as $(y,z) := \mathbf{A}x$ for $y \in \{0,1\}^M$ and $z \in \{0,1\}$. While the candidate was later shown vulnerable to a quasipolynomial-time rational-degree attack [BR17], we note that it still plausibly withstands subexponential attacks in the bounded-sample setting. A natural idea is therefore to conjecture the hardness of ABGKR14 in the quasi-polynomial query setting where the queries are *chosen*. It is easy to see, however, that the ABGKR14 candidate can be broken even with a *constant* number of chosen queries. Firstly, it holds $F_{\mathbf{A}}(x) = 0$ for $x = 0$. Secondly, it was shown by ABGKR14 that their weak PRF candidate always has a heavy Fourier coefficient (that is, of weight $1/\kappa^{\log^c \kappa}$ for some $c > 0$), and is therefore learnable in quasi-polynomial time using the Kushilevitz–Mansour algorithm [KM91] in the chosen-sample setting. Furthermore, for distinguishing a function with a heavy Fourier coefficient from a random function, it suffices to query a *constant* number of chosen

[1] Namely, if $M = nm$ one should take $n \approx \log M$ and $m \approx M/\log M$; this guarantees $\mathsf{TRIBES}_{m,n}$ is roughly unbiased. See Section 3 for further discussion.

samples.[2] Thus, to achieve subexponential security in the quasipolynomial sample setting, it is crucial to avoid heavy Fourier coefficients.

Motivated by this state of affairs, we propose the following candidate strong PRF. The key of our PRF consists of a secret matrix $\mathbf{A} \in \{0,1\}^{N \times N}$, sampled uniformly at random subject to being invertible over $\mathbb{F}_2$, together with a secret vector $\mathbf{b} \in \{0,1\}^N$ chosen uniformly at random, where $N = \ell mn$.

$$\begin{aligned} &F_{\mathbf{A},\mathbf{b}} : \{0,1\}^N \to \{0,1\} \\ &F_{\mathbf{A},\mathbf{b}}(x) = g(\mathbf{A}x \oplus \mathbf{b}), \text{ where } g(\mathbf{y}) = \bigoplus_{i=1}^{\ell} \mathsf{TRIBES}_{m,n}(\mathbf{y}_i) \ , \end{aligned} \tag{2}$$

and $\mathbf{y}_i \in \{0,1\}^{mn}$ denotes the vector $\mathbf{y}_i := (y_{mn\cdot(i-1)}, \ldots, y_{mn\cdot i-1})$.

Conjecture 1.1 *For ℓ, m, n, such that $\ell = \Theta(\lambda)$, $m = \Theta(\lambda/\log\lambda)$ and $n = \Theta(\log\lambda)$ there exists a $c > 0$ such that for any adversary running in time $T = 2^{\kappa^c}$ getting access to $\kappa^{\log^c \kappa}$ chosen samples, the adversary has advantage at most $1/T$ in distinguishing $F_{\mathbf{A},\mathbf{b}}$ as defined in Equation 4 from random (where κ denotes the key length $\kappa = N^2 + N$ for $N = \ell mn = \Theta(\lambda^2)$).*

To the best of our knowledge it's feasible that the above holds for any $c < \frac{1}{4}$, i.e., security holds against adversaries running in time $T = 2^{\ell^{1-\epsilon}} = 2^{(mn)^{1-\epsilon}}$ for any $\epsilon > 0$. For the query complexity, $c < \frac{1}{4}$ might be an overly conservative choice based on the natural property of Razborov-Smolenksy. In fact, security up to $\kappa^{\log^2 \kappa}$ queries appears plausible. Lastly, while concrete cryptanalysis is not our focus, we mention that for $\lambda = 128$,[3] it is consistent with our results that an adversary running in time $T = 2^{128}$ could have success probability at most $1/T$ given access to 2^{30} chosen queries.

Design Rationale. Firstly, the secret vector $\mathbf{b}$ is necessary to make the PRF unpredictable on input $x = 0$. Secondly, based on a conjecture by Akavia et al [ABG+14] (later reiterated in a slightly stronger form [BCG+21]), all functions in $\mathsf{AC}^0 \circ \mathsf{MOD2}$ – which is a subclass of $\mathsf{AC}^0[2]$ where XOR gates are only allowed at the bottom (input) layer[4] – have a heavy Fourier coefficient. As discussed above, functions with heavy Fourier coefficients *cannot* be pseudorandom. In Section 4.1, we prove that the outer XOR gate rules out heavy Fourier coefficients.

Efficiency. One downside of our construction is the large key size $\Theta(N^2)$. Assuming subexponential hardness of the *learning parity with noise (LPN)* problem, we can apply the trapdoored construction of Vaikuntanathan and Zamir [VZ25] to decrease the key size and speed up evaluation. A naïve implementation (with unstructured $\mathbf{A}$) evaluates our PRF via naïve matrix–vector multiplication in time $\Theta(N^2)$, and the key size is $\Theta(N^2)$. With the LPN-based trapdoor of [VZ25], matrix–vector multiplication, and hence evaluation of our PRF, can be performed in $\tilde{O}(N)$ time with key size $\tilde{O}(N)$. Note that this speedup also applies to the strong and weak PRF candidates of [ABG+14], [BIP+18], and [BCG+21] whose keys include a random matrix.

Distributed Setting. While the complexity of PRFs is usually measured relative to a hardwired key, for advanced cryptographic primitives – where the key may be secret-shared between parties – it is often essential to consider complexity as a function of the key. We want to highlight that our candidate falls into $\mathsf{AC}^0[2]$ even if considered as a function of the input *and key*. Further, we observe that for certain settings of distributed computation the speed-up to $\tilde{O}(N)$ key size and evaluation time outlined in the previous paragraph carries over, as the transformation can be applied to each secret-share of the key individually (i.e., each matrix secret share can be replaced by a trapdoored LPN matrix).

Cryptanalysis. To strengthen our conjecture, we show that our candidate exhibits the following security properties:

- **Fourier-analytic resilience:** The function lacks heavy Fourier coefficients, resisting adaptive learning attacks that identify such coefficients [KM91]. We also generalize the linear test framework popular in code-based cryptography [CRR21] to include chosen (but non-adaptive) samples and show that heavy Fourier coefficient finding attacks fall within this framework.

[2] This fact is implicit in the literature [O'D14, Theorem 1.30], and we explicitly discuss it in Section 4.2.
[3] Here, we are imagining $\ell = \lambda$ and $mn = \lambda$, where m and n are chosen to make $\mathsf{TRIBES}_{m,n}$ nearly balanced.
[4] Note that XOR gates with constant fan-in can be emulated in AC^0 and are therefore allowed at all levels.

- **Resistance to a natural adaptive attack:** We demonstrate that the candidate withstands a basic adaptive strategy in which the adversary samples inputs $x^{(i)}$ uniformly at random, selectively sums those for which $F(x^{(i)}) = 0$, and then checks whether the resulting vector $y = \sum_{i:F(x^{(i)})=0} x^{(i)}$ also satisfies $F(y) = 0$.
- **Resistance to algebraic attacks:** As a byproduct of adding the outer XOR-layer in our candidate, we can show that the candidate has high rational degree $\Omega(\kappa^c)$ for any $c < 1/4$ and therefore provides security against rational-degree attacks using 2^{κ^c} queries (see Section 2.3 for a more detailed discussion). Note that this does not help to boost the PRF security in the bounded-query setting (as here, even ABGKR14 resists rational-degree attacks). However, unlike ABGKR14, our candidate PRF plausibly attains full subexponential security as a weak PRF – that is, it remains secure even against adversaries with access to a subexponential number of random input/output pairs (as distinguishing algorithms based on Razborov–Smolensky require access to chosen samples). A similar result was achieved in [BCG+21]. Their candidate weak PRF, however, falls in $\mathsf{AC}^0 \circ \mathsf{MOD2}$ and therefore does not provide any security guarantees in the chosen sample setting for the same reason as ABGKR14.
- **Non-sparsity over $\mathbb{F}_2$:** The function, in expectation, is not sparse when represented as a polynomial over $\mathbb{F}_2$, and thus, in expectation, evades property testing attacks that target sparse polynomials [DLM+07].
- **Preservation of prior guarantees:** Our construction retains resistance to the same set of attacks demonstrated in the work of Akavia et al., including low correlation with fixed function families and inapproximability by low-degree polynomials over $\mathbb{F}_2$.

Universal Strong PRF in $\mathsf{AC}^0[2]$. Furthermore, we show that if a strong PRF in $\mathsf{AC}^0[2]$ exists, then we can construct a *universal PRF* in $\mathsf{AC}^0[2]$, a single, fixed function in $\mathsf{AC}^0[2]$ which serves as a PRF. This is notable because it yields a single target function for building distinguishers against, which is potentially easier to obtain than a full natural property. Moreover, if such a universal distinguisher could be transformed into a natural property, then the absence of PRFs in $\mathsf{AC}^0[2]$ would imply that $\mathsf{AC}^0[2]$ is learnable, addressing the long-standing question on the dichotomy between the existence of PRFs and learning for this class.

1.2 Additional Context

On the Hardness of Learning $\mathsf{AC}^0[p]$ with Bounded Queries. Our candidate's conjectured security would imply the hardness of learning $\mathsf{AC}^0[2]$ circuits using a fixed quasipolynomial number of membership queries. We formally state this implication, extended to all primes p[5], in the conjecture below:

Conjecture 1.2 *There exists a $c > 0$, such that $\mathsf{AC}^0[p]$ circuits cannot be learned with $\kappa^{\log^c \kappa}$ many membership queries and 2^{κ^c} time.*

Limitations of Natural Property-Based Learners for $\mathsf{AC}^0[p]$. Razborov and Smolensky showed, in 1987, that all circuits in $\mathsf{AC}^0[p]$ of size m and depth d can be approximated by a $O((p-1)^d \log^d(m))$-degree polynomial with $1/\mathsf{poly}(m)$ approximation error [Raz87,Smo87]. When combined with a property tester that distinguishes approximate low degree $\mathbb{F}_p$ polynomials from random functions [KL01,BKS+10], this yields a natural property for the circuit class $\mathsf{AC}^0[p]$. Introduced by Razborov and Rudich [RR94], a *natural property* (or *proof*) for a class $\mathcal{C}$ is a distinguishing algorithm between a function in $\mathcal{C}$ and a randomly chosen function with non-trivial distinguishing advantage.

A well-known result in coding theory [JPRZ09] shows that for any degree $k \geq p-1$, any testing algorithm[6] that distinguishes between a degree k polynomial in $\mathbb{Z}_p$ and a function ϵ-far from a degree k polynomial requires at least $\Omega(\frac{1}{\epsilon} + p^{\frac{k+1}{p-1}})$ queries. This implies that if the best polynomial approximation for a function in $\mathsf{AC}^0[p]$ has degree $\log^d m$, then any natural property based on low-degree polynomial testing must incur query complexity at least $\Omega(p^{\log^d m})$, which is quasipolynomial. This lower bound

[5] That is, in an $\mathsf{AC}^0[p]$ circuit one may use MOD_p gates, which output 0 iff the number of input wires carrying a 1 is 0 mod p; else a 1.

[6] In fact, this lower bound applies to the setting where one is directly interacting with a low-degree polynomial, not just an approximate low-degree polynomial; we discuss this distinction further later.

illustrates that finding a natural property that uses only polynomially many queries would require fundamentally new techniques.

In the thirty-five years since the Razborov-Smolensky natural property was introduced, no other natural properties for $\mathsf{AC}^0[p]$ have been discovered, and no fundamentally new techniques have emerged for constructing them. Despite the existence of this longstanding barrier, it is surprising that (to our knowledge) no PRF candidates in $\mathsf{AC}^0[p]$ in the bounded-query setting with subexponential security have been proposed.

Limitations of the CIKK Learning Algorithm. The CIKK learning algorithm [CIKK16] shows how to turn the natural property by Razborov-Smolenksy [Raz87, Smo87] into a learning algorithm for $\mathsf{AC}^0[p]$ that requires quasipolynomially many membership queries. More generally, their approach transforms any natural property for a complexity class $\mathcal{C}$ containing $\mathsf{AC}^0[p]$, defined for n-input circuits of size $s(n)$, denoted $\mathcal{C}[s(n)]$, into a learning algorithm for any function $f \in \mathcal{C}[s(n)]$. The learner outputs a circuit C such that, with high probability, $C(x) = f(x)$ for a large fraction of inputs $x \in \{0,1\}^n$.

Oliveira and Santhanam [OS16] further showed that a $2^{O(n)}$-time natural property for $\mathcal{C}[\mathsf{poly}(n)]$ can be converted into a learning algorithm for $\mathcal{C}[\mathsf{poly}(n)]$ running in time $O(2^{n^\epsilon})$ for any $\epsilon > 0$. Furthermore, they showed that a $2^{O(n)}$-time natural property for $\mathcal{C}[2^{n^\epsilon}]$ yields a learning algorithm for $\mathcal{C}[\mathsf{poly}(n)]$ running in time $O(2^{\log(n)^{O(1)}})$. Analogous results hold to speed up the query complexity of learning as well.

This raises a natural question: can the CIKK framework be used to speed up the Razborov–Smolensky natural property and thereby obtain a learning algorithm with lower query complexity for $\mathsf{AC}^0[p]$?

We show the answer is negative. Specifically, in Appendix A, we prove a lower bound on the query complexity of CIKK learners in terms of the circuit size of the Nisan-Wigderson function.

Theorem 1.3 *Let $\mathcal{C}$ be a circuit class. Suppose $\mathcal{C}[s(n)]$ has a natural property making $g(m)$ queries, where m is the size of the target circuit. Then any CIKK learner for $\mathcal{C}$ requires at least $g(t)$ queries where t is the size of the Nisan-Wigderson generator circuit.*

This result implies that CIKK learners cannot achieve fewer queries than those required by the underlying natural property unless the circuit size of the Nisan-Wigderson generator can be reduced significantly. Consequently, no known CIKK-style attack can learn circuits in $\mathsf{AC}^0[p]$ below the quasipolynomial bound implies by the Razborov-Smolensky approximation, unless either a fundamentally new natural property is discovered, or the circuit size of the Nisan-Wigderson generator is reduced.

1.3 Related Work

Table 1 provides a comparison of our candidate with existing constructions, which we describe in more detail below.

PRF Candidates from LPN. Yu and Steinberger [YS16] introduce a PRF candidate based on the LPN assumption with noise rate $\mu = n^{-c}$ for any constant $0 < c < 1$. Their construction is almost in $\mathsf{AC}^0[2]$: it uses only parity and unbounded fan-in AND/OR gates, but has super-constant depth $\omega(1)$, where $\omega(1)$ can be any superconstant function.

They show that if there exists a natural property against $\mathsf{AC}^0[2]$ circuits of depth d making $O(n^{d/3})$ queries in time T with advantage $O(dn^{d/3}\epsilon)$, then there is an algorithm for solving the LPN problem with noise rate n^{-c} in αn queries for any constant $\alpha > 1$, time T and advantage ϵ. Note that one can set d to be constant to obtain a PRF in $\mathsf{AC}^0[2]$, but this would limit the adversary to a fixed polynomial number of queries. In contrast, our candidate lies fully within $\mathsf{AC}^0[2]$ (with depth just 3) and we conjecture that it remains secure against $p(\lambda)$ queries for *any* polynomial p.

In related work, Ding et al. construct a weak PRF in NC^1 based on strong LPN and propose multiple strong PRFs in NC^1 based on variants of LPN [DJK25]. Their constructions additionally satisfy key homomorphism, a property useful for many symmetric cryptographic applications. Since NC^1 strictly contains $\mathsf{AC}^0[2]$, these results highlight known separations between the two classes.

PRF Candidates from Factoring, DDH, and Ring LWE. Candidate PRFs in TC^0 have been constructed based on standard cryptographic assumptions such as Factoring, the Diffie–Hellman problem [NR97], and Ring-LWE [BPR12]. Through a generic reduction by Viola [Vio13] from subexponentially secure PRFs in TC^0 to quasipolynomially secure PRFs in $\mathsf{AC}^0[p]$, these constructions yield candidate

Circuit Class	Reference	Lower Bounds
AC^0 $\mathsf{AC}^0 + O(1)$ XOR,MAJ	[LMN89] [Vio13]	*No* poly-*query PRF*
$\mathsf{DNF} \circ \mathsf{MOD2}$	[ABG+14]	*No* poly-*query PRF with* subexp-*security*
$\mathsf{AC}^0 \circ \mathsf{MOD2}$	[ABG+14]	*No* poly-*query PRF with* subexp-*security (conjectured)*
$\mathsf{AC}^0[p]$	[RR94, KL01, CIKK16]	*No* subexp-*query PRF with* subexp-*security*

Circuit Class	Reference	Query	Runtime	Assumption
$\mathsf{AC}^0[2]$	[Kha93, NR97, BPR12, Vio13]	quasipoly	quasipoly	Factor, DDH, LWE
	This work	quasipoly	subexp	Heuristic
almost-$\mathsf{AC}^0[2]$	[YS16]	quasipoly	subexp	Low-noise LPN
$\mathsf{NC}^0[m]$	[BIP+18]	exp	exp	Heuristic
quasilinear-TC^0	[MV12]	exp	exp	Heuristic
TC^0	[NR97, NRR00, BPR12] [AR16]	subexp	subexp	Factor, DDH, LWE, Goldreich's OWF

Table 1: Comparison of negative (top table) and positive (bottom table) results for low-depth PRFs (where we only consider strong PRFs, i.e., PRFs secure against chosen-query attacks). We consider the complexity of computing the output for any fixed key, where security level in terms of query complexity (Query) and adversarial runtime (Runtime) is with respect to the key length. Here, poly denotes an a-priori unbounded polynomial number of queries. We write AC^0 to denote the class of constant-depth circuits with NOT and AND gates of polynomial size with unbounded fan-in. We write DNF to denote the subclass of depth-2 circuits. We write $\mathsf{DNF} \circ \mathsf{MOD2}$ to denote the complexity class DNF with mod 2 gates at the bottom (=input) layer (and equivalently for $\mathsf{AC}^0 \circ \mathsf{MOD2}$). We write $\mathsf{AC}^0[m]$ to denote the class AC^0 with mod m gates at all levels. We denote by TC^0 the class of constant-depth circuits with NOT, AND and MAJORITY (or, equivalently, threshold) gates of polynomial size with unbounded fan-in. Factor refers to the hardness of factoring, DDH to the decisional Diffie-Hellman assumption, LWE refers to Learning with Errors, and LPN refers to Learning Parity with Noise.

PRFs in $\mathsf{AC}^0[p]$ that are secure against quasipolynomial-time adversaries, assuming the subexponential hardness of the underlying assumptions. However, the reductions to these assumptions do not hold in our bounded query setting. Viola's construction takes the form $f_s(h_t(x))$, where f_s is the TC^0 candidate on $\mathsf{poly} \log n$-bit inputs and h_t is a hash function. In the bounded query setting, this candidate admits a distinguisher with inverse quasipolynomial advantage by searching the key space. Because we aim for security against subexponential adversaries with non-negligible advantage, this motivates the need for a new candidate construction.

Other Low-Complexity PRF Candidates. Two weak PRF candidates in $\mathsf{AC}^0[2]$ that achieve subexponential security are proposed in [BCG+21]. However, neither candidate is suitable as a strong PRF. The first candidate lies in AC^0 layered on top of parity gates and has *heavy Fourier coefficients*, making it efficiently learnable using the Kushilevitz–Mansour algorithm [KM91]. The second candidate is a *sparse polynomial* over $\mathbb{F}_2$, which can be efficiently learned with membership queries using algorithms from [DLM+07].

Separately, Boneh et al. [BIP+18] introduce a strong PRF candidate with exponential security in $\mathsf{NC}^0[6]$, and conjecture that similar constructions may exist for $\mathsf{NC}^0[m]$ where $m = pq$ for distinct primes p and q. Inspired by this candidate, Chen et al. [CJL25] formalize the Learning with Alternating Moduli (LAM) Assumption which they show to be as hard as LWE for certain large moduli. They propose candidate weak PRFs based on the LAM assumption in $\mathsf{NC}^0[m]$ where $m = pq$ for distinct primes p and q with $2^{O(n/\log n)}$ security. Note that $\mathsf{AC}^0[p]$ does not contain $\mathsf{NC}^0[pq]$ since MOD_q gates cannot be computed in $\mathsf{AC}^0[p]$ [Smo87].

Applebaum and Raykov [AR16] also propose a weak PRF candidate in AC^0 with quasipolynomial security and a strong PRF candidate in TC^0 with subexponential security based on Goldreich's one-way function [Gol00].

1.4 Open Questions

We highlight several directions that remain open.

- **General linear tests.** Proving resistance of our candidate against *all* linear tests remains open. We describe a proof attempt in Appendix C, and note its shortcomings. Even in the weak-PRF regime under non-adaptive queries, establishing resistance to linear tests for the second PRF candidate of [BCG+21] is still unresolved.

- **Specific linear and adaptive tests.** For the special case in which the adversary issues Q *linearly independent* queries (rather than queries sampled independently at random), we were unable to bound the distinguishing bias. Bounding the resulting bias may be of independent interest in Boolean Fourier analysis. Identifying additional special cases of linear tests and adaptive that exhibit new Fourier-analytic phenomena would also be valuable.
- **Sparsity beyond expectation.** We currently show only that the multilinear expansion of our candidate is non-sparse *in expectation*. Proving non-sparsity *with high probability* remains open. Deriving a tail bound is challenging because the monomial-coefficient indicators are dependent, so standard Chernoff-type arguments do not apply.
- **Minimality of the construction.** It remains open whether our structure is essentially minimal. For example, does there exist a (strong) PRF candidate in $\mathsf{MOD}_2 \circ \mathsf{AC}^0$? Clarifying this would improve our understanding of the exact complexity required for PRF constructions in low-depth classes.
- **Dichotomy between Learning and PRFs.** Can we, in general, build PRFs in $\mathsf{AC}^0[2]$ assuming only the non-learnability of the class in the quasipolynomial-query regime?
- **Circuit-size lower bound for Nisan–Wigderson.** Our query lower bound in Theorem 1.3 for CIKK depends on the circuit size of the Nisan–Wigderson (NW) function (Definition A.4). Establishing a circuit-size lower bound for the NW function would therefore yield a correspondingly stronger statement on CIKK's query complexity.

2 Preliminaries

Throughout this work, we denote by λ the security parameter. We say a function $f(\lambda)$ is negligible in the security parameter λ if $f(\lambda) = \lambda^{-\omega(1)}$. We denote this by writing $f(\lambda) = \mathsf{negl}(\lambda)$. We write $\mathsf{poly}(\lambda)$ to denote a function that is bounded by a fixed polynomial in λ. We say an algorithm is efficient if it runs in probabilistic polynomial time (PPT) in the length of its input. For a positive integer n, we use $[n]$ to denote the set $\{1, \ldots, n\}$. For a set S, $x \leftarrow S$ denotes uniformly sampling from S; additionally, $x_1, \ldots, x_L \leftarrow S$ for $L \in \mathbb{N}$ denotes that $x_1, \ldots, x_L$ are all independent and uniformly sampled from S. We use $\mathcal{U}_n$ to denote the uniform distribution over $\{0,1\}^n$. For a vector $x \in \{0,1\}^N$, we let $|x| := |\{i \in [N] : x_i = 1\}|$ denote the (Hamming) weight of the vector.

2.1 Fourier Analysis of Boolean Functions

In this section (and this work), we will regularly identify $\{0,1\}$ with $\mathbb{F}_2$, the finite Galois field with two elements.

Characters and Fourier basis. For $s \in \{0,1\}^N$, define the *character*

$$\chi_s(x) \;=\; (-1)^{\langle s,x\rangle} \qquad (x \in \{0,1\}^N) \;.$$

The set $\{\chi_s : s \in \mathbb{F}_2^N\}$ forms an orthonormal basis for real-valued functions on the Boolean hypercube $\{0,1\}^N$ under the inner product

$$\langle f, g\rangle \;=\; \mathbb{E}_x\big[f(x)\, g(x)\big], \qquad f, g : \{0,1\}^N \to \mathbb{R} \;.$$

Here and elsewhere, the notation $\mathbb{E}_x$ means that this expectation is over a uniformly random $x \in \{0,1\}^N$. For this inner-product, the characters form an orthonormal basis.

Lemma 2.1 (Orthonormality) *For all $s, t \in \{0,1\}^N$,*

$$\mathbb{E}_x\big[\chi_s(x)\chi_t(x)\big] = \mathbf{1}_{s=t}.$$

Fourier coefficients and expansion. For $f : \{0,1\}^N \to \mathbb{R}$, the *Fourier coefficient* of f at $s \in \{0,1\}^N$ is

$$\widehat{f}(s) \;=\; \mathbb{E}_x\big[f(x)\,\chi_s(x)\big] \;.$$

Due to Lemma 2.1, every f admits the Fourier expansion

$$f(x) \;=\; \sum_{s\in\{0,1\}^N} \widehat{f}(s)\,\chi_s(x) \;.$$

Lemma 2.1 implies this change-of-basis transformation is orthogonal, at least when the norms are suitably defined. That is, we have the following.

Theorem 2.2 (Parseval's Theorem) *For all $f : \{0,1\}^N \to \mathbb{R}$,*

$$\mathbb{E}_x[f(x)^2] \;=\; \sum_{s\in\{0,1\}^N} \widehat{f}(s)^2.$$

In particular, when the range of f is $\{-1,1\}$, then as clearly $\mathbb{E}_x\big[f(x)^2\big] = 1$, we have

$$\sum_{s\in\{0,1\}^N} \widehat{f}(s)^2 \;=\; 1 \;.$$

Lastly, we record what it means for a function to have a concentrated Fourier spectrum.

Definition 2.3 (Fourier ε-concentration) *Let $\mathcal{F} \subseteq \mathbb{F}_2^N$ be a set. We say that $f : \{0,1\}^N \to \{-1,1\}$ has Fourier mass ε-concentrated on $\mathcal{F}$ if*

$$\sum_{s\in\mathcal{F}} \widehat{f}(s)^2 \;\geq\; 1-\varepsilon \;.$$

See [O'D14] for further background.[7]

2.2 Pseudorandom Functions

We give the definitions of weak pseudorandom functions (WPRFs) and of standard pseudorandom functions (PRFs, or *strong* PRFs). We note that, as is standard, the adversary's query complexity, running time, and success probability are all measured with respect to the *key length* of the PRF.

Definition 2.4 ((Weak) Pseudorandom Function) *Let $\lambda \in \mathbb{N}$ be the security parameter, and let $N = N(\lambda)$ and $\kappa = \kappa(\lambda)$ denote polynomially bounded functions specifying the input length and key length, respectively.*

A (weak) pseudorandom function is a family of functions $\mathcal{F} = \{f_\lambda : \{0,1\}^\kappa \times \{0,1\}^N \to \{0,1\}\}$, where $f_\lambda(k,x)$ can be evaluated in polynomial time given (k,x). We write $f_k(x)$ instead of $f_\lambda(k,x)$. For $Q = Q(\kappa)$, $T = T(\kappa)$ and $\varepsilon = \varepsilon(\kappa)$, we say that $\mathcal{F}$ is a (Q,T,ε)-secure (strong) pseudorandom function (PRF) *if for every $\lambda \in \mathbb{N}$ and every oracle circuit $\mathcal{A}$ of size $T(\kappa)$,*

$$\left| \Pr_{k\leftarrow\{0,1\}^\kappa}[\mathcal{A}^{f_k(\cdot)} = 1] - \Pr_{R}[\mathcal{A}^{R(\cdot)} = 1] \right| \leq \varepsilon(\kappa),$$

where $R : \{0,1\}^N \to \{0,1\}$ denotes a uniformly random function. We call $\mathcal{F}$ a (Q,T,ε)-secure weak PRF (WPRF) *if the same condition holds when $\mathcal{A}$ is only given input-output pairs $(x_i, f_k(x_i))$, where each $x_i \leftarrow \{0,1\}^n$ is sampled uniformly and independently. For brevity, we also say $\mathcal{F}$ is a* (Q,T)-secure (strong/weak) PRF *if it is $(Q,T,1/T)$-secure (in either the strong or weak sense).*

Remark 2.5 *We remark that it is a common observation that one can trade running time for success probability. That is, to increase success probability, one can run the adversary multiple times with independent random tapes, and take a majority vote. However, this amplification also requires the adversary to use additional queries. Hence, if a PRF is (Q,T,ε)-secure with $Q = \mathsf{poly}(\kappa)$ but $\varepsilon = 2^{-\kappa^{\Omega(1)}}$, amplification to constant success would demand exponentially many queries, it is unclear how to obtain an attack with constant success probability while maintaining polynomial query complexity.*

[7] We remark that there, Fourier characters/coefficients use a domain of $\{-1,1\}$.

2.3 Algebraic Attacks and Rational Degree

Bogdanov et al. [BR17] used an algebraic attack to break Akavia et al.'s weak PRF candidate [ABG+14]. The attack requires the candidate to have a low *rational degree*. We define rational degree below and later in Section 4.3 prove our candidate has high rational degree.

Definition 2.6 ((Primal/Dual) Rational Degree) *Let $F : \mathbb{F}_2^N \to \mathbb{F}_2$ be a boolean function, which we consider as an element of the Boolean ring $\mathcal{R}_N := \mathbb{F}_2[x_1, \ldots, x_N]/I$, where $I = \langle x_1^2 - x_1, \ldots, x_N^2 - x_N \rangle$ denotes the Boolean ideal. The* primal rational degree *of F is*

$$\min\{d : \exists G \in \mathcal{R}_N \text{ with } \deg(G) = d \text{ s.t. } GF \equiv 0 \pmod{I}\} .$$

The dual rational degree *of F is*

$$\min\{d : \exists G \in \mathcal{R}_N \text{ with } \deg(G) = d \text{ s.t. } G(F+1) \equiv 0 \pmod{I}\} .$$

The rational degree $\mathsf{rdeg}(F)$ *is the minimum of its primal and dual rational degree.*

If a function has rational degree d (either primal or dual), one can look to interpolate a degree-d multilinear polynomial G witnessing $GF \equiv 0 \pmod{I}$ (in the primal case) or $G(F+1) \equiv 0 \pmod{I}$ (in the dual case). The interpolation is done by viewing the $\binom{N+1}{d}$ coefficients of G as unknowns, and solving for them via Gaussian elimination. To ensure that there are enough equations, one must query $\Theta\left(\binom{N+1}{d}\right)$ samples. Note that random samples are sufficient.

This attack demonstrates that, if F has rational degree d, then it is *not* a $\big(N^{O(d)}, N^{O(d)}, \Omega(1)\big)$-weak PRF. Note that in the weak PRF setting, in order for our candidate to satisfy subexponential security *if the adversary is also given subexponentially many samples*, it is necessary for us to argue that our candidates rational degree is $\geq \kappa^{\Omega(1)}$ (where we recall κ denotes key-length). This is what we do in Section 4.3 (in fact, we show the rational degree is $\Omega\big(\kappa^{1/6}\big)$).

2.4 Linear Attacks

Linear attacks have been implicitly considered in much of the cryptanalytic literature for many years (especially in code-based cryptography), but was only recently formalized (see, e.g., [CRR21]). The framework has become popular as it encompasses many practical attacks on code-based assumptions, especially those based on information set decoding (ISD) [Pra62, Ste89, FS09, BLP11, MMT11, BJMM12, MO15, EKM17, BM18].

To provide context, we describe how linear attacks apply to a dual-LPN style assumption which states that given a matrix $\mathbf{A}$ sampled according to some distribution, $(\mathbf{A}, \mathbf{Ae})$ is indistinguishable from $(\mathbf{A}, \mathbf{b})$ where $\mathbf{b}$ is a uniformly random vector, even if $\mathbf{e}$ is random sparse vector. The linear attack proceeds by identifying a nonzero *attack vector* $\mathbf{u}$ such that $\langle \mathbf{u}, \mathbf{y} \rangle$ is biased when $\mathbf{y} = \mathbf{Ae}$ for a randomly chosen sparse vector $\mathbf{e}$. Equivalently, $\mathbf{uA}$ must have low Hamming weight; i.e., $\mathbf{u}$ corresponds to a message whose codeword under $\mathbf{A}$ has low weight. Detecting such a bias in $\mathbf{y}$ is already enough to distinguish $(\mathbf{A}, \mathbf{Ae})$ from uniform.

This generalizes to an attack on a weak PRF in the following way. Given randomly sampled inputs $x^{(1)}, \ldots, x^{(Q)}$, before seeing the values $F(k, x^{(1)})$, $\ldots$, $F(k, x^{(Q)})$, the adversary must choose an attack vector $\mathbf{u} \in \{0,1\}^Q \setminus \{\mathbf{0}\}$, and the adversary is said to succeed if the bias

$$\left| \mathbb{E}\left[(-1)^{\sum_{i=1}^{Q} u_i F_k(x^{(i)})} \right] \right|$$

is non-negligible, where the randomness is over the key k, the values $x^{(1)}, \ldots, x^{(Q)}$, and any other internal randomness of the adversary.[8]

A contribution of our work is to define a generalized notion of linear tests that even *strong PRF's* must satisfy. Essentially, we allow the adversary to non-adaptively choose a set of query points $x^{(1)}, \ldots, x^{(Q)}$. In particular, they need not be independent and uniform. Again, before submitting the queries to the oracle, the adversary must choose the attack vector $\mathbf{u}$, and the attack is said to succeed if the adversaries determines a non-trivial bias in the inner product of the PRF values and the vector $\mathbf{u}$. The formal definition follows.

[8] Since we typically will not be explicitly giving a running time bound on the adversaries, we will be able to assume the adversary is deterministic. But in full generality, we allow the adversary to toss random coins, and this is useful when we generalize to the non-adaptive setting.

Definition 2.7 (Non-adaptive linear test for Strong PRF) *Let* $\{f_k : \{0,1\}^n \to \{0,1\} \mid k \leftarrow \{0,1\}^\kappa\}$ *be a family of keyed functions, and let* $Q \in \mathbb{N}$ *be a query bound. Let* $x^{(1)}, \ldots, x^{(Q)} \in \{0,1\}^n$ *be distinct inputs* chosen non-adaptively *by the adversary, and let* $\mathbf{u} = (u_1, \ldots, u_Q) \in \{0,1\}^Q \setminus \{\mathbf{0}\}$ *be a vector* chosen by the adversary *(also known as the* attack vector*) prior to making any queries. The adversary is said to* succeed *in a linear attack if*

$$\text{bias}(\mathbf{u}) := \left| \mathbb{E}\left[(-1)^{\sum_{i=1}^{Q} u_i f_k(x^{(i)})} \right] \right| \geq \varepsilon(n)$$

for some non-negligible function $\varepsilon(n)$*, where the randomness is over the choice of the key* k *and any internal randomness of the adversary.*

In this setting of non-adaptive linear tests, it in fact suffices to consider the attack vector $\mathbf{u} = (1, 1, \ldots, 1)$. This is because the adversary must decide the attack vector $\mathbf{u}$ based on her queries $x^{(1)}, \ldots, x^{(Q)}$, and if she wishes to set $u_i = 0$, she can simply not submit the query $x^{(i)}$. (This implicitly assumes the adversary need not fix ahead of time exactly how many queries it will make to the PRF, but we consider that reasonable.) In Section 4.2, all the non-adaptive linear attacks we consider have attack vector $\mathbf{u} = (1, 1, \ldots, 1)$.

Additionally, we will even consider a natural *adaptive* attack that can be viewed as a linear attack (see Lemma 4.10). However, defining adaptive linear attacks in full generality appears to be quite unreasonable (at least, any "natural" definition seems to be impossible to fulfill, as it would imply information-theoretic security of the PRF).

3 Strong PRFs in $\mathsf{AC}^0[2]$

Let λ denote the security parameter. We first define an auxiliary function $G : \{0,1\}^N \to \{0,1\}$ as the product of ℓ independent TRIBES functions on disjoint input blocks:

$$G(x) = \bigoplus_{i=1}^{\ell} \mathsf{TRIBES}_{m,n}\big(x^{(i)}\big), \tag{3}$$

where $x^{(i)}$ denotes the i^{th} M-bit block of x, and where each $\mathsf{TRIBES}_{m,n}$ function is given by

$$\mathsf{TRIBES}_{m,n}(x) = \bigvee_{j=1}^{m} \bigwedge_{k=1}^{n} x_{j,k}.$$

where $x_{j,k}$ denotes the k^{th} bit in the j^{th} sub-block of x. Here:

- $m = \Theta(\lambda / \log \lambda)$ and $n = \Theta(\log \lambda)$ are the *balanced tribes* parameters, chosen so that $mn = \lambda$ and $m = \Theta(2^n)$,
- $M = mn$ is the input length to each TRIBES function and $N = \ell mn$ is the input length to G, where $\ell = \lambda$.

We clarify that $n = \Theta(\log \lambda)$ and $m = \Theta(\lambda / \log \lambda)$ can be chosen so that $\mathsf{TRIBES}_{m,n}$ is *unbiased*, i.e., so that $\Pr_x[\mathsf{TRIBES}_{m,n}(x) = 1] \approx 1/2$. Specifically, following [O'D14, Proposition 4.12] (and the surrounding discussion), we can set parameters such that $\Pr_x[\mathsf{TRIBES}_{m,n}(x) = 1] = 1 - (1 - 2^n)^m = 1/2 - O\left(\frac{\log \lambda}{\lambda}\right)$. In the sequel, we assume m and n are chosen in this way.

We now define our candidate strong PRF. Let the PRF key $\mathbf{A} \in \{0,1\}^{N \times N}$ be a randomly sampled invertible matrix over $\mathbb{F}_2$ and $\mathbf{b} \in \{0,1\}^N$ be sampled uniformly at random. We remark that the key length is $\kappa(\lambda) = \lambda^4 + \lambda^2$.

Candidate Strong PRF. Our candidate strong PRF $F_{\mathbf{A},\mathbf{b}} : \{0,1\}^N \to \{0,1\}$ is defined as

$$F_{\mathbf{A},\mathbf{b}}(x) \;=\; G\big(\mathbf{A}x \oplus \mathbf{b}\big)\;. \tag{4}$$

Throughout the rest of the paper, we use F and G as defined above. We remark as well that $F_{\mathbf{I},\mathbf{0}} = G$, i.e., when the key is set to $(\mathbf{I}, \mathbf{0})$ the PRF is this function G.

Changing the range. At certain points, it will be useful for us to analyze the version of our function where the range is $\{-1,1\}$ rather than $\{0,1\}$. That is, we will sometimes consider the function $f_{\mathbf{A},\mathbf{b}} : \{0,1\}^N \to \{-1,1\}$ defined as

$$f_{\mathbf{A},\mathbf{b}}(x) := (-1)^{F_{\mathbf{A},\mathbf{b}}(x)} .$$

We similarly, we define

$$g(x) := (-1)^{G(x)}, \qquad \mathsf{TRIBES}^{\pm}_{m,n}(x) := (-1)^{\mathsf{TRIBES}_{m,n}(x)}.$$

In particular, we then have $g(x) = \prod_{i=1}^{\ell} \mathsf{TRIBES}^{\pm}_{m,n}(x^{(i)})$ and $f_{\mathbf{A},\mathbf{b}}(x) = g(\mathbf{A}x \oplus \mathbf{b})$, where as above $x^{(i)}$ denotes the i^{th} M-bit block of x.

3.1 Universal PRF from Strong PRFs in $\mathsf{AC}^0[2]$

We also construct a *universal PRF* in $\mathsf{AC}^0[2]$ under the assumption that a strong PRF exists in $\mathsf{AC}^0[2]$ in Theorem 3.1. The proof is in Appendix B.1.

Theorem 3.1 *If there exists a PRF $f_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ secure against $\mathsf{poly}(n)$ time and query adversaries in $\mathsf{AC}^0[2]$ for some $b \geq 1$, then there exists a fixed, universal function $UG : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ in $\mathsf{AC}^0[2]$, which is a PRF secure against $\mathsf{poly}(n)$ time and query adversaries.*

4 Rationales for Security

In this section, we demonstrate that our candidate withstands several classes of attacks. Specifically, we show that it is resilient to Fourier-analytic attacks, linear tests, and sparse polynomial testing. In addition, we establish that our candidate inherits the same guarantees as the construction of Akavia et al. [ABG+14], including inapproximability by low-degree polynomials over finite fields and the absence of correlation with any fixed family of functions. In each subsection, we prove resistance against attacks running in time $O(2^{\lambda})$. This implies that the security is $2^{O(\kappa^{1/4})}$, sub-exponential in the key length κ (recall $\kappa(\lambda) = \lambda^4$).

4.1 Absence of Large Fourier Characters

We analyze the Fourier spectrum of our candidate function $f_{\mathbf{A},\mathbf{b}}$ as defined in Equation (4). Note here that we are explicitly looking at the $\{\pm 1\}$-valued version of the Boolean function. We prove that all Fourier coefficients are exponentially small in the security parameter λ.

First, before continuing, we record the Fourier-characters for the version of the TRIBES function with range $\{\pm 1\}$.

Proposition 4.1 (Fourier Coefficients of TRIBES [O'D14], Proposition 4.14) *We have*

$$\widehat{\mathsf{TRIBES}^{\pm}}_{m,n}(0) = 2(1-2^{-n})^m - 1 .$$

Else, if $z = (z^{(1)}, \ldots, z^{(m)}) \neq 0$ where $z^{(i)}$ denotes the i^{th} n-bit block of $z \in \{0,1\}^M$, if $p = |\{i \in [m] : z^{(i)} \neq 0\}|$, then

$$\widehat{\mathsf{TRIBES}^{\pm}}_{m,n}(z) = 2(-1)^{p+|z|}2^{-pn}(1-2^{-n})^{m-p} .$$

Note that, for a Boolean function $f : \{0,1\}^N \to \{-1,1\}$,

$$\hat{f}(0) = \mathbb{E}_x[f(x)\chi_0(x)] = \Pr[f(x) = 1] - \Pr[f(x) = -1] ;$$

hence, showing that $|\hat{f}(0)|$ is small is enough to argue that f is *unbiased*. Thus, the following proposition is a strengthening of the claim that our PRF candidate is unbiased.

Proposition 4.2 (No Large Fourier Coefficients)

$$\max_{z \in \{0,1\}^N} \left\{ |\widehat{f_{\mathbf{A},\mathbf{b}}}(z)| \right\} \leq O(2^{-\lambda}).$$

Proof. Our proof relies on two lemmas.

Lemma 4.3 (Product Lemma for Disjoint Inputs) *Let ℓ, M, N be positive integers such that $N = \ell M$. Let $h_1, \ldots, h_\ell : \{0,1\}^M \to \{-1,1\}$ be Boolean functions, and define $f : \{0,1\}^N \to \{-1,1\}$ by*

$$f(x) = \prod_{i=1}^{\ell} h_i(x^{(i)}),$$

where $x = (x^{(1)}, \ldots, x^{(\ell)})$ is partitioned into ℓ disjoint blocks of size M. Then

$$\widehat{f}(z) = \prod_{i=1}^{\ell} \widehat{h}_i(z^{(i)}) \quad \textit{and} \quad \max_{z \in N} |\widehat{f}(z)| = \prod_{i=1}^{\ell} \max_{z^{(i)} \in \{-1,1\}^M} |\widehat{h}_i(z^{(i)})|.$$

Proof. For $z = (z^{(1)}, \ldots, z^{(\ell)})$,

$$\begin{aligned}
\widehat{f}(z) &= \frac{1}{2^N} \sum_{x \in \{0,1\}^N} \prod_{i=1}^{\ell} h_i(x^{(i)}) \cdot (-1)^{\langle x, z \rangle} \\
&= \prod_{i=1}^{\ell} \left[\frac{1}{2^M} \sum_{x^{(i)} \in \{0,1\}^M} h_i(x^{(i)}) \cdot (-1)^{\langle x^{(i)}, z^{(i)} \rangle} \right] \\
&= \prod_{i=1}^{\ell} \widehat{h}_i(z^{(i)}).
\end{aligned}$$

Taking maxima over z yields the second claimed equality. □

In our construction, the blocks of inputs to the ℓ tribes functions are mixed by $\mathbf{A}$ (and shifted by $\mathbf{b}$), rather than being disjoint. The following lemma shows that an invertible linear transformation preserves the maximum magnitude of a Fourier coefficient.

Lemma 4.4 *Let f and g be functions such that $f(x) = g(\mathbf{A}x + \mathbf{b})$ with $\mathbf{A} \in \{0,1\}^{N \times N}$ invertible over $\mathbb{F}_2$ and $\mathbf{b} \in \{0,1\}^N$. Then*

$$\max_{z \in \{0,1\}^N} |\widehat{f}(z)| = \max_{z \in \{0,1\}^n} |\widehat{g}(z)|.$$

Proof. Let $y = \mathbf{A}x + \mathbf{b}$, so $x = \mathbf{A}^{-1}(y - \mathbf{b})$. Then

$$\begin{aligned}
\widehat{f}(z) &= \frac{1}{2^N} \sum_{x \in \{0,1\}^N} g(\mathbf{A}x + \mathbf{b}) \cdot (-1)^{\langle x, z \rangle} \\
&= \frac{1}{2^N} \sum_{y \in \{0,1\}^N} g(y) \cdot (-1)^{\langle \mathbf{A}^{-1}(y - \mathbf{b}), z \rangle} \\
&= (-1)^{\langle z, -\mathbf{A}^{-1}\mathbf{b} \rangle} \cdot \frac{1}{2^N} \sum_{y \in \{0,1\}^N} g(y) \cdot (-1)^{\langle y, (\mathbf{A}^{-1})^\top z \rangle} \\
&= (-1)^{\langle z, -\mathbf{A}^{-1}\mathbf{b} \rangle} \cdot \widehat{g}\big((\mathbf{A}^{-1})^\top z\big),
\end{aligned}$$

which implies $|\widehat{f}(z)| = |\widehat{g}((\mathbf{A}^{-1})^\top z)|$ and hence the maxima are equal. □

Now, we apply the lemmas to our candidate. From Proposition 4.1, we see that for our choice of parameters,

$$\max_z |\widehat{\mathsf{TRIBES}}^{\pm}(z)| < O\left(\tfrac{\log \lambda}{\lambda}\right) .$$

Indeed, following the discussion in Section 3, we see that or choice of m and n guarantees that $(1 - 2^{-n})^m \leq \frac{1}{2} + O\left(\frac{\log \lambda}{\lambda}\right)$, and so

$$|\widehat{\mathsf{TRIBES}}^{\pm}(0)| = \left|2(1 - 2^{-n})^m - 1\right| \leq O\left(\tfrac{\log \lambda}{\lambda}\right) .$$

Otherwise, for $z = (z^{(1)}, \ldots, z^{(m)}) \neq 0$ (where as in Proposition 4.1, $z^{(i)}$ denotes the i^{th} n-bit block of $z \in \{0,1\}^M$), we have $p = |\{i \in [m] : z^{(i)} \neq 0\}| \geq 1$, implying

$$|\widehat{\mathsf{TRIBES}}^{\pm}(z)| = \left|2(-1)^{p+|z|} 2^{-pn} (1 - 2^{-n})^{m-p}\right| \leq 2 \cdot 2^{pn} = O\left(\tfrac{1}{\lambda^p}\right) .$$

Therefore, by Lemma 4.3,

$$\max_z |\widehat{G}(z)| \le \left(O\left(\tfrac{\log \lambda}{\lambda}\right)\right)^{\ell}.$$

Since $\ell = \lambda$, this is certainly at most $O(2^{-\lambda})$, as desired. Finally, $f_{\mathbf{A},\mathbf{b}}(x) = g(\mathbf{A}x \oplus \mathbf{b})$, so Lemma 4.4 implies that $\max_z |\widehat{f}(z)| = \max_z |\widehat{g}(z)|$, completing the proof. □

This shows that the largest Fourier coefficient of $f_{\mathbf{A},\mathbf{b}}$ is exponentially small in λ, and thus $f_{\mathbf{A},\mathbf{b}}$ has no non-negligibly large Fourier characters.

4.2 Resistance Against Linear Tests

We show that our candidate $F_{\mathbf{A},\mathbf{b}}$ as defined in Equation (4) resists a specific set of non-adaptive linear attacks as defined in Definition 2.7. Additionally, recalling the discussion that followed the definition, we will always assume the attack vector $\mathbf{u} = (1, 1, \ldots, 1)$.

Again in this section we find it easier to work with $\{-1, 1\}$-valued versions of the functions, namely, $f_{\mathbf{A},\mathbf{b}}$ and g.

A Simple Linear-Sum Attack and Its Fourier Analysis We study the following *linear-sum* distinguishing attack, inspired by the Blum–Luby–Rubinfeld (BLR) linearity test.

Linear Sum Attack. Fix an integer $Q \ge 2$. The distinguisher samples $x^{(1)}, \ldots, x^{(Q-1)} \leftarrow \mathbb{F}_2^N$ and sets

$$x^{(Q)} := \sum_{i=1}^{Q-1} x^{(i)}.$$

It then outputs the result of $\sum_{i=1}^{Q} F(x^{(i)})$. The probability the distinguisher outputs 1 for F is equivalent to the following:

$$\Delta_Q(f) := \mathbb{E}\Big[\prod_{i=1}^{Q} f\big(x^{(i)}\big)\Big] = \mathbb{E}\Big[(-1)^{\sum_{i=1}^{Q} F\left(x^{(i)}\right)}\Big],$$

where the expectation is over the random choice of the $x^{(i)}$'s.

We obtain a closed form Fourier sum for the bias of the linear sum attack in Lemma 4.5.

Lemma 4.5 (Bias as a Fourier sum) *For any $Q \ge 2$ and any $h : \mathbb{F}_2^N \to \{-1, 1\}$,*

$$\Delta_Q(h) = \sum_{s \in \mathbb{F}_2^N} \widehat{h}(s)^Q.$$

Proof. Expand each $h(x^{(i)})$ into its Fourier series:

$$\prod_{i=1}^{Q} h(x^{(i)}) = \sum_{s_1, \ldots, s_Q \in \mathbb{F}_2^N} \Big(\prod_{i=1}^{Q} \widehat{h}(s_i)\Big)(-1)^{\sum_{i=1}^{Q} \langle s_i, x^{(i)} \rangle}.$$

Taking expectation over the jointly distributed $x^{(1)}, \ldots, x^{(Q)}$ with the constraint $x^{(Q)} = \sum_{i=1}^{Q-1} x^{(i)}$ gives

$$\mathbb{E}\Big[(-1)^{\sum_{i=1}^{Q} \langle s_i, x^{(i)} \rangle}\Big] = \prod_{i=1}^{Q-1} \mathbb{E}_{x^{(i)}}\Big[(-1)^{\langle s_i + s_Q,\, x^{(i)} \rangle}\Big],$$

which equals 1 iff $s_i = s_Q$ for every $i \in [Q-1]$, and equals 0 otherwise. Thus only the diagonal terms $s_1 = \cdots = s_Q = s$ survive, yielding $\Delta_Q(h) = \sum_{s \in \mathbb{F}_2^N} \widehat{h}(s)^Q$. □

We can first use this observation to argue that functions with their Fourier mass concentrated on a small set (as in Definition 2.3) are distinguished from uniform via a linear attack.

Corollary 4.6 (Distinguisher for functions with concentrated Fourier mass) *If $h : \{0,1\}^N \to \{-1, 1\}$ has its Fourier coefficients ε-concentrated on a set $\mathcal{F}$, then the linear-sum attack distinguishes g from a random function with advantage $\frac{1-\varepsilon}{|\mathcal{F}|}$.*

Proof. If h is ε-concentrated on a set $\mathcal{F}$, then by Cauchy-Schwarz, setting $Q = 4$ in Lemma 4.5 yields

$$\Delta_4(h) \geq \sum_{s\in\mathcal{F}} \widehat{h}(s)^4 \;\geq\; \frac{1}{|\mathcal{F}|}\sum_{s\in\mathcal{F}} \widehat{h}(s)^2 \;\;\geq\;\; \frac{1-\varepsilon}{|\mathcal{F}|} \;\;.$$

□

Remark 4.7 *The Kushilevitz–Mansour algorithm [KM91] learns any Boolean function whose Fourier mass is ε-concentrated on a set $\mathcal{F}$ of size K in time* $\mathsf{poly}(K, n, 1/\varepsilon)$. *For distinguishing a PRF from random, however, full learning is unnecessary. Under the same spectral-concentration assumption, a much simpler* distinguisher *suffices. In particular, one can phrase the attack as a linear test that queries the function on a constant number of inputs and checks the corresponding correlation. This highlights the strength of linear tests. Consequently, establishing resistance against* all *linear tests would, in particular, rule out Fourier-based attacks that rely on detecting heavy coefficients. But we emphasize that attacks detecting heavy Fourier coefficients constitute just a small subset of all possible linear attacks.*

We now prove the upper bound that we use for our candidates.

Proposition 4.8 (Small Fourier spectrum ⇒ negligible bias) *Let $h : \mathbb{F}_2^N \to \{-1, 1\}$ and suppose $\max_{s\in\mathbb{F}_2^N} |\widehat{h}(s)| \leq \tau$. Then for every $Q \geq 2$,*

$$|\Delta_Q(h)| \;=\; \tau^{\,Q-2} \;\;.$$

In particular, if $\tau = 2^{-\Omega(\lambda)}$ and $Q \geq 3$, then $|\Delta_Q(h)| = 2^{-\Omega(\lambda)}$.

Proof. The statement follows from Lemma 4.5, the triangle inequality, and Parseval through the series of inequalities below.

$$|\Delta_Q(h)| \;=\; \Big|\sum_s \widehat{h}(s)^{\,Q}\Big| \;\leq\; \sum_s |\widehat{h}(s)|^Q \;\leq\; \tau^{\,Q-2}\sum_s \widehat{h}(s)^2 \;=\; \tau^{\,Q-2} \;\;.$$

□

Applying the upper bound to $g(x)$ which Proposition 4.2 establishes has exponentially small Fourier coefficients, we obtain Corollary 4.9.

Corollary 4.9 *For $x^{(1)}, \ldots, x^{(Q-1)} \leftarrow \mathbb{F}_2^N$ and $x^{(Q)} = \sum_{i=1}^{Q-1} x^{(i)}$,*

$$\Delta_Q(g) \;=\; 2^{-\Omega(\lambda)}.$$

Finally, we show that our candidate $F_{\mathbf{A},\mathbf{b}}$ resists the linear sum attack.

Lemma 4.10 *Let $x^{(1)}, \ldots, x^{(Q-1)} \leftarrow \mathbb{F}_2^N$ and $x^{(Q)} = \sum_{i=1}^{Q-1} x^{(i)}$. Define*

$$\delta \;:=\; \mathbb{E}\Big[(-1)^{\sum_{i=1}^{Q} F_{\mathbf{A},\mathbf{b}}(x^{(i)})}\Big] \;=\; \Delta_Q(f) \;\;.$$

Then $|\delta| \;=\; 2^{-\Omega(\lambda)}$.

Proof. Let $y^{(i)} := \mathbf{A}x^{(i)} + \mathbf{b}$. Because $\mathbf{A}$ is invertible and $x^{(i)}$ are uniform and independent, the vectors $y^{(1)}, \ldots, y^{(Q-1)}$ are uniform and independent. Moreover

$$\begin{aligned} y^{(Q)} = \mathbf{A}x^{(Q)} + \mathbf{b} \;&=\; \mathbf{A}\Big(\sum_{i=1}^{Q-1} x^{(i)}\Big) + \mathbf{b} \\ &= \sum_{i=1}^{Q-1} (\mathbf{A}x^{(i)} + \mathbf{b}) + (Q \bmod 2)\,\mathbf{b} \\ &= \sum_{i=1}^{Q-1} y^{(i)} \;+\; (Q \bmod 2)\,\mathbf{b} \;\;. \end{aligned}$$

Hence:

- If Q is even, then $y^{(Q)} = \sum_{i=1}^{Q-1} y^{(i)}$, and by Lemma 4.5 (applied to g evaluated at the $y^{(i)}$'s) we get $\delta = \sum_s \widehat{g}(s)^Q$.
- If Q is odd, then $y^{(Q)} = \sum_{i=1}^{Q-1} y^{(i)} + \mathbf{b}$. Conditioning on $\mathbf{b}$ and repeating the same computation yields $\delta = \sum_s \widehat{g}(s)^Q (-1)^{\langle s,\mathbf{b}\rangle}$. In particular, $|\delta|$ equals $\left|\sum_s \widehat{g}(s)^Q\right|$.

In both cases, $|\delta| \le \sum_s |\widehat{g}(s)|^Q \le \left(\max_s |\widehat{g}(s)|\right)^{Q-2}$ by Proposition 4.8. Using the bound $\max_s |\widehat{g}(s)| \le 2^{-\Omega(\lambda)}$ from Proposition 4.2 completes the proof. □

Lemma 4.5 shows that the attack's bias is exactly a spectral sum of $\widehat{g}^{Q-2}$ Thus, whenever the spectrum has a *heavy* coefficient or is ε-concentrated on a small set $\mathcal{F}$, the bias is non-negligible. In addition, if all coefficients are small (as we established for $f_{\mathbf{A},\mathbf{b}}$ in Proposition 4.2), then the bias is negligible by Proposition 4.8. This shows why linear attacks subsume the usual "heavy Fourier" attacks and why our candidate resists them.

Adaptive tests We show resistance against a simple adaptive test, where the adversary chooses random independent samples $x^{(i)}$ and discards samples such that $F(x^{(i)}) = 1$. Finally, the adversary tests whether $F\left(\sum_{i=1}^{Q-1} x^{(i)}\right) = 0$. We show that our candidate PRF resists this adaptive test. Due to the same argument as in Lemma 4.10, it suffices to prove resistance for the function G with randomly chosen independent inputs in $\mathbb{F}_2^N$.

Lemma 4.11 *If $x^{(1)}, x^{(2)}, \ldots, x^{(Q-1)}$ are randomly chosen independent vectors in $\mathbb{F}_2^N$ such that $G(x^{(i)}) = 0$ for $i \in [Q-1]$, $x^{(Q)} = \sum_{i=1}^{Q-1} x^{(i)}$, and $\delta = \Pr\left[G(x^{(Q)}) = 0 | G(x^{(i)}) = 0 \text{ for all } i \in [Q-1]\right] - \frac{1}{2}$, then $|\delta| = O(2^{-\lambda})$.*

Proof. Note that $g(x) = (-1)^{G(x)}$ also satisfies $g(x) = 1 - 2G(x)$, which equivalently gives $G(x) = \frac{1-g(x)}{2}$. We calculate the probability as follows:

$$\begin{aligned}&\Pr\left[G(x^{(Q)}) = 0 | G(x^{(i)}) = 0 \text{ for all } i \in [Q-1]\right] \\ &= \frac{\Pr\left[G(x^{(i)}) = 0 \text{ for all } i \in [Q]\right]}{\Pr\left[G(x^{(i)}) = 0 \text{ for all } i \in [Q-1]\right]}.\end{aligned}$$

Because $x^{(i)}$ are independent for $i \in [Q-1]$, the denominator is equal to $\left(\frac{1}{2} + \epsilon\right)^{Q-1}$ for negligible $\epsilon(\lambda)$. The numerator can be calculated as follows:

$$\begin{aligned}\mathbb{E}\left[\prod_{i=1}^{Q} 1 - G(x^{(i)})\right] = \mathbb{E}\left[\prod_{i=1}^{Q} \frac{1+g(x^{(i)})}{2}\right] &= \frac{1}{2^Q} + \frac{1}{2^Q}\mathbb{E}\left[\prod_{i=1}^{Q} g(x^{(i)})\right] \\ &= \frac{1}{2^Q} + \frac{1}{2^Q} \cdot O(2^{-\lambda}) \qquad \text{(Corollary 4.9)}\end{aligned}$$

Thus, the overall probability is upper bounded by

$$\frac{\frac{1}{2^Q} + \frac{1}{2^Q} \cdot O(2^{-\lambda})}{\left(\frac{1}{2} + \epsilon\right)^{Q-1}} \le \frac{\frac{1}{2^Q} + \frac{1}{2^Q} \cdot O(2^{-\lambda})}{\frac{1}{2^{Q-1}}} \le \frac{1}{2} + O(2^{-\lambda}) \ ,$$

and lower bounded by

$$\begin{aligned}\frac{\frac{1}{2^Q} + \frac{1}{2^Q} \cdot O(2^{-\lambda})}{\left(\frac{1}{2} + \epsilon\right)^{Q-1}} &\ge \frac{\frac{1}{2^Q}}{\frac{1}{2^{Q-1}}(1+2\epsilon)^{Q-1}} \ge \frac{1}{2} e^{-2\epsilon(Q-1)} \\ &\ge \frac{1}{2}\left(1 - 2\epsilon(Q-1)\right) \ge \frac{1}{2} - O(2^{-\lambda}) \ ,\end{aligned}$$

where we use the standard inequality $1 + x \le e^x$ for all $x \in \mathbb{R}$. This gives $|\delta| = O(2^{-\lambda})$. □

4.3 High Rational Degree

We show that our candidate has high rational degree as defined in Definition 2.6, thereby avoiding the algebraic attack that breaks the weak PRF candidate of [ABG+14]. For $N = lmn \in \mathbb{N}$, let

$$\mathcal{R}_N \;:=\; \mathbb{F}_2[X_1,\ldots,X_N]/\big(X_1^2+X_1,\ldots,X_N^2+X_N\big)$$

be the ring of multilinear polynomials over $\mathbb{F}_2$. Every $F \in \mathcal{R}_N$ has a unique multilinear representative

$$F \;=\; \sum_{S\subseteq[N]} c_S \prod_{i\in S} X_i \qquad (c_S \in \mathbb{F}_2),$$

and multiplication satisfies

$$\Big(\prod_{i\in S} X_i\Big)\cdot\Big(\prod_{i\in T} X_i\Big) \;=\; \prod_{i\in S\cup T} X_i \quad \text{in } \mathcal{R}_N$$

for all $S, T \subseteq [N]$. In particular, each $F \in \mathcal{R}_N$ is idempotent: $F^2 = F$.[9]

Let I denote the ideal $(X_{ijk}^2 + X_{ijk})_{(i,j,k)}$ in $\mathbb{F}_2[X_1,\ldots,X_N]$. We use the following elementary properties about the rational degree.

Lemma 4.12 *Let $F \in \mathcal{R}_N \setminus \{0\}$.*

1. (Algebraic-degree upper bound) *The (primal) rational degree of F is at most the algebraic degree of its multilinear representative in $\mathbb{F}_2[X_1,\ldots,X_N]$.*
2. (Affine invariance) *If $\mathbf{A} \in \mathbb{F}_2^{N\times N}$ is invertible and $\mathbf{b} \in \mathbb{F}_2^N$, then the rational degree of F equals that of $F(\mathbf{AX}+\mathbf{b})$.*

Proof. (1) Since $F^2 + F \equiv 0 \pmod I$, taking $G := F + 1$ shows that $GF \equiv 0 \pmod I$ and $\deg(G) = \deg(F)$, proving the claimed upper bound.

(2) Suppose $G(\mathbf{X})\,F(\mathbf{AX}+\mathbf{b}) \equiv 0$ in $\mathcal{R}_N$. Evaluating on $\mathbb{F}_2^N$ and substituting $\mathbf{Y} = \mathbf{AX}+\mathbf{b}$ shows $G'(\mathbf{Y})\,F(\mathbf{Y}) = 0$ for all $\mathbf{Y} \in \mathbb{F}_2^N$, where $G'(\mathbf{Y}) := G(\mathbf{A}^{-1}\mathbf{Y} + \mathbf{A}^{-1}\mathbf{b})$ has the same algebraic degree as G. Thus any annihilator of $F(\mathbf{AX}+\mathbf{b})$ of degree d induces one of F of the same degree (and vice versa), for both primal and dual definitions. Hence rational degree is preserved. □

Write the input variables as $\{X_{ijk}\}_{i\in[\ell],\, j\in[m],\, k\in[n]}$ so that block i consists of m tribes, each of width n. Let

$$H_i(X_i) := \prod_{j=1}^{m}\Big(1+\prod_{k=1}^{n} X_{ijk}\Big) \qquad \text{and} \qquad \tilde{F}(X) \;:=\; \sum_{i=1}^{\ell} H_i(X_i) \in \mathcal{R}_N.$$

Up to adding a constant (i.e., replacing F by $F+1$), this agrees with the standard multilinearization of the TRIBES construction; rational degree is invariant under $F \mapsto F+1$ (it swaps the primal and dual notions). By Lemma 4.12, it suffices to analyze the unshifted, identity case $F_{\mathbf{A},\mathbf{b}}(\mathbf{X}) = F(\mathbf{X})$ with $\mathbf{A} = I$ and $\mathbf{b} = \mathbf{0}$. We work with $\tilde{F}$ defined above since $\mathsf{rdeg}(F_{\mathbf{A},\mathbf{b}}) = \mathsf{rdeg}(\tilde{F})$. Each H_i has algebraic degree mn, hence $\deg(\tilde{F}) = mn$. By Lemma 4.12, we have the following corollary.

Corollary 4.13 (Trivial upper bound) *For the function $\tilde{F}$ with parameters ℓ, m, n,*

$$\mathsf{rdeg}\Big(\tilde{F}\Big) \;\le\; mn.$$

We now prove the main lower bound on the rational degree of our candidate.

Proposition 4.14 (Lower bound) *For the function $\tilde{F}$ with parameters ℓ, m, n,*

$$\mathsf{rdeg}\Big(\tilde{F}\Big) \;\ge\; \min\{\ell,\, mn\}.$$

Consequently, for any choice of invertible $\mathbf{A} \in \mathbb{F}_2^{N\times N}$ and $\mathbf{b} \in \mathbb{F}_2^N$, we also have $\mathsf{rdeg}(F_{\mathbf{A},\mathbf{b}}) \ge \min\{\ell, mn\}$.

[9] This follows from the fact that two multilinear polynomials agree as elements of $\mathcal{R}_N$ iff they agree pointwise on $\mathbb{F}_2^N$, and $F(x)^2 = F(x)$ for all $x \in \mathbb{F}_2^N$, as $y^2 = y$ for all $y \in \mathbb{F}_2$.

Proof. We show that no nonzero multilinear G of degree $< t := \min\{\ell, mn\}$ can satisfy either $G\,\tilde{F} \equiv 0 \pmod I$ (primal) or $G(\tilde{F}+1) \equiv 0 \pmod I$ (dual).

Write $G = \sum_{S\in\mathcal{G}} \mathbf{X}^S$, where each monomial $\mathbf{X}^S := \prod_{(i,j,k)\in S} X_{ijk}$ is square-free and $\mathcal{G} \subseteq 2^{[\ell]\times[m]\times[n]}$. Let $S^\star \in \mathcal{G}$ be of maximum size, so $|S^\star| = \deg(G) < t$. Because $|S^\star| < \ell$, there exists an index $i^\star \in [\ell]$ such that no variable from block $i^\star$ appears in $S^\star$, i.e., $(i^\star, j, k) \notin S^\star$ for all $j \in [m], k \in [n]$. Consider the monomial

$$M^\star \;:=\; \Big(\prod_{(i,j,k)\in S^\star} X_{ijk}\Big)\cdot\Big(\prod_{j=1}^{m}\prod_{k=1}^{n} X_{i^\star jk}\Big).$$

In the expansion of

$$G\cdot\tilde{F} = \sum_{i=1}^{\ell} G\cdot H_i = \sum_{i=1}^{\ell}\sum_{T\subseteq[m]}\sum_{S\in\mathcal{G}}\Big(\prod_{(i',j',k')\in S} X_{i'j'k'}\Big)\Big(\prod_{j\in T}\prod_{k=1}^{n} X_{ijk}\Big)\;,$$

the term with $i = i^\star$ and $T = [m]$ produces $M^\star$. Because $S^\star$ is disjoint from block $i^\star$, the variables in $M^\star$ are pairwise distinct, hence

$$\deg(M^\star) \;=\; |S^\star| + mn.$$

We claim that $M^\star$ appears *exactly once* (and hence, with coefficient 1) in $G\cdot\tilde{F}$ modulo I. Indeed, suppose another triple (i, T, S) contributes the same monomial $M^\star$. Then we must have

$$S \;\cup\; \big(\{i\}\times T\times[n]\big) \;=\; S^\star \;\cup\; \big(\{i^\star\}\times[m]\times[n]\big)\;.$$

If $i \neq i^\star$, the right-hand side would force $\{i^\star\}\times[m]\times[n] \subseteq S$, implying $|S| \geq mn$, which contradicts $|S| < t \leq mn$. Thus $i = i^\star$, and the disjointness shows $T = [m]$ and $S = S^\star$. Hence $M^\star$ survives with coefficient 1 in $G\tilde{F}$, so $G\tilde{F} \not\equiv 0 \pmod I$.

The above argument also establishes the dual rational degree. Indeed, we need to argue that $G\tilde{F} \not\equiv G \pmod I$. We have just shown that even after reducing modulo I, it must be that $\deg(G\tilde{F}) = \deg(G) + \deg(\tilde{F})$, which is strictly greater than $\deg(G)$ since $\deg(\tilde{F}) = mn \geq 1$. Hence, we cannot have $G\tilde{F} \equiv G \pmod I$. We conclude that no nonzero G of degree $< t$ annihilates $\tilde{F}$ or $\tilde{F}+1$, so $\mathsf{rdeg}(\tilde{F}) \geq t$. □

Corollary 4.15 *Let $N = \ell mn$. For the candidate $F_{\mathbf{A},\mathbf{b}}$ viewed as an element of $\mathcal{R}_N$,*

$$\min\{\ell, mn\} \;\leq\; \mathsf{rdeg}(F_{\mathbf{A},\mathbf{b}}) \;\leq\; mn.$$

This implies there is no rational degree attack on our candidate with less than $2^{\Omega(\lambda)}$ queries, which is impossible in our bounded query setting.

4.4 Non-sparsity over $\mathbb{F}_2$

We show that the expected number of nonzero monomials in the multilinear expansion of our candidate is subexponential. Throughout this section, let

$$U \;:=\; [\ell]\times[m]\times[n]\;.$$

We write our candidate as

$$F_{\mathbf{A},\mathbf{b}}(X) \;=\; \sum_{i=1}^{\ell}\left(1 \;+\; \prod_{j=1}^{m}\Big(1+\prod_{k=1}^{n}(\mathbf{A}\cdot X+\mathbf{b})_{(i,j,k)}\Big)\right). \tag{5}$$

As described in Section 4.3, the degree of $F_{\mathbf{A},\mathbf{b}}$ satisfies

$$\deg\big(F_{\mathbf{A},\mathbf{b}}\big) \leq M = mn. \tag{6}$$

We show that this polynomial is, in expectation, *not sparse*: over the choice of the key $(\mathbf{A}, \mathbf{b})$, the expected number of nonzero coefficients is very large. Our proof proceeds by (i) writing each coefficient as an explicit polynomial in the entries of $(\mathbf{A}, \mathbf{b})$, (ii) proving that this polynomial is not identically zero, and (iii) invoking the minimum distance of Reed–Muller codes to lower bound the fraction of $(\mathbf{A}, \mathbf{b})$ for which the coefficient is nonzero.

We repeatedly use the identity, valid over any ring, that

$$\prod_{t\in T}(\alpha_t+\beta_t) \;=\; \sum_{R\subseteq T}\left(\prod_{t\in R}\alpha_t\right)\left(\prod_{t\in T\setminus R}\beta_t\right). \tag{7}$$

Fix $i\in[\ell]$ and expand the ith block in (5):

$$\begin{aligned}
&\prod_{j=1}^{m}\Big(1+\prod_{k=1}^{n}(\mathbf{A}\cdot X+\mathbf{b})_{(i,j,k)}\Big)\\
&=\sum_{S\subseteq[m]}\prod_{j\in S}\prod_{k=1}^{n}(\mathbf{A}\cdot X+\mathbf{b})_{(i,j,k)}\\
&=\sum_{S\subseteq[m]}\prod_{(j,k)\in S\times[n]}\left(\sum_{v\in U}\mathbf{A}_{(i,j,k),v}\,X_v\;+\;\mathbf{b}_{(i,j,k)}\right)\\
&\overset{(7)}{=}\sum_{S\subseteq[m]}\sum_{R\subseteq S\times[n]}\left(\prod_{(j,k)\in R}\sum_{v\in U}\mathbf{A}_{(i,j,k),v}\,X_v\right)\left(\prod_{(j,k)\in(S\times[n])\setminus R}\mathbf{b}_{(i,j,k)}\right)\\
&=\sum_{S\subseteq[m]}\sum_{R\subseteq S\times[n]}\sum_{\varphi:R\to U}\left(\prod_{(j,k)\in R}\mathbf{A}_{(i,j,k),\varphi(j,k)}\right)\\
&\qquad\qquad\cdot\left(\prod_{(j,k)\in(S\times[n])\setminus R}\mathbf{b}_{(i,j,k)}\right)\prod_{(j,k)\in R}X_{\varphi(j,k)}\;.
\end{aligned}$$

Summing over i and adding the outer $+1$ yields a full expansion of $F_{\mathbf{A},\mathbf{b}}(X)$.

For a nonempty set $T\subseteq U$ with $|T|\le M$, let $X^T:=\prod_{u\in T}X_u$. Define $C_T(\mathbf{A},\mathbf{b})\in\mathbb{F}_2$ to be the coefficient of X^T in the expansion of $F_{\mathbf{A},\mathbf{b}}(X)$. From the expansion above,

$$\begin{aligned}
C_T(\mathbf{A},\mathbf{b}) \;=\; \sum_{i=1}^{\ell}1+\Bigg(&\sum_{S\subseteq[m]}\sum_{R\subseteq S\times[n]}\sum_{\substack{\varphi:R\to U\\ \varphi(R)=T}}\left(\prod_{(j,k)\in R}\mathbf{A}_{(i,j,k),\varphi(j,k)}\right)\\
&\cdot\left(\prod_{(j,k)\in(S\times[n])\setminus R}\mathbf{b}_{(i,j,k)}\right)\Bigg)\;.
\end{aligned}\tag{8}$$

In the following lemma, we show that each coefficient is not identically zero, when viewed as a polynomial in entries of $(\mathbf{A},\mathbf{b})$.

Lemma 4.16 *For every nonempty $T\subseteq U$ with $|T|\le M$, the polynomial $C_T(\mathbf{A},\mathbf{b})$ in the variables $\{\mathbf{A}_{u,v}\}_{u,v\in U}$ and $\{\mathbf{b}_u\}_{u\in U}$ is* not *the zero polynomial. Moreover, viewed as a polynomial in the $\mathbf{A}$–variables alone, C_T has a monomial of degree exactly $M=mn$.*

Proof. Fix any $i^\star\in[\ell]$ and consider in (8) the single summand with

$$S=[m],\qquad R=S\times[n]=[m]\times[n],$$

and ranging over all maps $\varphi:R\to U$ with image $\varphi(R)=T$. This contribution equals

$$\sum_{\substack{\varphi:[m]\times[n]\to U\\ \varphi([m]\times[n])=T}}\;\prod_{(j,k)\in[m]\times[n]}\mathbf{A}_{(i^\star,j,k),\,\varphi(j,k)}\;,$$

which is a sum of distinct degree mn monomials in the $\mathbf{A}$–variables. Distinctness holds because each factor is indexed by *ordered* pairs $\big((i^\star,j,k),\varphi(j,k)\big)$, so two different functions $\varphi\neq\varphi'$ yield different sets of variables and hence different monomials. Therefore this degree–mn portion cannot cancel internally, and no other term in (8) can cancel it either because terms with $R'\subsetneq[m]\times[n]$ have strictly smaller $\mathbf{A}$–degree, and terms with $i\neq i^\star$ involve a disjoint set of $\mathbf{A}$ variables in the first index. Hence C_T is a nonzero polynomial and has $\mathbf{A}$–degree exactly mn. □

Lemma 4.17 (Lower bound for coefficient bias) *Fix $T \subseteq U$ with $1 \le |T| \le M$. Then*

$$\Pr_{\mathbf{A} \leftarrow \mathbb{F}_2^{N \times N},\, \mathbf{b} \leftarrow \mathbb{F}_2^N} \big[\, C_T(\mathbf{A}, \mathbf{b}) = 1 \,\big] \;\ge\; 2^{-M}.$$

Proof. By Lemma 4.16, C_T is a nonzero polynomial of total degree at most M in the $N^2 + N$ variables $(\mathbf{A}, \mathbf{b})$. The evaluation vector of C_T over $\mathbb{F}_2^{N^2+N}$ is a nonzero codeword of the binary Reed–Muller code $\mathrm{RM}(M, N^2 + N)$, whose minimum distance is $2^{\,N^2+N-M}$ [MS77]. Thus, at least a 2^{-M} fraction of inputs $(\mathbf{A}, \mathbf{b}) \in \mathbb{F}_2^{N^2+N}$ make C_T equal to 1, proving the claim. □ Finally, applying Lemma 4.17 and Stirling's approximation gives us the main result of this section.

Proposition 4.18 (Sub-exponential expected number of monomials) *Let $S(\mathbf{A}, \mathbf{b})$ be the number of nonzero coefficients in the multilinear expansion of $F_{\mathbf{A},\mathbf{b}}(X)$ supported on monomials of degree at most M. Then*

$$\mathbb{E}_{\mathbf{A},\mathbf{b}}\big[S(\mathbf{A}, \mathbf{b})\big] = 2^{\Omega(\lambda \log \lambda)} \;.$$

Proof. Summing over all possible subsets T of size t and applying Lemma 4.17 gives us the following.

$$\begin{aligned}
\mathbb{E}_{\mathbf{A},\mathbf{b}}\big[S(\mathbf{A}, \mathbf{b})\big] \;&\ge\; \sum_{t=1}^{M} \binom{N}{t} \cdot 2^{-M} \;\ge\; \binom{N}{M} \cdot 2^{-M} \\
&\ge\; \frac{N^M}{M!} \cdot 2^{-M} \;\ge\; \frac{N^M}{(M/e)^M} \qquad \text{(Stirling)} \\
&= 2^{\Omega(\lambda \log \lambda)} \;.
\end{aligned}$$

□

4.5 Inapproximability by Low-Degree $\mathbb{F}_2$ Polynomials

We show that the correlation between a low degree $\mathbb{F}_2$ polynomial and our candidate is small. We measure correlation between a Boolean function and low-degree $\mathbb{F}_2$ polynomials as follows. For a function $f : \{0,1\}^M \to \{-1, 1\}$ and an integer $d \ge 0$, define

$$\mathsf{Corr}_d(f) \;:=\; \max_{\substack{p:\{0,1\}^M \to \{0,1\} \\ \deg(p) \le d}} \Big|\, \mathbb{E}_{x \leftarrow \{0,1\}^M} \big[f(x) \cdot (-1)^{p(x)} \big] \,\Big| \;.$$

We show that our candidate satisfies the same inapproximability guarantee as in [ABG+14]. The proof is in Appendix B.2.

Theorem 4.19 *Let $\alpha > 0$ be a constant and $K = K(\alpha)$. Let $M = mn$ be the input size for $\mathsf{TRIBES}_{m,n}$, and assume $\alpha n 2^n \le M \le (1-\alpha) n 2^n$. Then, for every $d \le \log M - \log\big(K \log M\big)$ there exists $c = c(\alpha) > 0$ such that*

$$\mathsf{Corr}_d\big(F_{\mathbf{A},\mathbf{b}}\big) \;\le\; 1 - c \cdot 2^{-d}.$$

In particular, taking parameters $\ell = \Theta(\lambda), m = \Theta(\lambda / \log(\lambda))$ and $n = \Theta(\log \lambda)$ satisfies the regime above, and the same bound holds uniformly for all invertible $\mathbf{A}$ and all $\mathbf{b}$.

4.6 Lack of Correlation with Fixed Function Families

Similarly to Section 4.5, we define the *correlation* of two functions (not necessarily polynomials) $g, h : \mathbb{F}_2^N \to \{-1, 1\}$ as

$$\mathsf{Corr}(g, h) := \mathop{\mathbb{E}}_{x \in \mathbb{F}_2^N} [g(x) h(x)] \;,$$

where x is sampled uniformly at random from $\mathbb{F}_2^N$.

Considering any family $\mathcal{H}$ of functions $h : \mathbb{F}_2^N \to \{-1, 1\}$, we show that, with high probability over the choice of secret key $(\mathbf{A}, \mathbf{b})$, $f_{\mathbf{A},\mathbf{b}}$ will not noticeably correlate with any function $h \in \mathcal{H}$, assuming $\mathcal{H}$ is not too large. We prove the following theorem, which is analogous to Corollary 17 from [ABG+14], showing that our candidate has low correlation with fixed function families. The proof is in Appendix B.3.

Theorem 4.20 *Let $\eta, \gamma > 0$ be such that $1-\gamma-2\eta > 0$. Let $\mathcal{H}$ be any set of functions $h : \mathbb{F}_2^N \to \{-1, 1\}$ of size $|\mathcal{H}| \leq 2^{\gamma N}$. Let $\mu_g := \mathbb{E}_x[f_{\mathbf{I},\mathbf{0}}(x)]$ and $\mu_h := \mathbb{E}_x[h(x)]$. Then with probability at least $1 - 2^{-(1-\gamma-2\eta)N}$ over the choice of key $(\mathbf{A}, \mathbf{b})$, for all $h \in \mathcal{H}$, we have*

$$-2^{-\eta N} \leq \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) - \mu_g \mu_h \leq 2^{-\eta N} \ .$$

In particular, for all $h \in \mathcal{H}$,

$$-|\mu_g| - 2^{-\eta N} \leq \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) \leq |\mu_g| + 2^{-\eta N}$$

Remark 4.21 *Note that for our choice of parameters, we have $|\mu_g| \leq 2^{-\ell} = O\left(2^{-\lambda}\right)$, so the total bound on the correlation is negligible in the security parameter λ, and thus also in the key-length $\kappa(\lambda) = O(\lambda^4)$.*

5 Acknowledgements

We thank Peter Scholl for helpful discussions and for pointing out that the improved efficiency achieved via the trapdoored LPN construction extends to the distributed setting. The work of L. Kohl is funded by NWO Talent Programme Veni (VI.Veni.222.348) and by NWO Gravitation project QSC. Part of this work was conducted while the author was visiting the Simons Institute for the Theory of Computing. N. Resch is supported by an NWO (Dutch Research Council) grant with number C.2324.0590. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2039655 as well as the National Science Foundation CAREER Award #2443735. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors(s) and do not necessarily reflect the views of the National Science Foundation.

References

[ABG+14] Adi Akavia, Andrej Bogdanov, Siyao Guo, Akshay Kamath, and Alon Rosen. Candidate weak pseudorandom functions in AC^0 *o* MOD_2. In Moni Naor, editor, *ITCS 2014*, pages 251–260. ACM, January 2014.

[AKK+03] Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing low-degree polynomials over GF(2(. In Sanjeev Arora, Klaus Jansen, José D. P. Rolim, and Amit Sahai, editors, *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2003 and 7th International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM 2003, Princeton, NJ, USA, August 24-26, 2003, Proceedings*, volume 2764 of *Lecture Notes in Computer Science*, pages 188–199. Springer, 2003.

[AR16] Benny Applebaum and Pavel Raykov. Fast pseudorandom functions based on expander graphs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 27–56. Springer, Berlin, Heidelberg, October / November 2016.

[ARS+15] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 430–454. Springer, Berlin, Heidelberg, April 2015.

[BCG+20] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Correlated pseudorandom functions from variable-density LPN. In *61st FOCS*, pages 1069–1080. IEEE Computer Society Press, November 2020.

[BCG+21] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Low-complexity weak pseudorandom functions in $\mathtt{AC0}[\mathtt{MOD2}]$. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 487–516, Virtual Event, August 2021. Springer, Cham.

[BHI+20] Marshall Ball, Justin Holmgren, Yuval Ishai, Tianren Liu, and Tal Malkin. On the complexity of decomposable randomized encodings, or: How friendly can a garbling-friendly PRF be? In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 86:1–86:22. LIPIcs, January 2020.

[BIP+18] Dan Boneh, Yuval Ishai, Alain Passelègue, Amit Sahai, and David J. Wu. Exploring crypto dark matter: New simple PRF candidates and their applications. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 699–729. Springer, Cham, November 2018.

[BJMM12] Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^{n/20}$: How $1 + 1 = 0$ improves information set decoding. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 520–536. Springer, Berlin, Heidelberg, April 2012.

[BKS+10] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *51st FOCS*, pages 488–497. IEEE Computer Society Press, October 2010.

[BLP11] Daniel J. Bernstein, Tanja Lange, and Christiane Peters. Smaller decoding exponents: Ball-collision decoding. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 743–760. Springer, Berlin, Heidelberg, August 2011.

[BM18] Leif Both and Alexander May. Decoding linear codes with high error rate and its impact for LPN security. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, pages 25–46. Springer, Cham, 2018.

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 719–737. Springer, Berlin, Heidelberg, April 2012.

[BR17] Andrej Bogdanov and Alon Rosen. Pseudorandom functions: Three decades later. Cryptology ePrint Archive, Report 2017/652, 2017.

[CCF+16] Anne Canteaut, Sergiu Carpov, Caroline Fontaine, Tancrède Lepoint, María Naya-Plasencia, Pascal Paillier, and Renaud Sirdey. Stream ciphers: A practical solution for efficient homomorphic-ciphertext compression. In Thomas Peyrin, editor, *Fast Software Encryption - 23rd International Conference, FSE 2016, Bochum, Germany, March 20-23, 2016, Revised Selected Papers*, volume 9783 of *Lecture Notes in Computer Science*, pages 313–333. Springer, 2016.

[CIKK16] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Learning algorithms from natural proofs. In Ran Raz, editor, *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, volume 50 of *LIPIcs*, pages 10:1–10:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[CJL25] Yilei Chen, Liheng Ji, and Wenjie Li. Learning with alternating moduli, arora-ge over composite moduli, and weak prfs. *IACR Cryptol. ePrint Arch.*, page 968, 2025.

[CRR21] Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 502–534, Virtual Event, August 2021. Springer, Cham.

[DJK25] Youlong Ding, Aayush Jain, and Ilan Komargodski. A new approach for LPN-based pseudorandom functions: Low-depth and key-homomorphic. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 1898–1909. ACM Press, June 2025.

[DLM+07] Ilias Diakonikolas, Homin K. Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A. Servedio, and Andrew Wan. Testing for concise representations. In *48th FOCS*, pages 549–558. IEEE Computer Society Press, October 2007.

[EKM17] Andre Esser, Robert Kübler, and Alexander May. LPN decoded. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 486–514. Springer, Cham, August 2017.

[FS09] Matthieu Finiasz and Nicolas Sendrier. Security bounds for the design of code-based cryptosystems. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 88–105. Springer, Berlin, Heidelberg, December 2009.

[Gol00] Oded Goldreich. Candidate one-way functions based on expander graphs. Cryptology ePrint Archive, Report 2000/063, 2000.

[JPRZ09] Charanjit S. Jutla, Anindya C. Patthak, Atri Rudra, and David Zuckerman. Testing low-degree polynomials over prime fields. *Random Struct. Algorithms*, 35(2):163–193, 2009.

[Kha93] Michael Kharitonov. Cryptographic hardness of distribution-specific learning. In *25th ACM STOC*, pages 372–381. ACM Press, May 1993.

[KL01] Matthias Krause and Stefan Lucks. On the minimal hardware complexity of pseudorandom function generators. In Afonso Ferreira and Horst Reichel, editors, *STACS 2001, 18th Annual Symposium on Theoretical Aspects of Computer Science, Dresden, Germany, February 15-17, 2001, Proceedings*, volume 2010 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2001.

[KM91] Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the Fourier sprectrum (extended abstract). In *23rd ACM STOC*, pages 455–464. ACM Press, May 1991.

[LMN89] Nathan Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, Fourier transform, and learnability. In *30th FOCS*, pages 574–579. IEEE Computer Society Press, October / November 1989.

[MJSC16] Pierrick Méaux, Anthony Journault, François-Xavier Standaert, and Claude Carlet. Towards stream ciphers for efficient FHE with low-noise ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 311–343. Springer, Berlin, Heidelberg, May 2016.

[MMT11] Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $\tilde{\mathcal{O}}(2^{0.054n})$. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 107–124. Springer, Berlin, Heidelberg, December 2011.

[MO15] Alexander May and Ilya Ozerov. On computing nearest neighbors with applications to decoding of binary linear codes. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 203–228. Springer, Berlin, Heidelberg, April 2015.

[MS77] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The theory of error-correcting codes*, volume 16. Elsevier, 1977.

[MV12] Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 68–85. Springer, Berlin, Heidelberg, August 2012.

[NR97] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *38th FOCS*, pages 458–467. IEEE Computer Society Press, October 1997.

[NRR00] Moni Naor, Omer Reingold, and Alon Rosen. Pseudo-random functions and factoring (extended abstract). In *32nd ACM STOC*, pages 11–20. ACM Press, May 2000.

[NW88] Noam Nisan and Avi Wigderson. Hardness vs. randomness (extended abstract). In *29th FOCS*, pages 2–11. IEEE Computer Society Press, October 1988.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of computer and System Sciences*, 49(2):149–167, 1994.

[O'D14] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[OS16] Igor C. Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds and pseudorandomness. *CoRR*, abs/1611.01190, 2016.

[Pra62] Eugene Prange. The use of information sets in decoding cylcic codes. *IRE Transactpons on Information Theory*, 8(5):5–9, 1962.

[Raz87] Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mat. Zametki*, 41(4):598–607, 1987.

[RR94] Alexander A. Razborov and Steven Rudich. Natural proofs. In *26th ACM STOC*, pages 204–213. ACM Press, May 1994.

[RST15] Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for boolean circuits. *CoRR*, abs/1504.03398, 2015.

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In Alfred Aho, editor, *19th ACM STOC*, pages 77–82. ACM Press, May 1987.

[Ste89] Jacques Stern. A method for finding codewords of small weight. *Coding Theory and Applications*, 388(10):106–113, 1989.

[Vio13] Emanuele Viola. The communication complexity of addition. In Sanjeev Khanna, editor, *24th SODA*, pages 632–651. ACM-SIAM, January 2013.

[VZ25] Vinod Vaikuntanathan and Or Zamir. Improving algorithmic efficiency using cryptography. *CoRR*, abs/2502.13065, 2025.

[Wil13] Ryan Williams. Natural proofs versus derandomization. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 21–30. ACM Press, June 2013.

[YS16] Yu Yu and John P. Steinberger. Pseudorandom functions in almost constant depth from low-noise LPN. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 154–183. Springer, Berlin, Heidelberg, May 2016.

A Limitations of Natural Property-Based Learners

In this section, we provide a more detailed discussion on the rationales for the hardness of learning in $\mathsf{AC}^0[p]$ (Conjecture 1.2). The CIKK learning algorithm [CIKK16] for $\mathsf{AC}^0[p]$ uses a natural property, which we define below.

A.1 Natural Properties

Let F_n be the collection of all Boolean functions on n variables. Λ and Γ denote complexity classes. A combinatorial property is a sequence of subsets of F_n for each n.

Definition A.1 (Natural Property [Raz87]) *A combinatorial property $R_n \subseteq F_n$ is Γ-natural against Λ with density δ_n if it satisfies the following three conditions:*

1. ***Constructivity:*** *The predicate $f_n \overset{?}{\in} R_n$ is computable in Γ. We call D a* natural property distinguisher *if it computes the predicate $f_n \overset{?}{\in} R_n$.*
2. ***Largeness:*** *$|R_n| \geq \delta_n \cdot |F_n|$.*
3. ***Usefulness:*** *For any sequence of functions f_n, if $f_n \in \Lambda$, then $f_n \notin R_n$.*

Remark A.2 *Note that this implies the natural property distinguisher D can distinguish any Boolean function $f_n \in \Lambda$ from a randomly chosen function in F_n with advantage δ_n.*

A.2 Limitations of Natural Properties Against $\mathsf{AC}^0[p]$

In order to invalidate the security of a PRF candidate, it suffices to distinguish it from random, and a natural property precisely promises such a distinguisher. For $\mathsf{AC}^0[2]$, Razborov [Raz87] and Smolensky [Smo87] argue that $\mathsf{AC}^0[2]$ can be approximated by polynomials of polylogarithmic-degree. Krause and Lucks [KL01] point out that so long as one can *tolerantly test* low-degree polynomials – namely, given oracle access to a function which is either promised to be close to a low-degree polynomial, or is far from any low-degree polynomial – then one has an effective distinguisher for any PRF candidate computable in $\mathsf{AC}^0[2]$. Indeed, by Razborov-Smolensky [Raz87, Smo87] if one has oracle access to the PRF candidate then it is close to a low-degree polynomial, while if the function is random then it is far from any low-degree polynomial (with very high probability).

It is reasonable to consider the sample complexity of such an attack on our PRF. We note that the depth D of our candidate is 3 (at least; it is even more if we view the key as an explicit input), and the

size S is at least $N(\lambda) = \sqrt{\kappa(\lambda)}$, where we recall again that λ is the security parameter, N the input size, and κ the key-size. Hence, even if one is happy with a very large approximation error ε the degree of the Razborov-Smolensky approximating polynomial is at least $\log(S/\varepsilon)^D \geq \log^3(N) = \log^3(\sqrt{\kappa}) = \log^3(\kappa)/8$. We thus have a polylogarithmic degree d for the Razborov-Smolensky approximating polynomial.

Now, how many samples would one need to take to determine if the oracle is implementing a function which is close to a degree-d polynomial? As mentioned above, this is a *tolerant testing* problem, but we can consider the formally easier task where, in the yes case, the oracle is promised to be *exactly implementing* a degree-d polynomial. This problem has seen much attention [AKK+03, JPRZ09] and in particular it is known that $\Omega(2^{d-1})$ queries are *necessary*, at least if one wishes to distinguish with constant advantage. This follows from the fact that the family of degree-d multivariate polynomials forms of $\Omega(2^{d-1})$-wise independent hash function family; thus, with less queries it is information-theoretically impossible to distinguish the case that the function is implementing a uniformly random degree-d polynomial from the case it is implementing a uniformly random function. This implies that if one hopes for a distinguisher with constant advantage, then $2^{\log^3(\kappa)/8}$ samples – that is, a quasi-polynomially large number – must be taken.

Lastly, we would like to point out that this also means that this recipe cannot yield a distinguisher with, say, $\mathsf{poly}(\kappa)$ query complexity. Indeed, by the reasoning above seeing only $\mathsf{poly}(\kappa)$ values one cannot learn anything in general about a degree $d = \Omega(\log^3 \kappa)$ polynomial.

Thus, to speed up this attack one would have find a better approximating polynomial for our PRF (or, more ambitiously, improve on Razborov-Smolensky, which hasn't been done in the last 35 years). In any case, in light of the current state-of-affairs, we do not see any reason to suspect that distinguishing attacks based on natural properties defined in terms of approximating polynomials could invalidate our candidate's conjectured security (Conjecture 1.1).

A.3 Limitations of Learning in $\mathsf{AC}^0[p]$

We recall the CIKK [CIKK16] learner for classes $\mathcal{C}$ that contain $\mathsf{AC}^0[p]$ and admit a natural property. Besides the natural property, the learner uses a family of low-intersection sets (a Nisan-Wigderson design).

A.4 Nisan–Wigderson designs and generator

Definition A.3 (NW design [NW94]) *For $n, m, L \in \mathbb{N}$, a sequence of sets $S_1, \ldots, S_L \subseteq [m]$ forms an NW design if*

- *$|S_i| = n$ for all $1 \leq i \leq L$ and*
- *$|S_i \cap S_j| \leq \log L$ for all $1 \leq i \neq j \leq L$.*

Definition A.4 (NW generator [NW94]) *Let $f : \{0,1\}^n \to \{0,1\}$, set $m = n^2$, and let $L(m) < 2^n$ be a stretch function. For an NW design $S_1, \ldots, S_L \subseteq [m]$, define the NW pseudorandom generator $G_f : \{0,1\}^m \to \{0,1\}^{L(m)}$ as:*

$$G_f(z) = f(z|_{S_1}) f(z|_{S_2}) \ldots f(z|_{S_{L(m)}}),$$

where $z|_S$ is the restriction of z to coordinates in S. We call g_z the NW function, where $g_z(i) = f(z|_{S_i})$ is the i^{th} truth table entry of G_f.

Theorem A.5 (NW function circuit size [CIKK16]) *There exists an NW design such that the function $MX_{NW} : \{0,1\}^\ell \times \{0,1\}^m \to \{0,1\}^n$, defined by*

$$MX_{NW}(i, z) = z|_{S_i} \ ,$$

is computable by an $\mathsf{AC}^0[p]$ circuit of size $O(n^3 \log(n) \log(L))$.

On the other hand, providing a lower bound on the size of the NW function g_z remains an open question.

A.5 The CIKK Learner

We now formalize a CIKK-style learner. Our definition is deliberately liberal so that the CIKK query lower bound applies to a broad class of learning algorithms. In particular, we place no restrictions on how the hypothesis h is obtained; the only requirement is that the procedure eventually invokes a natural property distinguisher on an NW function.

Definition A.6 (CIKK Learner) *Let $\mathcal{C}$ be a circuit class, and suppose there exists a* natural property distinguisher D *for $\mathcal{C}$. A* CIKK learner for $\mathcal{C}$ with respect to D *for a function f is a probabilistic algorithm $\mathcal{A}^D$ such that*

- *$\mathcal{A}^D$ runs D on the Nisan-Wigderson function $g_z(i)$ (whose truth table is the string $G_f(z)$) of size L for some stretch function $L(m) > m$.*
- *$\mathcal{A}^D$ outputs a hypothesis $h : \{0,1\}^n \to \{0,1\}$ such that*

$$\Pr_{x \sim U_n}[h(x) = f(x)] \;\geq\; \tfrac{1}{2} + \tfrac{1}{\mathsf{poly}(n)}.$$

Algorithm 1 CIKK Learning Algorithm

Input: Oracle to function $f : \{0,1\}^n \to \{0,1\}$, natural property distinguisher D
Output: Circuit C approximating f

1. Randomly sample $i \leftarrow \{1, \ldots, L\}$.
2. For each $1 \leq j \leq n^2$ with $j \notin S_i$, randomly sample $z_j \leftarrow \{0,1\}$.
3. For each $k < i$, query f on all $2^{|S_i \cap S_k|} \leq L$ strings that might be $z|_{S_k}$ consistent with $z'_j s$. Store these in table T.
4. Randomly sample $b_k \leftarrow \{0,1\}$ for all $i \leq k \leq L$.

Output the following circuit C: on input x, set $z|_{S_i} = x$, while other bits of z are fixed to random.

- Set $b_k = f(z|_{S_k})$ for all $k < i$ using the table T.
- If $D(b_1, \ldots, b_L) = 1$, output b_i. Else, output $1 - b_i$.

The expensive step is recovering the true values $f(z|_{S_k})$ for indices $k < i$. Because each S_k overlaps previously recovered sets in at most $\log L$ positions, the number of oracle calls needed to compute b_k is at most $2^{\log L} = L$ for each k. Summing over k gives $Q(n) = \Theta(L^2)$ queries in the standard implementation.

A.6 A query lower bound for CIKK

Oliveira and Santhanam [OS16] (Lemma 6) show two speedup implications for CIKK:

1. A $2^{O(n)}$-time natural property for $\mathcal{C}[\mathsf{poly}(n)]$ gives a learning algorithm for $\mathcal{C}[\mathsf{poly}(n)]$ running in time $O(2^{n^\varepsilon})$ for any $\varepsilon > 0$.
2. A $2^{O(n)}$-time natural property for $\mathcal{C}[2^{n^\varepsilon}]$ yields a learning algorithm for $\mathcal{C}[\mathsf{poly}(n)]$ running in time $2^{(\log n)^{O(1)}}$.

Analogous statements hold when one plugs CIKK into the standard reductions to improve the *query* complexity of learning.

However, these speedups disappear once circuit size is accounted for. The statements above parameterize running time solely by the distinguisher/learner input length n, suppressing the circuit-size parameter. Rewriting them in terms of the circuit size m makes this explicit:

1. If $m = n^c$ for some $c > 0$, then a $2^{O(m^{1/c})}$-time natural property for $\mathcal{C}[m]$ yields a learner for $\mathcal{C}[m]$ running in time $O\big(2^{m^{\varepsilon/c}}\big)$ for any $\varepsilon > 0$. In other words, we start and end with subexponential-time procedures giving no asymptotic improvement.
2. If $m = 2^{n^\varepsilon}$ (with $\varepsilon, c > 0$), then a $2^{O\left((\log m)^{1/\varepsilon}\right)}$-time natural property for $\mathcal{C}[m]$ gives a learner for $\mathcal{C}[s = n^c]$ running in time $2^{(\log s)^{O(1)}}$. Here both sides are quasipolynomial.

The apparent advantage comes from viewing the property as a function of n alone. Once the necessary size threshold is enforced, the overall learning complexity does not actually improve.

Moreover, for classes such as $AC^0[p]$ any meaningful natural property must depend on circuit size since every Boolean function has a (potentially exponential-size) AC^0 implementation (e.g., as a 2-DNF).

Lemma A.7 (CIKK learners are limited by the NW circuit size for size-dependent natural properties) *Let $\mathcal{C}$ be a typical circuit class. Suppose $\mathcal{C}$ admits a CIKK learner A^D as in Definition A.6, which uses a natural property distinguisher D making $g(s)$ queries, where s is the size of the target circuit. Let t be the circuit size of the NW function g_z. Then A^D requires at least $g(t)$ queries.*

Proof. By the definition of a CIKK learner (Definition A.6), A^D must run D on the NW function g_z. Thus, D must run on circuits of size at least t, requiring at least $g(t)$ queries. From the definition of a CIKK learner (Definition A.6) which runs D on G_f, since D requires at least $\Omega(t)$ queries, then A^D also requires at least $g(t)$ queries. □

Applying this lemma to the best known natural property (Razborov-Smolensky) for $\mathsf{AC}^0[p]$, which uses $\Omega(2^{\log^d(m)})$ queries, where m is the circuit size and d is the depth of the circuit and the smallest known NW circuit size $t = \Omega(n^3 \log n \log L)$ from Theorem A.5 gives us the following corollary.

Corollary A.8 *Suppose a CIKK learner uses the Razborov-Smolensky natural property for $\mathsf{AC}^0[p]$ and an NW circuit size of $\Omega(n)$. Then the learner requires at least quasipolynomially many queries.*

Remark A.9 *Note that the above result* does not *extend to the regime where one measures the query complexity of the distinguisher D solely as a function of the input length n. The subtlety is that the NW function g_z is parameterized by an input length of $\Theta(\log L)$. Hence its* input length *can be as small as $O(\log n)$ even when the relevant* circuit size *parameters might remain as large as $\Omega(n)$.*

To obtain a polynomial-query learner, we would require a natural property that uses polynomially many queries in the circuit size or a smaller circuit size for the NW function, which is currently not known for $\mathsf{AC}^0[p]$.

B Proofs from Section 3 and Section 4

We give the proofs from Section 3 and Section 4 here.

B.1 Proof of Theorem 3.1

Before we begin, we say that a PRF $\{f_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}\}_{n\in\mathbb{N}}$ in $\mathsf{AC}^0[2]$ is constructable in time $t(n)$ if there is a Turing Machine running in time $t(n)$ that outputs an $\mathsf{AC}^0[2]$ circuit computing f_n in time $t(n)$.

We prove Theorem 3.1 by first establishing Lemma B.3: if a PRF exists in $\mathsf{AC}^0[2]$, then there exists one in $\mathsf{AC}^0[2]$ computable in $O(n^2)$ time. In turn, Lemma B.3 follows from the PRF domain shrinkage (Lemma B.1) and domain extension (Lemma B.2) lemmas below.

Lemma B.1 (PRF Domain Shrinkage) *If $f_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ is a PRF, secure against $\mathsf{poly}(n)$ time and query adversaries, that can be represented as a circuit in $AC^0[2]$ constructable in time and size $O(n^c)$ for some $b, c \geq 1$ and depth $d = O(1)$, then there exists a PRF $f'_n := f_{n^\epsilon} : \{0,1\}^{n^{b\epsilon}} \times \{0,1\}^{n^\epsilon} \to \{0,1\}$, for some $\epsilon \in (0,1)$, secure against arbitrary $\mathsf{poly}(n)$ time and query adversaries, which can be represented as an $AC^0[2]$ circuit constructable in time and size $O(n)$ and depth d.*

Proof. Let $\epsilon = 1/c$. If f_{n^ϵ} admitted an $O(n^{c_2})$ time or query attack, then f_n would admit a $O(n^{c_2/\epsilon})$ time or query attack, a contradiction. So f_{n^ϵ} is also secure against arbitrary $\mathsf{poly}(n)$ time/query adversaries. Furthermore, f'_n is constructable in time and size $O((n^\epsilon)^c) = O(n)$. The depth of f'_n remains the same as f_n. □

Lemma B.2 (PRF Domain Extension) *If $f'_n : \{0,1\}^{n^{b\epsilon}} \times \{0,1\}^{n^\epsilon} \to \{0,1\}$ is a PRF, secure against $\mathsf{poly}(n)$ time and query adversaries, that can be represented as a circuit in $AC^0[2]$ constructable in time and size $O(n)$ for some $b \geq 1$ and depth d for some $\epsilon \in (0,1)$, then there exists a PRF $g_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$, secure against arbitrary $\mathsf{poly}(n)$ time and query adversaries, which can be represented as an $AC^0[2]$ circuit computable in time $O(n^2)$ and depth $2d+1$.*

Proof. We define $g_n(k, x)$ as follows:

1. Let k_1 be the first $n^{b\epsilon}$ bits of k and k_2 be the second $n^{b\epsilon}$ bits of k.
2. Let $l = \frac{n^\epsilon}{2}$ and $m = \frac{n}{l}$. Partition the function input x into m blocks of size l: x_1, x_2, …, x_m.

3. For each $i \in [m]$, we define $y_i \in \{0,1\}^{n^\epsilon}$ such that its j^{th} bit is $f'_n(k_1, x_i || n^\epsilon \cdot (i-1) + j)$ for $j \in [n^\epsilon]$, where $||$ denotes the concatenation. Note that $n^\epsilon \cdot (i-1) + j$ can be represented by $\log(3n) < \frac{n^\epsilon}{2}$ bits for large enough n.

4. Let $z = \bigoplus_{i=1}^{m} y_i$.

5. Output $f'_n(k_2, z)$.

First, we verify that g_n is constructable in time and size $O(n^2)$. The function g_n contains one XOR gate and $m \cdot n^\epsilon + 1 = 2n + 1$ functions f'_n, which are each constructable in time and size $O(n)$, which means g_n is constructable in time and size $O(n^2)$. The depth is $2d + 1$.

Let $\mathcal{A}$ be an adversary against the PRF security of g_n. Let $q^{(1)}, q^{(2)}, \ldots, q^{(t(n))}$ are its distinct queries for some polynomial $t(n)$. Let $z^{(1)}, z^{(2)}, \ldots, z^{(t(n))}$ be the outputs of step 4 for the queries, before they are plugged into f'_n. We wish to show that these $z^{(i)}$ are all distinct with high probability. If they are distinct, since f'_n is a PRF, then the distribution $\{f'_n\left(k_2, z^{(i)}\right)\}_{i \in [t(n)], k_2 \in \{0,1\}^{n^b \epsilon}}$ is indistinguishable from $\mathcal{U}_{t(n)}$ (with some negligible advantage δ) which means g_n is a PRF by a hybrid argument.

We upper bound $\Pr\left[z^{(a)} = z^{(b)}\right]$ for distinct $a, b \in [t(n)]$. We know that $q^{(a)}$ and $q^{(b)}$ must differ in at least one block i^*, so the distribution $y_{i^*}^{(a)} || y_{i^*}^{(b)}$ (the outputs from step 3) is indistinguishable (with some negligible advantage δ') from $\mathcal{U}_{2n^\epsilon}$ since f'_n is a PRF. Thus,

$$\Pr\left[z^{(a)} = z^{(b)}\right] < \left(\frac{1}{2}\right)^{n^\epsilon} + \delta'$$

By the Union Bound, we have

$$\Pr\left[z^{(i)} \text{ are not distinct}\right] < t(n)^2 \left(\frac{1}{2}\right)^{n^\epsilon} + t(n)^2 \delta'.$$

Finally, we obtain the advantage of the PRF game for g_n is at most

$$\delta + t(n)^2 \left(\frac{1}{2}\right)^{n^\epsilon} + t(n)^2 \delta',$$

which is negligible if δ and δ' are negligible and $t(n)$ is polynomial. □

Lemma B.3 *If $f_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ is a PRF, secure against $\mathsf{poly}(n)$ time and query adversaries, that can be represented as a circuit in $AC^0[2]$ constructable in time and size $O(n^c)$ for some $b, c \geq 1$ and depth d, then there exists a PRF $g_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$, secure against $\mathsf{poly}(n)$ time and query adversaries, that can be represented as a circuit in $AC^0[2]$ constructable in time and size $O(n^2)$ and depth $2d + 1$.*

Proof. Let $\epsilon = \frac{1}{c}$ and apply Lemma B.1 to obtain f'_n. The statement follows via Lemma B.2. □

Theorem B.4 *If there exists a PRF $f_n : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ secure against $\mathsf{poly}(n)$ time and query adversaries in $\mathsf{AC}^0[2]$ for some $b \geq 1$, then there exists a fixed, universal function $UG : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ in $\mathsf{AC}^0[2]$, which is a PRF secure against $\mathsf{poly}(n)$ time and query adversaries.*

Proof. Take any enumeration of all Turing Machines M_i. We enumerate all possible function families that take $\{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$ in $\mathsf{AC}^0[2]$ as follows

1. For each $i \in \mathbb{N}$, run the Turing Machine M_i for n^2 steps.

2. Interpret the resulting output string from M_i as an $\mathsf{AC}^0[2]$ circuit (as all circuits can be represented by binary strings) and call this circuit F_i.

If a PRF exists in $\mathsf{AC}^0[2]$, Lemma B.3, shows that a there must be a PRF computable in time $O(n^2)$. In addition, since PRFs must have a constant-size description with respect to n, then for some large enough n, the truncated list

$$F_1, F_2, F_3, \ldots, F_{\sqrt{n}}$$

must contain a PRF. Note that $\sqrt{n}$ is chosen arbitrarily, we could also choose any $\mathsf{poly}(n)$.

$$UF(k_1, \ldots, k_{\sqrt{n}}, x) = \bigoplus_{i=1}^{\sqrt{n}} F_i(k_i, x).$$

We show that UF is a PRF secure against $\mathsf{poly}(n)$ time and query adversaries. Assume UF is not secure against $\mathsf{poly}(n)$ time or query adversaries for the sake of contradiction. Then there is a polynomial-time or query distinguisher $\mathcal{D}$ such that

$$\left|\Pr_{\{k_1,\ldots,k_{\sqrt{n}}\}\leftarrow\{0,1\}^{n\times\sqrt{n}}}\left[\mathcal{D}^{UF_k}=1\right] - \Pr_{R\leftarrow RF_n}\left[\mathcal{D}^{R}=1\right]\right| > \frac{1}{n^c}$$

where RF_n is the set of all functions that map $\{0,1\}^n \to \{0,1\}$ and for some constant c. Since we know there exists a PRF in the truncated list, let j be the first index $1 \le j \le \sqrt{n}$ such that F_j is a PRF. We construct a distinguisher $\mathcal{A}$ for F_j as follows:

- Sample $k_i \leftarrow \{0,1\}^n$ randomly for all $1 \le i \ne j \le \sqrt{n}$.
- Let $\mathcal{O}$ be the oracle representing F_j or function sampled randomly from RF_n.
- We construct the oracle $\mathcal{U}$ from $\mathcal{O}$ as follows:

$$\mathcal{U}(x) = \bigoplus_{i=1}^{j-1} F_i(k_i, x) \bigoplus \mathcal{O}(x) \bigoplus_{i=j+1}^{\sqrt{n}} F_i(k_i, x).$$

- Run the distinguisher $\mathcal{D}$ with $\mathcal{U}$ as the oracle and output the outcome.

If $\mathcal{O}$ represents F_j, then $\mathcal{U}$ has the same distribution as UF_k where $\{k_1, \ldots, k_{\sqrt{n}}\} \leftarrow \{0,1\}^{n\times\sqrt{n}}$. Otherwise, if $\mathcal{O}$ is a function sampled randomly from RF^n, then $\mathcal{U}$ also has the same distribution as a function sampled randomly from RF^n. Thus,

$$\left|\Pr_{k_j\leftarrow\{0,1\}^n}\left[\mathcal{A}^{F_j}=1\right] - \Pr_{R\leftarrow RF_n}\left[\mathcal{A}^{R}=1\right]\right| > \frac{1}{n^c}.$$

If $\mathcal{D}$ uses $\mathsf{poly}(n)$ time or queries, then $\mathcal{A}$ also uses $\mathsf{poly}(n)$ time or queries. This is a contradiction that F_j is a PRF secure against $\mathsf{poly}(n)$ time and query adversaries, so we have shown UF is a PRF secure against $\mathsf{poly}(n)$ time and query adversaries. The size of UF is at most $\sqrt{n}\cdot\mathsf{poly}(n)$ and the depth is at most $d+1$, where d is the maximum depth of the enumerated functions F_i. This ensures UF stays in the class $\mathsf{AC}^0[2]$.

Finally, to construct a PRF UG with key and domain sizes which are the same as the original PRF f_n, we proceed in two steps. First, by applying the domain shrinkage technique from Lemma B.1, we obtain the PRF $UF' : \{0,1\}^{n^{3b\epsilon/2}} \times \{0,1\}^{n^\epsilon} \to \{0,1\}$ for some $\epsilon \in (0,1)$. Then, using the domain extension technique from Lemma B.2, we obtain the PRF $UG : \{0,1\}^{n^b} \times \{0,1\}^n \to \{0,1\}$. □

B.2 Proof of Theorem 4.19

We recall the correlation result from [ABG+14] for the TRIBES function, and prove Lemma B.6 and Lemma B.7 to show that the same correlation result holds for our candidate $F_{\mathbf{A},\mathbf{b}}$.

Theorem B.5 ([ABG+14], Thm. 3) *Let $\alpha > 0$ be a constant and let $K = K(\alpha)$ be a corresponding constant. Suppose the TRIBES parameters m, n satisfy $M = mn$ and $\alpha n 2^n \le M \le (1-\alpha)n2^n$. Then for every degree $d \le \log M - \log\big(K \log M\big)$ there exists a constant $c = c(\alpha) > 0$ such that*

$$\mathsf{Corr}_d\big(\mathsf{TRIBES}_{m,n}\big) \;\le\; 1 - c \cdot 2^{-d}.$$

The bound above is robust under natural operations that we use in our construction.

Lemma B.6 *Let $\mathbf{A} \in \mathbb{F}_2^{M\times M}$ be invertible and $\mathbf{b} \in \mathbb{F}_2^M$. For any $f : \{0,1\}^M \to \{-1,1\}$ and any d,*

$$\mathsf{Corr}_d\big(x \mapsto f(\mathbf{A}x+\mathbf{b})\big) \;=\; \mathsf{Corr}_d(f).$$

Proof. The map $x \mapsto \mathbf{A}x+\mathbf{b}$ is a bijection on $\{0,1\}^M$, hence preserves the uniform distribution. For any polynomial p with $\deg(p) \le d$, the composition $x \mapsto p(\mathbf{A}x+\mathbf{b})$ also has degree $\le d$. Therefore

$$\max_{\deg(p)\le d}\big|\mathbb{E}_x f(\mathbf{A}x+\mathbf{b})(-1)^{p(x)}\big| \;=\; \max_{\deg(q)\le d}\big|\mathbb{E}_y f(y)(-1)^{q(y)}\big|,$$

after the change of variables $y = \mathbf{A}x+\mathbf{b}$. This equals $\mathsf{Corr}_d(f)$. □

Lemma B.7 *Let $f_1,\dots,f_\ell$ be functions on disjoint input blocks: $f_i : \{0,1\}^M \to \{-1,1\}$, and write $f'(x^{(1)},\dots,x^{(\ell)}) := \prod_{i=1}^{\ell} f_i(x^{(i)})$, which corresponds to the XOR of the $\{0,1\}$-valued versions on disjoint inputs. Then for every d,*

$$\mathsf{Corr}_d(f') \;\le\; \max_{i\in[\ell]} \mathsf{Corr}_d(f_i).$$

Proof. Fix any polynomial p of degree at most d over all variables, and set $g := \prod_{i=1}^{\ell} f_i(x^{(i)}) \cdot (-1)^{p(x^{(1)},\dots,x^{(\ell)})}$. As a function of $x^{(t)}$, the factor $(-1)^{p(\cdot)}$ is a polynomial of degree at most d (times a $\{\pm1\}$ constant coming from the fixed blocks). Hence

$$\Big|\mathbb{E}_{x^{(t)}}\big[f_t(x^{(t)})\cdot(-1)^{\tilde p(x^{(t)})} \,\big|\, \{x^{(i)}\}_{i\neq t}\big]\Big| \;\le\; \mathsf{Corr}_d(f_t),$$

for some polynomial $\tilde p$ of degree at most d. Taking total expectation and then the maximum over p,

$$\mathsf{Corr}_d(f') \;=\; \max_{\deg(p)\le d}\big|\mathbb{E}_x g\big| \;\le\; \max_{t\in[\ell]} \mathsf{Corr}_d(f_t).$$

□

Combining Lemmas B.6 and B.7 with Theorem B.5 yields the desired inapproximability guarantee for our candidate, since $F_{\mathbf{A},\mathbf{b}}$ is the XOR (i.e., product in $\{-1,1\}$) of ℓ disjoint copies of $\mathsf{TRIBES}_{m,n}$ applied to the affine image $\mathbf{A}x+\mathbf{b}$, where $\mathbf{A} \in \mathbb{F}_2^{N\times N}$ is invertible and $\mathbf{b} \in \mathbb{F}_2^N$.

Corollary B.8 *Let $\alpha > 0$ be a constant and $K = K(\alpha)$. Let $M = mn$ be the input size for $\mathsf{TRIBES}_{m,n}$, and assume $\alpha n 2^n \le M \le (1-\alpha)n2^n$. Then, for every $d \le \log M - \log\big(K\log M\big)$ there exists $c = c(\alpha) > 0$ such that*

$$\mathsf{Corr}_d\big(F_{\mathbf{A},\mathbf{b}}\big) \;\le\; 1 - c\cdot 2^{-d}.$$

In particular, taking parameters $\ell = O(\lambda), m = O(\lambda/\log(\lambda)$ and $n = O(\log\lambda)$ satisfies the regime above, and the same bound holds uniformly for all invertible $\mathbf{A}$ and all $\mathbf{b}$.

B.3 Proof of Theorem 4.20

Recall for $\mathbf{A} \leftarrow \mathbb{F}_2^{N\times N}$ and $\mathbf{b} \leftarrow \mathbb{F}_2^N$, we denote $f_{\mathbf{A},\mathbf{b}} = (-1)^{F_{\mathbf{A},\mathbf{b}}}$ denote the ±1 version of our PRF candidate. We also denote $g := (-1)^G$, where we recall $G = F_{\mathbf{I},\mathbf{0}}$. Define $\mu_g := \mathbb{E}_x[g(x)]$ and $\mu_h := \mathbb{E}_x[h(x)]$. Considering any family $\mathcal{H}$ of functions $h : \mathbb{F}_2^N \to \{-1,1\}$, we show that, with high probability over the choice of secret key $(\mathbf{A},\mathbf{b})$, $f_{\mathbf{A},\mathbf{b}}$ will not noticeably correlate with any function $h \in \mathcal{H}$, assuming $\mathcal{H}$ is not too large. This will follow from the following lemma, combined with a union bound.

Lemma B.9 *Fix any $h : \mathbb{F}_2^N \to \{-1,1\}$. Then, for all $\varepsilon > 0$,*

$$\Pr\left[\left|\mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) - \mu_g\cdot\mu_h\right| \ge \varepsilon\right] \le 2^{-N}/\varepsilon^2 \ .$$

Proof. Consider the random variable (over the randomness of the key $(\mathbf{A},\mathbf{b})$)

$$Z(\mathbf{A},\mathbf{b}) := \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) = \mathop{\mathbb{E}}_x[f_{\mathbf{A},\mathbf{b}}(x)h(x)] \ .$$

Firstly, observe that, if y is also distributed uniformly over $\mathbb{F}_2^N$,

$$\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})] = \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[\mathop{\mathbb{E}}_x[f_{\mathbf{A},\mathbf{b}}(x)h(x)]\right]$$

$$
\begin{aligned}
&= \mathop{\mathbb{E}}_{x}\left[\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[f_{\mathbf{A},\mathbf{b}}(x)h(x)]\right] \\
&= \mathop{\mathbb{E}}_{x}\left[h(x)\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[g(\mathbf{A}x+\mathbf{b})]\right] \\
&= \mathop{\mathbb{E}}_{x}\left[h(x)\mathop{\mathbb{E}}_{y}[g(y)]\right] \\
&= \mu_h \cdot \mu_g \ ,
\end{aligned}
$$

where we used that $\mathbf{A}x+\mathbf{b}$ is distributed uniformly and independently of x. Hence, the bound in the statement is equivalent to the following tail bound for $Z(\mathbf{A},\mathbf{b})$:

$$
\Pr\left[\left|Z(\mathbf{A},\mathbf{b}) - \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})]\right| \geq \varepsilon\right] \leq 2^{-N}/\varepsilon^2 \ .
$$

Chebyshev's inequality says that

$$
\Pr\left[\left|Z(\mathbf{A},\mathbf{b}) - \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})]\right| \geq \varepsilon\right] \leq \mathop{\mathrm{Var}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})]/\varepsilon^2 \ .
$$

Hence, the desired bound follows from upper bounding

$$
\mathop{\mathrm{Var}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})] = \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[Z(\mathbf{A},\mathbf{b})^2\right] - \left(\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})]\right)^2 \leq 2^{-N}.
$$

Towards establishing this, we upper bound the second moment as follows (below, x, x' are independent and uniform over $\mathbb{F}_2^N$):

$$
\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[Z(\mathbf{A},\mathbf{b})^2\right] &= \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[\mathop{\mathbb{E}}_{x}[g(\mathbf{A}x+\mathbf{b})\cdot h(x)]\cdot\mathop{\mathbb{E}}_{x'}[g(\mathbf{A}x'+\mathbf{b})\cdot h(x')]\right] \\
&= \mathop{\mathbb{E}}_{x,x'}\left[h(x)\cdot h(x')\cdot\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[g(\mathbf{A}x+\mathbf{b})\cdot g(\mathbf{A}x'+\mathbf{b})]\right] \\
&= \frac{1}{2^{2N}}\sum_{x_0} h(x_0)^2\cdot\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[g(\mathbf{A}x_0+\mathbf{b})^2\right] \\
&\quad + \frac{1}{2^{2N}}\sum_{x_0\neq x_0'} h(x_0)h(x_0')\cdot\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[g(\mathbf{A}x_0+\mathbf{b})\cdot g(\mathbf{A}x_0'+\mathbf{b})] \ .
\end{aligned}
$$

Now, if $x_0 \neq x_0'$ are fixed vectors, then $\mathbf{A}x_0+\mathbf{b}$ and $\mathbf{A}x_0'+\mathbf{b}$ are independent and uniform over $\mathbb{F}_2^N$. That is, if y and y' are independent and uniform over $\mathbb{F}_2^N$,

$$
\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[g(\mathbf{A}x_0+\mathbf{b})\cdot g(\mathbf{A}x_0'+\mathbf{b})] = \mathop{\mathbb{E}}_{y}[g(y)]\cdot\mathop{\mathbb{E}}_{y'}[g(y')] = \mu_g^2 \ .
$$

Hence, since $h(x_0)^2 = g(\mathbf{A}x_0+\mathbf{b})^2 = 1$ for all $x_0 \in \mathbb{F}_2^N$, it follows that

$$
\begin{aligned}
&\frac{1}{2^{2N}}\sum_{x_0} h(x_0)^2\cdot\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[g(\mathbf{A}x_0+\mathbf{b})^2\right] \\
+&\frac{1}{2^{2N}}\sum_{x_0\neq x_0'} h(x_0)h(x_0')\cdot\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[g(\mathbf{A}x_0+\mathbf{b})\cdot g(\mathbf{A}x_0'+\mathbf{b})] \\
&= \frac{1}{2^{2N}}\cdot 2^N + \frac{1}{2^{2N}}\sum_{x_0\neq x_0'} h(x_0)h(x_0')\cdot\mu_g^2 \leq 2^{-N} + \mu_g^2\cdot\mu_h^2 \ .
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\mathop{\mathrm{Var}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})] &= \mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}\left[Z(\mathbf{A},\mathbf{b})^2\right] - \left(\mathop{\mathbb{E}}_{\mathbf{A},\mathbf{b}}[Z(\mathbf{A},\mathbf{b})]\right)^2 \\
&\leq 2^{-N} + \mu_g^2\cdot\mu_h^2 - \mu_g^2\cdot\mu_h^2 = 2^{-N} \ ,
\end{aligned}
$$

as desired. □

From here, a simple union bound establishes the following theorem, which is analogous to Corollary 17 from [ABG+14].

Corollary B.10 *Let $\eta, \gamma > 0$ be such that $1-\gamma-2\eta > 0$. Let $\mathcal{H}$ be any set of functions $h : \mathbb{F}_2^N \to \{-1,1\}$ of size $|\mathcal{H}| \le 2^{\gamma N}$. Then with probability at least $1 - 2^{-(1-\gamma-2\eta)N}$ over the choice of key $(\mathbf{A}, \mathbf{b})$, for all $h \in \mathcal{H}$, we have*

$$-2^{-\eta N} \le \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) - \mu_f \mu_h \le 2^{-\eta N} \ .$$

In particular, for all $h \in \mathcal{H}$,

$$-|\mu_g| - 2^{-\eta N} \le \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) \le |\mu_g| + 2^{-\eta N}$$

Proof. The first part is immediate by a union bound over $h \in \mathcal{H}$ combined with Lemma B.9, where we set $\varepsilon = 2^{-\eta N}$. For the second part, simply note that the first inequality is equivalent to

$$\mu_f \mu_h - 2^{-\eta N} \le \mathsf{Corr}(f_{\mathbf{A},\mathbf{b}}, h) \le \mu_f \mu_h + 2^{-\eta N} \ ,$$

and since $\mu_h \in [-1,1]$ it follows that $-|\mu_g| \le \mu_g \mu_h \le |\mu_g|$. □

C Attempt at Proving Resistance Against General Linear Tests for a Random Matrix A

We conjecture that our candidate is resistant against more general *non-adaptive linear tests* if the matrix key $\mathbf{A}$ is chosen randomly. Recall, we define $g : \{0,1\}^N \to \{-1,1\}$ to be the signed version of G: $g(x) \coloneqq (-1)^{G(x)}$. For integers $Q \ge 1$ and $k \ge 1$, fix a matrix $U = (u_{ij}) \in \mathbb{F}_2^{Q \times k}$ and consider

$$x^{(i)} \ = \ \sum_{j=1}^{k} u_{ij}\, v_j \ \in \ \mathbb{F}_2^N \qquad (i = 1, \ldots, Q),$$

where the $v_1, \ldots, v_k \in \mathbb{F}_2^N$ are drawn uniformly at random conditioned on being linearly independent. We require that the resulting queries $x^{(1)}, \ldots, x^{(Q)}$ are distinct and nonzero.[10] The test's bias, as described in Section4.2 is

$$\delta \ \coloneqq \ \mathbb{E}_{v_1,\ldots,v_k}\Big[\,(-1)^{\sum_{i=1}^{Q} G(x^{(i)})}\Big] = \mathbb{E}_{v_1,\ldots,v_k}\Big[\prod_{i=1}^{Q} g\big(x^{(i)}\big)\Big] \ . \tag{9}$$

Our goal is to show $|\delta| \le 2^{-\Omega(\lambda)}$.

Lemma C.1 (Reduction to a spectral sum) *Let $U \in \mathbb{F}_2^{Q \times k}$ have rank r, and write $S \coloneqq (s_1, \ldots, s_Q) \in (\mathbb{F}_2^N)^Q$. Then*

$$\delta \ = \ \sum_{S \in (\mathbb{F}_2^N)^Q :\, S \cdot U = 0} \ \prod_{i=1}^{Q} \widehat{g}(s_i),$$

where S is written as an $N \times Q$ matrix.

Proof. Expanding each $g(x^{(i)})$ into its Fourier expansion and plugging it into 9 gives:

$$\delta = \sum_{s_1,\ldots,s_Q} \Big(\prod_{i=1}^{Q} \widehat{g}(s_i)\Big) \mathbb{E}_{v_1,\ldots,v_k}\Big[\,(-1)^{\sum_{i=1}^{Q} \langle s_i,\, x^{(i)} \rangle}\Big].$$

Using $x^{(i)} = \sum_{j=1}^{k} u_{ij} v_j$ and independence of the v_j,

$$\mathbb{E}_{v_1,\ldots,v_k}\Big[(-1)^{\sum_i \langle s_i,\, x^{(i)} \rangle}\Big] = \prod_{j=1}^{k} \mathbb{E}_{v_j}\Big[(-1)^{\langle \sum_i u_{ij} s_i,\, v_j \rangle}\Big] = \prod_{j=1}^{k} \mathbf{1}\Big\{\sum_i u_{ij} s_i = 0\Big\},$$

since $\mathbb{E}_v[(-1)^{\langle t, v \rangle}] = \mathbf{1}\{t = 0\}$ for v uniform. This yields the claim. □

Since the number of possible $(s_1, \ldots, s_Q)$ in the spectral sum from Lemma C.1 is bounded by $2^{N(Q-r)} \le 2^{N(Q - \log Q)}$ and from Proposition 4.2, we have $|\widehat{g}(s)| \le O(2^{-\ell})$, using the Triangle Inequality, this gives the naive bound

$$|\delta| \ \le \ O(2^{Q(N-\ell) - N \log Q}) \ .$$

However, this bound is clearly too weak whenever $Q = \Omega(\lambda)$. We leave proving a better bound on the sum in Lemma C.1 as an open question.

[10] This excludes degenerate tests where repeated queries make the XOR trivially cancel.