

Algebraic Attacks against Random Local Functions and Their Countermeasures

Benny Applebaum* Shachar Lovett†

May 15, 2016

Abstract

Suppose that you have n truly random bits $x = (x_1, \dots, x_n)$ and you wish to use them to generate $m \gg n$ pseudorandom bits $y = (y_1, \dots, y_m)$ using a local mapping, i.e., each y_i should depend on at most $d = O(1)$ bits of x . In the polynomial regime of $m = n^s$, $s > 1$, the only known solution, originates from (Goldreich, ECCC 2000), is based on *Random Local Functions*: Compute y_i by applying some fixed (public) d -ary predicate P to a random (public) tuple of distinct inputs $(x_{i_1}, \dots, x_{i_d})$. Our goal in this paper is to understand, for any value of s , how the pseudorandomness of the resulting sequence depends on the choice of the underlying predicate. We derive the following results:

(1) We show that pseudorandomness against \mathbb{F}_2 -linear adversaries (i.e., the distribution y has low-bias) is achieved if the predicate is (a) $k = \Omega(s)$ -resilient, i.e., uncorrelated with any k -subset of its inputs, and (b) has algebraic degree of $\Omega(s)$ even after fixing $\Omega(s)$ of its inputs. We also show that these requirements are necessary, and so they form a tight characterization (up to constants) of security against linear attacks. Our positive result shows that a d -local low-bias generator can have output length of $n^{\Omega(d)}$, answering an open question of Mossel, Shpilka and Trevisan (FOCS, 2003). Our negative result shows that a candidate for pseudorandom generator proposed by the first author (computational complexity, 2015) and by O'Donnell and Witmer (CCC 2014) is insecure. We use similar techniques to refute a conjecture of Feldman, Perkins and Vempala (STOC 2015) regarding the hardness of planted constraint satisfaction problems.

(2) Motivated by the cryptanalysis literature, we consider security against *algebraic attacks*. We provide the first theoretical treatment of such attacks by formalizing a general notion of algebraic inversion and distinguishing attacks based on the Polynomial Calculus proof system. We show that algebraic attacks succeed if and only if there exist a degree $e = O(s)$ non-zero polynomial Q whose roots cover the roots of P or cover the roots of P 's complement. As a corollary, we obtain the first example of a predicate P for which the generated sequence y passes all linear tests but fails to pass some polynomial-time computable test, answering an open question posed by the first author (Question 4.9, computational complexity 2015).

*Tel-Aviv University, bennyap@post.tau.ac.il. Supported by the European Union's Horizon 2020 Programme (ERC-StG-2014-2020) under grant agreement no. 639813 ERC-CLC, ISF grant 1155/11, GIF grant 1152/2011, and the Check Point Institute for Information Security.

†University of California, San Diego, slovett@ucsd.edu. Supported by an NSF CAREER award 1350481 and a Sloan fellowship.

1 Introduction

The efficiency of cryptographic constructions is of both theoretical and practical interest. At the extreme, one would aim for an ultimate level of efficiency at the form of *constant*-parallel time implementation. Namely, the goal is to have cryptographic constructions in which each bit of the output depends only on a small constant number d of input bits, and so, each output can be individually computed with complexity that does not grow with the total input length or the level of security. This strong efficiency requirement seems hard to get as, at least intuitively, such form of locality may lead to algorithmic attacks.

The feasibility of locally-computable cryptography was established in [7]: It was shown that many cryptographic tasks admit a local implementation. This result is proven by “encoding” log-space computable cryptographic functions into *specially crafted* local functions whose input-output dependency graph has a very specific form. Despite this development, the hardness of local functions is far from being understood. Although we have basic feasibility results showing that *some* local functions are strong enough to be used in cryptographic applications, one may conjecture that *most* local functions carry some cryptographic hardness.

Assumption 1 (Random Local Functions are Hard (RLFH)). *For $d \geq 3$, let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a random d -local function that each of its output bits is computed by applying some fixed d -ary predicate P to a random set of d distinct inputs. Then, for a properly chosen predicate P and a properly chosen output length $m = m(n)$, the function f is likely to be a one-way function or even a pseudorandom generator (PRG).*

Assumption 1, whose one-wayness version was originally introduced by Goldreich [32], suggests that cryptographic hardness is not an isolated phenomenon that applies only to some special subclass of local functions, but rather it is common among all local functions. In the last fifteen years, the RLFH assumption was extensively studied (cf. [1, 3, 6, 8, 12, 13, 17, 24, 25, 30, 38, 39]), and was used to derive important applications in cryptography and complexity. Notable examples include secure computation protocols with constant computational overhead [35], a new public-key encryption scheme and hardness results for learning juntas [5], and strong inapproximability results for combinatorial problems [3].

More generally, the problem of inverting a random local function can be formulated as the problem of finding a satisfying assignment for a random constraint satisfaction problem (CSP) with a “planted solution”. Similarly, as discussed in [4, Section 4.3], the problem of breaking the pseudorandomness of random local functions is closely related to the algorithmic tasks of approximation and refutation of random CSP’s.¹ Such problems are well studied in complexity theory and are known to be “hard” in several aspects. In particular, the Cook-Levin theorem [18, 37] shows that it is NP-hard to exactly solve CSPs, while the PCP theorem [9, 10] shows that it is NP-hard even to approximate their solution. The study of the RLFH assumption complements this picture by providing a strong and clean formulation for the conjectured average-case hardness of CSPs under a natural distribution.² The validity of the RLFH assumption is therefore a central question in the study of computational intractability.

¹The former requires to approximate the maximum fraction of satisfiable constraints, and the latter requires to certify unsatisfiability.

²Other (related) conjectures about the average-case hardness of CSP problems were suggested by Feige [28] and were later extended by Barak et al. [11].

1.1 How to Choose a Hard Predicate?

Our goal in this paper is to understand how the pseudorandomness of a random local function depends on the choice of the underlying predicate P . In an attempt to understand the limitations of the RLF assumption, we will be interested in the most aggressive setting of parameters where $m = n^s$ for some constant $s > 1$. (This setting is also motivated by some of the aforementioned applications). We say that a predicate P is s -pseudorandom if, when sampling a random local function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ with the predicate P according to the procedure described in Assumption 1, the resulting function is likely to fool all efficiently computable tests.³ We ask:

Given an output length $m = n^s$, which predicates are s -pseudorandom?

Before presenting our results, let us review two important hardness criteria that were studied in previous works. (A more detailed account of related works is given in Section 1.5).

Resiliency. A key requirement for achieving pseudorandomness is high *resiliency*. A d -ary predicate P is k -resilient (also known as k -wise independent) if it has no nontrivial correlation with any linear combination of less than or equal to k of its inputs.⁴ For example, the Parity function on d bits has resilience of $d - 1$ which is the largest possible. Predicates with high resiliency were shown to resist a large class of attacks. For example, the results of [6] imply that, for every constant s , a predicate with resiliency 2 is s -pseudorandom against constant-depth (AC^0) circuits with sub-exponential size (see [4, Proposition 4.10]). O'Donnell and Witmer [39] proved that resiliency larger than $2s - 1$ yields s -pseudorandomness against attacks which are based on a large class of semidefinite programs. Similar bounds were proven by Feldman, Perkins and Vempala [30] for a wide family of *statistical algorithms* [29, 36]. Interestingly, the results of [39] and [30] are *tight* since algorithms from the corresponding classes can break pseudorandomness when the resiliency is smaller than $2s - 1$.

Algebraic Degree. The above results cover a rich family of attacks including most standard algorithmic approaches for solving CSPs (see discussion in [30, 39]), excluding one important technique: Gaussian Elimination. Indeed, a random local function instantiated with the d -ary parity, (whose resilience is $d - 1$) resists all the above attacks, but can be easily broken even for $m = n + 1$ by detecting a linear dependency in the outputs. More generally, as observed by [38], if the predicate P computes a *degree* ℓ -polynomial over \mathbb{F}_2 , then P cannot be ℓ -pseudorandom.

Resiliency + Degree are Sufficient for Pseudorandomness? In light of this, it is natural to conjecture that *high resiliency defeats all attacks except for Gaussian elimination*. Indeed, such views were taken by several researchers [30, 39] (see also [11] for a similar conjecture in a related context). In particular, it was suggested that s -pseudorandomness is achieved by any predicate with

³Formally, we require that for $1 - o(1)$ fraction of the functions, every distinguisher \mathcal{A} has no more than negligible distinguishing advantage, i.e., $|\Pr[\mathcal{A}_{x \leftarrow \{0, 1\}^n}(f(x)) = 1] - \Pr_{y \leftarrow \{0, 1\}^m}[\mathcal{A}(y)]| < n^{-\omega(1)}$. The adversary may depend on the description of the function, which in our case consists of the m tuples (I_1, \dots, I_m) that describe the input-output dependencies and the predicate P . In the cryptographic terminology, this corresponds to a *collection* of pseudorandom generators. See Section 2 for formal definitions.

⁴The special case of 0-resiliency corresponds to *balanced* predicates which are satisfied by exactly half of all possible assignments in $\{0, 1\}^d$. See Section 2 for formal definitions.

sufficiently large resiliency $k = k(s)$ and sufficiently large algebraic degree $\ell = \ell(s)$. Concretely, the $(k - 1)$ -resilient ℓ -degree predicate

$$(z_1 \oplus \cdots \oplus z_k) \oplus (z_{k+1} \wedge \cdots \wedge z_{k+\ell}),$$

denoted as XOR-AND $_{k,\ell}$, was suggested by the first author [4] and by O’Donnell and Witmer [39], as a candidate for achieving s -pseudorandomness for any $s < \min(k/2, \ell)$.

Applebaum, Bogdanov, and Rosen [6] confirmed this conjecture for the limited class of \mathbb{F}_2 -linear tests and for $s < 1.25$. They showed that any predicate whose resiliency and degree are at least 2, is s -pseudorandom against linear tests. That is, $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ is likely to be an ε -biased generator with sub-exponential bias of $\varepsilon = \exp(-n^{\Omega(1)})$. This result was later extended by [39] to $s < 1.5$ for the special case of the XOR-AND $_{3,2}$. It is important to mention that all known distinguishing attacks against random local function can be expressed as linear tests. Indeed, so far there was no example for a predicate which is 1.01-pseudorandom against linear tests but not 1.01-pseudorandom against general polynomial-time tests. It was therefore suggested that perhaps local functions are too simple to “separate” between these two different notions (cf. [24], [4, Question 4.9]).

1.2 Our Results

In this paper, we discuss two classes of distinguishing attacks, *linear* and *algebraic*, and develop necessary and sufficient conditions for achieving s -pseudorandomness against them.

1.2.1 Linear Tests

We revisit the case of linear distinguishers and show that in order to defeat such adversaries the predicate P must have, in addition to resiliency, high *bit-fixing degree*. Specifically, we say that P has r -bit fixing degree e if, by fixing r of the input bits of P , the minimal \mathbb{F}_2 -degree of a restriction is e . For example, 0-bit fixing degree is simply standard algebraic degree, the 1-bit fixing degree of d -ary AND is 0 (just fix one input to zero), and the r -bit fixing degree of d -ary Majority is at least $d - 2r$ for any $r < d/2$. We prove the following theorem:

Theorem 1.1. *Let P be a predicate with resiliency k and r -bit fixing degree e . Then for any $s > 1$ the following hold:*

1. *If $k < 2s - 1$ or $r + e < s$ then P is not s -pseudorandom against \mathbb{F}_2 -linear tests. In particular, for $1 - o(1)$ fraction of the functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ there will be an \mathbb{F}_2 -linear test with constant distinguishing advantage.*
2. *For any $\varepsilon > 0$, if $k, r, e \geq \Omega(s/\varepsilon)$ then P is s -pseudorandom against \mathbb{F}_2 -linear tests. In particular, $1 - o(1)$ fraction of the functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ are $\exp(-\delta n^{1-\varepsilon})$ -biased generators, where $\delta = \delta(d) > 0$.*

A more quantitative version of the theorem can be found in Sections 3 and 4. The theorem shows that resiliency k and r -bit fixing degree e with $k, r, e = \Theta(s)$ are both necessary and sufficient for fooling linear tests. In other words, the property of being s -pseudorandom against linear tests is characterized (up to the constants in the Θ notation) by the resiliency and bit fixing degree of the predicate. Our characterization is the first to hold for an arbitrary value of s . Previous sufficient conditions for s -pseudorandomness against linear tests were limited to $s < 1.5$ [5, 39].

We can also use Theorem 1.1 to improve the tradeoff between the locality and the output length of low-biased generators. Mossel, Shpilka and Trevisan [38] asked what is the best output length $m = n^s$ that can be achieved by any d -local low-bias generator (not necessarily based on random functions). They described a construction which achieves $s > \Omega(\sqrt{d})$ and posed the possibility of improving this to $\Omega(d)$ as an open question. We provide an affirmative answer to this question.

Corollary 1.2. *For every integer constant $s > 1$ and $\varepsilon > 0$ there exists a $\Theta(s/\varepsilon)$ -local function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ which is $\exp(-n^{1-\varepsilon})$ -biased generator. In particular, this holds for a random local function instantiated with the XOR-MAJ $_{k,\ell}$ predicate $(z_1 \oplus \dots \oplus z_k) \oplus \text{MAJ}(z_{k+1}, \dots, z_{k+\ell})$, with $k, \ell = \Theta(s/\varepsilon)$.*

Proof. It is not hard to verify that XOR-MAJ $_{k,\ell}$ is k -resilient and has r -bit fixing degree e for any r and e for which $r + e \leq \ell/2$. To see the latter, observe that given any fixing of r variables in a Majority, we can fix another r variables to complementary values, and obtain Majority over $\ell - 2r \geq 2e$ inputs, which has \mathbb{F}_2 -degree $\geq e$. (In general, the algebraic degree of t -wise majority is $2^{\lceil \log t \rceil}$.) The corollary then follows from the second part of Theorem 1.1. \square

One can also derive interesting insights from the first part (lower-bound) of Theorem 1.1. In particular, it shows that, in contrast to prior beliefs, high resiliency and high degree are not sufficient for achieving pseudorandomness, even against linear distinguishers.

Corollary 1.3. *For every constants k and ℓ there exists a k -resilient ℓ -degree predicate P which is not 2.01-pseudorandom against linear tests. In particular, this holds for the XOR-AND $_{k,\ell}$ predicate.*

Proof. It is not hard to verify that XOR-AND $_{k,\ell}$ has 1-bit fixing degree of 1 (simply fix any of the “AND” inputs to zero). The corollary therefore follows from the first part of Theorem 1.1. \square

We use similar techniques to refute a conjecture of Feldman et al. [30, Conjecture 1] regarding the average-case hardness of planted CSPs. For a predicate P , they consider the experiment where a planted solution x is chosen uniformly at random, and then m i.i.d. clauses are sampled uniformly from all clauses C for which the restriction of x to the clause C satisfies the predicate P (a formal definition appears in Section 6). [30] conjecture that if the predicate is $(r - 1)$ -resilient and its algebraic degree is larger than $r/2$ then it is hard to recover the planted assignment given $m = n^s$ clauses for any $s < r/2$. In Section 6 we refute this conjecture.

1.2.2 Algebraic Attacks

Motivated by the cryptanalysis literature, we consider security against *algebraic attacks*. Algebraic attacks against a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ start with an output y (presumably in the image of f) and use it to initialize a system of polynomial equations over the hidden input variables $x = (x_1, \dots, x_n)$. The system is further manipulated and extended (e.g., by multiplying the polynomials by some low-degree polynomial) until a solution is found (e.g., via linearization and Gaussian elimination or by computing a Gröbner basis of the expanded system).

We provide the first theoretical treatment of algebraic attacks by formalizing a general notion of algebraic inversion and distinguishing attacks. Roughly speaking, we parameterize an algebraic attack by an efficient *scheduling algorithm* which generates in each step a new polynomial equation via the derivation rules of the *Polynomial Calculus* proof system [15]. That is, by adding a pair of existing equations or by multiplying an existing equation by a formal variable x_i . We consider

two different variants of attacks: algebraic *Inversion* attacks in which the adversary attempts to recover the preimage of y (i.e., recover a consistent solution for the input variables x); and algebraic *Refutation* attacks in which the adversary tries to prove that y is not in the image of f (i.e., to show that the system of equations has no solution). We observe that our first variant indeed captures standard algebraic inversion algorithms (e.g., [22, 26, 27]). (See Section 5 for more details).

We characterize s -pseudorandomness against algebraic inversion and refutation attacks via the notion of *rational degree*. The *rational degree* of a predicate is the smallest integer e for which there exist degree e polynomials Q, R , not both zero, such that $PQ = R$.⁵ It can be shown that a rational degree e implies r -bit fixing degree of at least $e - r$ for any $r < e$, but the converse does not necessarily hold. As a simple example, for any e and r , the $\text{OR}_e \circ \text{XOR}_r$ predicate

$$\bigvee_{i \in [e]} \bigoplus_{j \in [r]} z_{i,j}$$

has an $(r - 1)$ -bit fixing degree e , but has a rational degree of 1 since its roots satisfy the linear equation $\bigoplus_{i \in [k], j \in [\ell]} z_{i,j} = 0$. On the other extreme, the d -ary majority predicate has a rational degree of $\lceil d/2 \rceil$ which is the largest possible [23]. (See Section 2).

The following theorem shows that rational degree characterizes s -pseudorandomness against algebraic attacks. (See Section 5 for a more quantitative version).

Theorem 1.4. *Let P be a predicate with rational degree e , then for any s the following hold:*

1. *If $e < s$ then P is not s -pseudorandom against algebraic attacks. In particular, for $1 - o(1)$ fraction of the functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$ there is an efficient algebraic refutation attack that, for all but $o(1)$ -fraction of the string $y \notin \text{Im}(f)$, certifies that y has no preimage.*
2. *If $e > 8s + 1$ then P is s -pseudorandom against algebraic tests. In particular, for $1 - o(1)$ fraction of the functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n^s}$, inversion/refutation algebraic attacks with sub-exponential time-complexity of $\exp(n^{1-O(s/e)})$ fail to invert/refute any $y \in \{0, 1\}^{n^s}$.*

We can show that algebraic attacks are incomparable to linear attacks. (See Section 1.3 for further discussion).

Corollary 1.5. *For every integer constant $s > 1$ there exists a predicate P_1 which is s -pseudorandom against linear tests but is not even 2.01-pseudorandom against algebraic attacks and a predicate P_2 which is s -pseudorandom against algebraic attacks but is not even 1.01-pseudorandom against linear tests.*

Proof. For integers k, r , and e , let $P_1 : \{0, 1\}^{k+1+(r+1)e}$ be the predicate

$$(w_1 \oplus \cdots \oplus w_{k+1}) \oplus (\text{OR}_e \circ \text{XOR}_{r+1}(z)), \quad \text{where } z = (z_{i,j})_{i \in [r+1], j \in [e]}.$$

It is not hard to verify that P_1 has resiliency k and r -bit fixing degree e , hence, for sufficiently large $k, r, e = \Omega(s)$, it is s -pseudorandom against linear tests. However, P_1 has rational degree of 2 since any root of $P_1(w, z) = 0$ is also a root of the degree-2 polynomial $(\bigoplus_{i \in [k+1]} w_i \oplus 1) \cdot (\bigoplus_{i \in [k], j \in [\ell]} z_{i,j})$. (To see this note that the XOR part must be equal to the $\text{OR} \circ \text{XOR}$ part, and so either the

⁵An equivalent definition is the minimal degree of a non-zero polynomial Q whose roots cover the roots of P or cover the roots of P 's complement. This is called *algebraic immunity* in the applied cryptographic literature. see Section 2.

XOR part equals to one and the first multiplicand vanishes or the OR \circ XOR equals to zero and the second multiplicand vanishes).

The predicate P_2 is simply MAJ_ℓ for $\ell = 16s + 4$. This predicate has a rational degree $\lceil \ell/2 \rceil$ and so it is s -pseudorandom against algebraic attacks, but is not even 1-resilient and so it fails to be 1.01-pseudorandom against linear tests. \square

The corollary provides the first example of a predicate P which is s -pseudorandom against linear tests but not s -pseudorandom against polynomial-time computable tests, answering an open question posed by the first author [4, Question 4.9].

1.3 Discussion and Conclusion

Our results indicate that linear-algebraic attacks are more powerful than it appears. Even in the simple setting of random local functions, they can be used in unexpected ways and defeat intuitive countermeasures like high algebraic degree or low-bias. Quoting Cook, Etesami, Miller and Trevisan [16]:

[An] interesting goal would be to show that no “variation of Gaussian elimination” can invert Goldreichs function when nonlinear predicates are used. Unfortunately, it is not clear how to even formalize such a statement.

Our new results provide more tools (rational degree, algebraic attacks) to address this question, and can be seen as a first step toward a better theoretical understanding of the power of linear-algebraic attacks.

We leave several challenges for future works. First, it will be interesting to find a unified class of adversaries that capture both linear attacks and algebraic attacks (recall that these classes are incomparable). Roughly speaking, algebraic attacks can exploit relatively complicated (non-linear) relations among the outputs of a function f as long as these relations hold for *all* possible inputs. Unfortunately, such attacks completely miss relations that hold for *most* inputs, even if these relations are simple (e.g., linear).⁶ Can we enrich the model of algebraic attacks in a way that will take into account relations that hold with high probability over a random input? It will also be interesting to study pseudorandomness against low-degree distinguishers which, in a sense, interpolate between algebraic attacks and linear attacks.

Our results show that any k -resilient predicate with rational degree e is s -pseudorandom against linear attacks and algebraic attacks as long as $k \geq 5s$ and $e > 18s$. Are there efficient attacks against such predicates? As a concrete challenge, one may try to break the s -pseudorandomness of the XOR- $\text{MAJ}_{a,b}$ predicate

$$(z_1 \oplus \cdots \oplus z_a) \oplus \text{MAJ}(z_{a+1}, \dots, z_{a+b}),$$

with $a \geq 5s$ and $b > 36s$.

Finally, one may ask what is the smallest arity d for which there exists a predicate $P : \{0, 1\}^d \rightarrow \{0, 1\}$ with a rational degree e and resiliency k . The XOR- $\text{MAJ}_{k,2e}$ predicate shows that $d \leq k + 2e$. For the case of algebraic degree ℓ , Siegenthaler [42] proved a tight lower-bound of $d \geq \ell + k + 1$ for every $\ell > 1$. An improved lower-bound on d (in terms of k and e) would lead to better upper-bounds on the number of pseudorandom bits n^s that can be generated by a random d -local function.

⁶Indeed, this is exactly what happens for predicates with high rational degree and low resiliency (such as majority) which are pseudorandom against algebraic attacks but can be broken by linear attacks.

1.4 Techniques

In this section we briefly sketch the main ideas behind the proofs of our main results.

1.4.1 Linear Tests (Theorem 1.1)

Let $m = n^s$ and consider a random local function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ instantiated with a predicate P , and recall that f is described by m (random) d -tuples I_1, \dots, I_m where each tuple consists of d distinct elements (input variables) from $[n]$. In the the first part of Theorem 1.1, we show that if P either has (i) small resiliency; or (ii) small bit fixing degree, then linear tests can distinguish the output of f from random. That is, $f(x)$ is not a small-bias generator.

Consider first the case that P has resiliency $k < 2s - 1$. In this case P is correlated to some linear combination of $\ell \leq k + 1 < 2s$ of its inputs, say the first ℓ . So $\hat{P}(\gamma^*) \neq 0$ for $\gamma^* = 1^\ell 0^{d-\ell}$. As P is d -local we in fact have $|\hat{P}(\gamma^*)| \geq 2^{-d}$. By a birthday paradox calculation (as $m \gg n^{\ell/2}$), with high probability there exist indices $i, j \in [m]$ such that I_i, I_j agree on their first ℓ coordinates, and where all the remaining coordinates of I_i, I_j are distinct. A Fourier-based calculation then shows that

$$\mathbb{E}_x[(-1)^{f(x)_i + f(x)_j}] = \mathbb{E}_x[(-1)^{P(x|_{I_i}) + P(x|_{I_j})}] = \hat{P}(\gamma^*)^2 \geq 2^{-2d}.$$

Full details are given in Lemma 3.1.⁷

The second case is where P has r -bit fixing degree e where $r + e < s$. Assume for simplicity that by fixing the first r bits of P to zero, we obtain a polynomial Q of degree e . Then

$$P(z_1, \dots, z_d) = z_1 Q_1(z) + \dots + z_r Q_r(z) + R(z_{r+1}, \dots, z_d).$$

Let $A = \{i \in [m] : (I_i)_1 = 1, \dots, (I_i)_r = r\}$ be all output bits, whose first r input bits are x_1, \dots, x_r . With high probability, $|A| = \Theta(n^{s-r}) \gg n^e$. Note that whenever $x_1 = \dots = x_r = 0$ then we have $f(x)_i = R(x|_{I_i})$, which is a polynomial of degree e . As $|A| \gg n^e$ the polynomials $\{R(x|_{I_i}) : i \in A\}$ are linearly dependent, and hence satisfy some linear equations (in fact, many linear equations). Each of these equations defines a linear test over the outputs. All these tests output zero simultaneously whenever $x_1 = \dots = x_r = 0$, which is an event which occurs with probability 2^{-r} . Assuming, that we have more than $r + 1$ such tests, this means that their joint distribution is not uniform and so we can find a linear combination of these (linear) tests that has a constant bias. Full details are given in Lemma 3.2.

In the second part of Theorem 1.1 (given formally in Corollary 4.2), we show that when P has resiliency k and r -bit fixing degree e for $k, r, e = \Omega(s)$ then with high probability over the choice of f , the distribution of $f(x)$ fools all linear tests (that is, it is a small bias distribution with sub-exponentially small bias). Concretely, we show that this is the case when the underlying hypergraph G corresponding to the choice of I_1, \dots, I_m is sufficiently expanding (which holds with probability $1 - o(1)$ over a random choice of I_1, \dots, I_m).

Following [38], we distinguish two types of linear tests: *light tests*, which have at most t nonzero coefficients, and *heavy tests*, which have more than t nonzero coordinates where $t = n^{1-\varepsilon}$ is an expansion parameter. The case of light tests was essentially handled by [6]. If $A \subset [m]$ and $|A| \leq t$, then in fact

$$\mathbb{E}_x \left[(-1)^{\sum_{i \in A} f(x)_i} \right] = \mathbb{E}_x \left[(-1)^{\sum_{i \in A} P(x|_{I_i})} \right] = 0.$$

⁷A similar birthday-type argument can be used to invert the function using spectral techniques. See [39].

This follows from the fact that if the underlying hypergraph is expanding, then there is some $i \in A$ such that I_i has very low intersection with all other $I_j, j \in A$. Concretely, $|I_i \cap (\cup_{j \in A, j \neq i} I_j)|$ is upper-bounded by the resiliency parameter k . Hence the sum $\sum_{i \in A} P(x|_{I_i})$ consists of at least one independent unbiased coin, and so the test has zero bias. The details are given in Lemma 4.3.

Ruling out heavy tests is more subtle. Fix $A \subset [m]$ of size $|A| \geq t$, where our goal is to bound the parity of $\oplus_{i \in A} f(x)_i$. We follow a sequence of steps in which we fix some input variables to random values. At the end of this process, some subset $K \subset [n]$ of variables is still non-fixed. We furthermore obtain a subset $A^* \subset A$ of output variables such that

- For each $i \in A^*$, the sets $K_i = I_i \cap K$ are disjoint and of size $|K_i| \geq d - r$.
- For each $i \in A \setminus A^*$, we have $|I_i \cap K| < e$.

The expansion properties of the underlying hypergraph allows to argue that $|A^*| \geq \Omega_d(|A|)$. We may thus rewrite $\sum_{i \in A} f(x)_i$, for any fixing of the input variables outside K , as $\sum Q_i(x^i) + R(x)$, where $\{x^i : i \in A^*\}$ are disjoint subsets of variables; each Q_i is a polynomial obtained by fixing at most r inputs in P , and hence by our assumption $\deg(Q_i) \geq e$; and R is a polynomial whose degree is less than e . We then apply the Gowers-Cauchy-Schwarz inequality [33] to show that the bias is exponentially small in A^* . The details are given in Lemma 4.4.

1.4.2 Algebraic Attacks (Theorem 1.4)

The first part of Theorem 1.4 is based on a simple intuition. Given the output y of a random local function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, we can write m polynomial equations over the formal input variables $x = (x_1, \dots, x_n)$

$$P(x|_{I_i}) = y_i \quad \forall i \in [m].$$

Now, assume that any root x of P also satisfies the equation $Q(x) = 0$ for some degree e polynomial Q . Then, we can replace any equation for which $y_i = 0$, with a degree e equation of the form $Q(x|_{I_i}) = 0$. Intuitively, $m = \Omega(n^e)$ such equations should suffice to solve the system via “linearization”: replace each degree e monomial $\prod_{i \in S} x_i$ with a new formal variable X_S , solve the linear system over the X ’s, and recover the original solution (say by looking at the monomials x_i). The success of this approach is based on the assumption that the linearized system has a unique solution.

Unfortunately, this assumption is overoptimistic. Our linear equations come from a probability distribution whose structure strongly depends on the d -ary predicate Q . As a result, the system may have a large number of solutions which do not correspond to solutions of the original system. For example, consider the predicate $Q(z_1, \dots, z_d) = z_1 z_2 z_3 + z_1 z_2 + z_1 z_3 + z_2 z_3$, which translates into linear equations of the form $X_{\{i,j,k\}} + X_{\{i,j\}} + X_{\{i,k\}} + X_{\{j,k\}} = 0$. This means that we will always get equations that contain correlated monomials (with common variables) and that some monomials will never appear (e.g., degree 1 monomials). Consequently, the system will not have a unique solution (e.g., any solution with $X_{\{i,j,k\}} = X_{\{i,j\}} + X_{\{i,k\}} + X_{\{j,k\}}$ would be a valid solution).

We bypass this limitation by exploiting the information given in the original P -equations (whose RHS was 1), and by resorting to refutation attack as opposed to inversion. In particular, we show that, in order to identify that the system $f(x) = y$ is unsatisfiable, it suffices to examine a few P equations together with the linearized equations that are spanned by the Q -constraints. Furthermore, by focusing on a small set of input variables, we can certify the unsatisfiability of the system $f(x) = y$ via an efficient algebraic attack. See Theorem 5.1 for details.

The second part of Theorem 1.4 asserts that if P has a high rational degree then it resists refutation and inversion algebraic attacks. This part is based on the work of Alekhovich and Razborov [2] who showed that, under some conditions, the unsatisfiability of a system of polynomial equations $\{P_i(x|_{I_i}) = 0\}$ does not have a short proof in the polynomial calculus (PC) proof system.⁸ Specifically, for their lower-bound one has to assume that the set system (I_1, \dots, I_m) has good expansion, and that each of the predicates P_i is *e-immune* in the sense that there is no degree- e non-zero polynomial Q_i which satisfies $P_i(x) = 0 \Rightarrow Q_i(x) = 0$. These conditions are met in our setting. Indeed, the tuples I_i 's are chosen at random and so they are likely to be expanding. Also, the system $\{P(x|_{I_i}) = y_i\}$ can be written as $\{P_i(x|_{I_i}) = 0\}$ where P_i is either P or its complement, and therefore a rational degree of e translates into e -immunity a-la Alekhovich and Razborov. Since the time-complexity of a refutation algebraic attack is lower-bounded by the size of the smallest PC proof for unsatisfiability, we derive a lower-bound against algebraic refutation attacks. The resulting lower-bound is quite strong as it holds for almost all functions f and for every $y \notin \text{Im}(f)$.

To handle algebraic inversion attacks we reduce inversion to refutation and apply the previous lower-bound. Specifically, we show that algebraic refutation of a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ follows from algebraic inversion of the function $f' : \{0, 1\}^n \rightarrow \{0, 1\}^{m-1}$ that computes the $(m-1)$ -prefix of f . Indeed, an attack that recovers a preimage $b \in \{0, 1\}^n$ of (y_1, \dots, y_{m-1}) under f' can be used to show that the system $f(x) = (y, 1 - f_m(b))$ has no solution (here f_m denotes the last output of f). First, recover b from the prefix y , and then derive the equation $f_m(x) = f_m(b)$ which contradicts the last output equation $f_m(x) = 1 - f_m(b)$. This general transformation from inversion to refutation has some overhead. The complexity of the refutation attack grows (as it takes time to derive the equation $f_m(x) = f_m(b)$), and the fraction of outputs y 's for which it applies is tiny (at best $|\text{Im}(f')|/2^m$). Still, we show that in our case where f is sufficiently simple⁹ and the refutation lower-bound holds for any y , it is possible to derive a strong inversion lower-bound as stated in Theorem 1.4.

1.5 Other Related Works

Pseudorandomness of Random Local Functions. The study of locally computable PRGs was initiated by Cryan and Miltersen [24] who asked whether such generators exist. Mossel, Shpilka and Trevisan gave the first construction of low-bias generators with constant locality (of 5) by “plugging” the XOR-AND_{3,2} predicate to a specially crafted dependency graph. They were also the first to identify the importance of algebraic degree and resiliency in this context. Alekhovich [1] conjectured that random local functions instantiated with a noisy 3-parity predicate (which flips its output with some small constant probability) generate a pseudorandom sequence. Applebaum, Ishai and Kushilevitz [7] constructed PRGs $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with constant locality based on the intractability of standard cryptographic assumptions (e.g., the intractability of factoring, discrete logarithms, and lattice problems). However, the stretch of this construction $m - n$ is

⁸The paper actually proves a degree lower-bound, however, a general theorem of Impagliazzo et al. [34, Corollary 6.3] allows to conclude a size lower-bound.

⁹More generally, the efficiency loss is proportional to the size of the smallest arithmetic *skew* circuit that computes the last output of f , where an arithmetic circuit (over the binary field) is skew if each of its multiplication gates involves at least one argument that is an input variable. It is known that any language computable in log-space (or even non-deterministic log-space) has a polynomial-size skew-circuit [43], and therefore for such functions the reduction incurs only a polynomial overhead.

inherently sub-linear (i.e., $m - n = o(n)$). The stretch was improved in [8] to linear ($m - n = \Omega(n)$) by derandomizing Alekhnovich’s proposal. A local PRG with polynomial stretch ($m - n = n^{1+\varepsilon}$) and an inverse polynomial distinguishing advantage (as opposed to negligible) was later constructed by [3]. It is also showed in [3] that, for a large family of predicates, the one-wayness version of the RLF assumption implies its pseudorandomness version. A similar result was first proved by [5] for the special case of the noisy parity predicate. The one-wayness of random local functions was studied in [12, 13, 17]. A detailed survey on the cryptographic hardness of random local functions appears in [4].

Algebraic attacks and Polynomial Calculus. Algebraic attacks were originally presented by Patarin [40] and were later extended and abstracted by Courtois and Meier [19, 20, 23]. As already mentioned the basic idea (which can be traced back to Shannon’s work [41]) is to recover the secret x by solving a system of multivariate algebraic equations. The hope is that by recognizing low-degree input/output relations it would be possible to construct an over-defined system of low-degree equations. While solving such a system is NP-hard in general, it may be solvable in practice via the use of Gröbner basis or by linearizing the system (cf. [21, 22, 26, 27]). Unfortunately, these methods are typically not amenable to a formal analysis. In Theorem 1.4 we derive a formal analysis of such attacks based on the specific structure of random local functions and on the framework of the Polynomial Calculus (PC) proof system [15]. The connection between PC and algebraic attacks is natural in retrospect, but, to the best of our knowledge, was never exploited before. We further mention that in the applied cryptography literature high rational degree is considered to be a necessary condition (referred to as *algebraic immunity*) for resisting algebraic attacks against stream ciphers (see the survey of [14]). Our results provide a concrete (yet different) setting in which this requirement can be rigorously justified.

Acknowledgement

We would like to thank Claude Carlet for helpful discussions.

2 The Model

For a finite set S we let S_d denote the set of ordered d -tuples with distinct elements in S . In particular, $[n]_d$ denotes the set of ordered d -tuples with distinct elements in the set $[n] := \{1, \dots, n\}$. For $x \in \{0, 1\}^n$ and $I = (i_1, \dots, i_d) \in [n]_d$, let $x|_I = (x_{i_1}, \dots, x_{i_d})$ denote the restriction of x to I . For a d -ary predicate $P : \{0, 1\}^d \rightarrow \{0, 1\}$, and a tuple $G = (I_1, \dots, I_m)$, where each $I_i \in [n]_d$, we let $f_{G,P} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ denote the d -local function whose i -th output bit is set to be $P(x|_{I_i})$, that is

$$f_{G,P}(x) = (P(x|_{I_1}), \dots, P(x|_{I_m})).$$

We will treat I_i either as ordered tuples, or unordered sets, depending on the context: when we evaluate $x|_{I_i}$ we treat them as ordered sets; but when we discuss intersections or unions of them, we treat them as unordered sets. We sometimes view G as a hypergraph over n vertices with m ordered hyperedges I_1, \dots, I_m each of cardinality d . Correspondingly, we refer to G as the input-output dependency hypergraph of $f_{G,P}$. We say that the hypergraph $G = (I_1, \dots, I_m)$ is (r, c) expanding if for every set $A \subseteq [m]$ of cardinality at most r , it holds that $|\cup_{i \in A} I_i| > c|A|$. We define

a probability distribution $\mathcal{G}_{n,m,d}$ over hypergraphs $G = (I_1, \dots, I_m)$ by choosing each $I_i \in [n]_d$ independently and uniformly at random. We let $\mathcal{F}_{P,n,m}$ denote the probability distribution over d -local functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ obtained by sampling $G \leftarrow \mathcal{G}_{n,m,d}$ and letting $f = f_{G,P}$.

We define a few notions required of a strong predicate. The Fourier coefficients of P are

$$\hat{P}(\gamma) = \mathbb{E}_{z \in \{0,1\}^d} \left[(-1)^{P(z) + \langle z, \gamma \rangle} \right],$$

where $\gamma \in \{0, 1\}^d$.

Definition 1 (resilience). *The predicate P has resilience k if it has no nontrivial correlation with any linear combination of less than or equal to k of its inputs. That is, $\hat{P}(\gamma) = 0$ for all $\gamma \in \{0, 1\}^d$ of hamming weight $\leq k$.*

For example, the parity function on d bits has resilience $d - 1$, while the majority function is not even 1-resilient (i.e., has resilience 0). For a boolean function $Q : \{0, 1\}^k \rightarrow \{0, 1\}$, its \mathbb{F}_2 -degree is the degree of the \mathbb{F}_2 polynomial it computes. We denote it by $\deg_{\mathbb{F}_2}(Q)$.

Definition 2 (bit-fixing degree). *The predicate P has r -bit fixing degree e if, taking the minimum over all restrictions of P by fixing r bits, the minimal \mathbb{F}_2 -degree of a restriction is e . That is,*

$$\min \left(\deg_{\mathbb{F}_2}(P|_{z_{i_1}=b_1, \dots, z_{i_r}=b_r}) : i_1, \dots, i_r \in [d], b_1, \dots, b_r \in \{0, 1\} \right) = e.$$

Definition 3 (rational degree). *The predicate P has rational degree e if there exist $Q, R : \{0, 1\}^d \rightarrow \{0, 1\}$ of degrees $\deg_{\mathbb{F}_2}(Q), \deg_{\mathbb{F}_2}(R) \leq e$, where $R \neq 0$, such that*

$$P(z)R(z) + Q(z) = 0 \quad \forall z \in \{0, 1\}^d.$$

Remarks:

1. As observed by [23] the rational degree is at most $\lceil d/2 \rceil$. Indeed, if the \mathbb{F}_2 -degree of P is larger than $\lceil d/2 \rceil$ (the other case is trivial) then we can always find a degree $\lceil d/2 \rceil$ polynomial R whose product with P has degree of at most $\lceil d/2 \rceil$. It is not hard to prove (e.g., by induction on d) that d -wise majority achieves this bound.
2. An equivalent condition to rational degree $\leq e$ is that there exists a nonzero $Q : \{0, 1\}^d \rightarrow \{0, 1\}$ with $\deg_{\mathbb{F}_2}(Q) \leq e$ such that $P(z) = 0 \Rightarrow Q(z) = 0$ or $P(z) = 1 \Rightarrow Q(z) = 0$. For one direction, note that if $P(z)R(z) + Q(z) = 0$ then $P(z) = 0 \Rightarrow Q(z) = 0$ and $P(z) = 1 \Rightarrow R(z) + Q(z) = 0$. It can be that one of Q or $R + Q$ is identically zero, but not both. For the other direction, if $P(z) = 0 \Rightarrow Q(z) = 0$ then $P(z)Q(z) + Q(z) = 0$, and if $P(z) = 1 \Rightarrow Q(z) = 0$ then $P(z)Q(z) = 0$.
3. We note that a rational degree of e implies an r -bit fixing degree of at least $e - r$ for any $r < e$. Indeed, assume that by fixing the inputs $(z_{i_1}, \dots, z_{i_r})$ of P to (b_1, \dots, b_r) the polynomial P simplifies to a polynomial Q of degree $\ell < e - r$. Then, every root of P is also a root of the polynomial

$$\left(1 - \prod_{j=1}^r (z_{i_j} - b_{i_j}) \right) \cdot Q(z)$$

whose degree is $\ell + r < e$, in contrast to our assumption regarding the rational degree of P .

3 Linear Attacks

We prove the first part of Theorem 1.1 via the following two lemmas. Similar techniques are used to refute the FPV conjecture regarding the hardness of planted CSPs in Section 6.

Lemma 3.1. *Assume that P has resilience strictly smaller than k , and let $m = an^{k/2}$ for some constant $a > 0$. Then, except with probability $\exp(-\Omega(a^2))$ over the choice of $f \leftarrow \mathcal{F}_{P,n,m}$, there exist distinct $i, j \in [m]$ such that*

$$\mathbb{E}_{x \in \{0,1\}^n} \left[(-1)^{f(x)_i + f(x)_j} \right] \geq 2^{-2d}.$$

In particular, the output of f is 2^{-2d} biased.

The special case of $k = 2$ was proven in [6] via a different argument. We further mention that in the above regime one can efficiently invert the function, see [4].

Proof. Assume without loss of generality that P is correlated to the sum of its first $\ell \leq k$ inputs. That is, for $\gamma^* = 1^\ell 0^{d-\ell}$ we have $\hat{P}(\gamma^*) \neq 0$. Since $\hat{P}(\gamma^*)$ is the average of 2^d elements in $\{-1, 1\}$, we in fact have $|\hat{P}(\gamma^*)| \geq 2^{-d}$.

Let $G = (I_1, \dots, I_m)$ be the dependency hypergraph of a function f . Call a pair of outputs $i, j \in [m]$ *good* if I_i, I_j agree on the first ℓ coordinates, while the last coordinates of I_i and I_j are completely distinct. That is,

$$\begin{aligned} I_i &= (p_1, \dots, p_\ell, q_1, \dots, q_{d-\ell}) \\ I_j &= (p_1, \dots, p_\ell, r_1, \dots, r_{d-\ell}) \end{aligned}$$

where $\{p_1, \dots, p_\ell, q_1, \dots, q_{d-\ell}, r_1, \dots, r_{d-\ell}\}$ are all distinct.

First, we show that if i and j is a good pair, then f_i and f_j are correlated. Indeed,

$$\begin{aligned} \mathbb{E}_x [(-1)^{f(x)_i + f(x)_j}] &= \mathbb{E}_x \left[(-1)^{P(x|_{I_i}) + P(x|_{I_j})} \right] \\ &= \sum_{\gamma_1, \gamma_2 \in \{0,1\}^d} \hat{P}(\gamma_1) \hat{P}(\gamma_2) \mathbb{E}_x \left[(-1)^{\langle \gamma_1, x|_{I_i} \rangle + \langle \gamma_2, x|_{I_j} \rangle} \right]. \end{aligned}$$

Observe that the term $\mathbb{E}_x \left[(-1)^{\langle \gamma_1, x|_{I_i} \rangle + \langle \gamma_2, x|_{I_j} \rangle} \right]$ is zero, unless $\gamma_1 = \gamma_2 = \gamma$, and moreover γ is supported on the first ℓ bits. However, in this case $\hat{P}(\gamma) = 0$, unless $\gamma = \gamma^*$. So, we conclude that

$$\mathbb{E}_x [(-1)^{f(x)_i + f(x)_j}] = \hat{P}(\gamma^*)^2 \geq 2^{-2d}.$$

Next, we show that a randomly chosen $f \leftarrow \mathcal{F}_{P,n,m}$ contains a good pair with probability $1 - \exp(-\Omega(a^2))$. Assume that G is sampled in two steps: (1) the first ℓ elements of each hyperedge are chosen; and (2) for each hyperedge the remaining $d - \ell$ elements are sampled. The birthday paradox guarantees that, with probability $1 - \exp(-\Omega(a^2))$, after the first step there will be a pair of outputs which share the first ℓ entries. Such a pair turns to be good with probability at least $1 - o(1)$ over the choice of the second step. The lemma follows. \square

Remark 3.1. *While the lemma is restricted to the case where a is a constant, it implies that when $a = a_n$ grows with n , a random function $f \leftarrow \mathcal{F}_{P,n,a_n n^{k/2}}$ is likely to be 2^{-2d} -biased with probability at least $1 - \exp(-\Omega(a_n))$. Indeed, in this case one can view f as a concatenation of a random functions selected from $\mathcal{F}_{P,n,n^{k/2}}$, and apply the lemma to each of them separately.*

Lemma 3.2. *Assume that P has r -bit fixing degree e . Let $m \geq \Omega(n^{r+e})$. Then, with probability $1 - o(1)$ over the choice of $f \leftarrow \mathcal{F}_{P,n,m}$, the distribution of f is $2^{-(r+1)}$ biased.*

Proof. Assume without loss of generality that by fixing the first r input bits of P to zero, its degree reduces to e . That is,

$$P(z_1, \dots, z_d) = z_1 Q_1(z) + \dots + z_r Q_r(z) + R(z_{r+1}, \dots, z_d),$$

where $\deg_{\mathbb{F}_2}(R) \leq e$.

Let $G = (I_1, \dots, I_m)$ be the random dependency hypergraph of $f \leftarrow \mathcal{F}_{P,n,m}$. Define $A \subset [m]$ by

$$A = \{i \in [m] : (I_i)_1 = 1, \dots, (I_i)_r = r\}.$$

By standard concentration bounds, with high probability over the choice of G , we have that $|A| \geq \Omega(m/n^r) \geq \Omega(n^e)$. From now on, we restrict our attention to $i \in A$.

The set of polynomials $\{R(x|_{I_i}) : i \in A\}$ are polynomials of degree e in n variables. Hence, they span a linear subspace of dimension $O(n^e)$. So, there must be a linear combination of $\{R(x|_{I_i}) : i \in A\}$ equal to zero. In fact, by choosing the unspecified constants above so that $|A| \geq 2n^e$, say, there are many linearly independent such linear combinations. Specifically, for $T = r + 1$, we have vectors $\lambda_t = (\lambda_{t,i} : i \in A) \in \mathbb{F}_2^A$ for $t \in [T]$, linearly independent, such that

$$\langle \lambda_t, R(x) \rangle = \sum_{i \in A} \lambda_{t,i} R(x|_{I_i}) = 0,$$

where we denote $R(x) = (R(x|_{I_i}) : i \in A)$. This implies that

$$\langle \lambda_t, f(x) \rangle = x_1 H_{t,1}(x) + \dots + x_r H_{t,r}(x),$$

for some polynomials $H_{t,1}, \dots, H_{t,r}$. Observe that when we restrict to $x_1 = \dots = x_r = 0$, we get that $\langle \lambda_t, f \rangle$ is the zero polynomial. So

$$\Pr_x [\langle \lambda_1, f(x) \rangle = \dots = \langle \lambda_T, f(x) \rangle = 0] \geq \Pr_x [x_1 = \dots = x_r = 0] = 2^{-r}.$$

We can equivalently write this probability as

$$\Pr_x [\langle \lambda_1, f(x) \rangle = \dots = \langle \lambda_T, f(x) \rangle = 0] = 2^{-T} \sum_{a \in \{0,1\}^T} \mathbb{E}_x \left[(-1)^{\langle \sum a_i \lambda_i, f(x) \rangle} \right].$$

The term $a = 0^T$ contributes $2^{-|T|}$ to the sum. So, since we have $T = r + 1$, there must exist some $\lambda = \sum a_i \lambda_i$ where a_i are not all zero (and hence $\lambda \neq 0$), such that

$$\mathbb{E} \left[(-1)^{\langle \lambda, f(x) \rangle} \right] \geq 2^{-(r+1)}.$$

This shows that the distribution of f is $2^{-(r+1)}$ biased. \square

4 Proving Small-bias

The goal in this section is to find sufficient conditions for the distribution of $f(x)$ to have small bias, for $m = n^s$. We showed in Lemma 3.1 that a necessary condition is that P has resilience $k \geq 2s$, and in Lemma 3.2 that any restriction of P obtained by fixing $r \leq s$ bits should have \mathbb{F}_2 -degree more than $s - r$. We show here that, up to constants, these bounds are tight. We assume below that n is large enough as a function of d (concretely, at least exponentially larger than d ; see the technical lemmas for more specific details).

We begin by showing that $f_{G,P}$ has low bias when the input-output dependency hypergraph G has sufficiently good expansion (as a function of the resiliency and bit fixing degree of the predicate).

Theorem 4.1. *Let P be a d -ary predicate which has resilience k , and has r -bit fixing degree e . Let $G = (I_1, \dots, I_m)$ be a d -uniform hypergraph which is $(t, d - \Delta)$ -expanding for $\Delta \leq \min((k + 1)/2, r/4, e/2)$. Then, for some $\delta = \delta(d) > 0$, the distribution of $f(x)$ is $\exp(-\delta t)$ -biased.*

The proof of Theorem 4.1 is deferred to Section 4.1. A standard calculation shows that for every $d \geq 3$, $\Delta > 2s$, and $m = n^s$, a random $G \leftarrow \mathcal{G}_{n,m,d}$ is likely to be an $(n^{1-\varepsilon}, d - \Delta)$ expander for $\varepsilon = 2(s - 1)/(\Delta - 2)$ (cf., [2, Lemma 4.1]). We therefore derive the following corollary by fixing $s \geq 1, \varepsilon > 0$ and setting the values of $\Delta, k, r, e = \Omega(s/\varepsilon)$ appropriately.

Corollary 4.2. *Fix $s \geq 1, \varepsilon > 0$. Let $m = n^s$. Let P be a predicate which has resilience k , and has r -bit fixing degree e where $k, r, e \geq \Omega(s/\varepsilon)$. Then, with probability at least $1 - o(1)$ over the choice of \mathcal{F} , the distribution of $f(x)$ is $\exp(-\delta n^{1-\varepsilon})$ -biased where $\delta = \delta(d) > 0$.*

4.1 Proof of Theorem 4.1

Following [38], we distinguish two types of linear tests: *light tests*, which have at most t nonzero coefficients, and *heavy tests*, which have more than t nonzero coordinates. We first rule out light tests. The following lemma essentially appears in [6] and is given here for completeness.

Lemma 4.3 (light tests). *Let $f = f_{G,P}$ where $G = (I_1, \dots, I_m)$ is $(t, d - (k + 1)/2)$ expanding and P is k -resilient. Then for any $\alpha \in \{0, 1\}^m$ of hamming weight $|\alpha| \leq t$ we have*

$$\mathbb{E}_x \left[(-1)^{\langle \alpha, f(x) \rangle} \right] = 0.$$

Proof. Let $A = \{i : \alpha_i = 1\}$. Since G is expanding, we have that $|\cup_{i \in A} I_i| > (d - \frac{k+1}{2})|A|$. This implies that for some $i \in A$ we have $|I_i \cap (\cup_{j \in A, j \neq i} I_j)| \leq k$. Otherwise, each set I_i contributes at most $d - (k + 1)$ unique elements to $\cup_{i \in A} I_i$, and so $|\cup_{i \in A} I_i| \leq |A|(d - (k + 1) + \frac{k+1}{2}) \leq |A|(d - \frac{k+1}{2})$, contradicting our expansion assumption.

Let $J = \cup_{j \in A, j \neq i} I_j$. Let $z = x|_{I_i^c}$ and fix all the other inputs $x|_{I_i^c} = v$. We will show that for any such fixing, $\langle \alpha, f(x) \rangle$ has zero bias (when averaging only over z). Indeed, averaging over the variables in z , we have that

$$\mathbb{E}_x \left[(-1)^{\langle \alpha, f(x) \rangle} | x|_{I_i^c} = v \right] = \mathbb{E}_z \left[(-1)^{P(z) + Q(z)} \right],$$

where $Q(z)$ depends only on the variables in $I_i \cap J$, which is at most k variables. Expanding P, Q into their Fourier decompositions, we obtain that

$$\mathbb{E}_z \left[(-1)^{P(z) + Q(z)} \right] = \sum_{\gamma, \gamma' \in \{0, 1\}^d} \hat{P}(\gamma) \hat{Q}(\gamma') \mathbb{E}_z \left[(-1)^{\langle \gamma + \gamma', z \rangle} \right] = \sum_{\gamma \in \{0, 1\}^d} \hat{P}(\gamma) \hat{Q}(\gamma).$$

Now, as Q depends on only k variables, we have that $\hat{Q}(\gamma) = 0$ whenever $|\gamma| > k$. By our assumption that P has resilience k , we have that $\hat{P}(\gamma) = 0$ whenever $|\gamma| \leq k$. This concludes the proof. \square

We handle heavy tests via the following lemma.

Lemma 4.4 (heavy tests). *Let $f = f_{G,P}$ where P has r -bit fixing degree e and $G = (I_1, \dots, I_m)$ is $(t, d - \Delta)$ -expanding for $\Delta \leq \min(r/4, e/2)$. Then for any $\alpha \in \{0, 1\}^m$ of hamming weight $|\alpha| \geq t$ we have*

$$\left| \mathbb{E}_x \left[(-1)^{\langle \alpha, f(x) \rangle} \right] \right| \leq \exp(-\delta t),$$

where $\delta = 1/(12d^24^d)$.

The proof strategy is, by fixing most variables, to obtain a collection of output variables with essentially disjoint inputs, and where all the other output variables depend on only a few inputs. We then use the properties of low degree polynomials to argue that the resulting bias is small. Concretely the proof employs the following algebraic claim and combinatorial claim whose proofs are deferred to Sections 4.2 and 4.3.

Claim 4.5. *Let $x = (x^1, \dots, x^\tau) \in \{0, 1\}^{n\tau}$ be τ disjoint sets of n binary variables. Define a polynomial*

$$F(x) = Q_1(x^1) + \dots + Q_\tau(x^\tau) + R(x^1, \dots, x^\tau),$$

where Q_1, \dots, Q_τ are nonzero polynomials of degrees $e \leq \deg_{\mathbb{F}_2}(Q_i) \leq d$ and $\deg_{\mathbb{F}_2}(R) < e$. Then

$$\left| \mathbb{E}_{x \in \{0, 1\}^{n\tau}} \left[(-1)^{F(x)} \right] \right| \leq \exp(-\tau/2^{d+e}).$$

Claim 4.6. *Assume that $G = (I_1, \dots, I_m)$ is $(t, d - \Delta)$ expanding. Then for every set of hyperedges $A \subset [m]$ of size $|A| \geq t$, there exists a subset of hyperedges $A^* \subseteq A$ of size $t/(12d^2)$ and a set of nodes $K \subset [n]$ such that*

- For each $i \in A^*$, the sets $K_i = I_i \cap K$ are disjoint and of size $|K_i| \geq d - 4\Delta$.
- For each $i \notin A^*$, we have $|I_i \cap K| < 2\Delta$.

Based on Claims 4.5 and 4.6, we analyze the bias of heavy tests.

Proof of Lemma 4.4. Let $A = \{i : \alpha_i = 1\}$ where $|A| \geq t$. Let A^* and K be the sets of hyperedges and nodes promised by Claim 4.6. Let $x \in \{0, 1\}^n$, where we consider some fixing $x|_{K^c} = v$ of the variables outside K . Let $y = x|_K$ denote the unfixed variables, and let $Q_i(y) = P(x|_{I_i})$ be the polynomial computed at the i -th bit, restricted to the unfixed variables y . Note that the polynomials $\{Q_i(y) : i \in A^*\}$ are evaluated on disjoint inputs, and by our assumptions, $\deg_{\mathbb{F}_2}(Q_i) \geq e \geq 2\Delta$ for all $i \in A^*$ (this is since we fixed at most $4\Delta \leq r$ variables in $P(x|_{I_i})$, and we assume that any such fixing attains a polynomial of degree at least e). On the other hand, for $i \notin A^*$ we have that $\deg_{\mathbb{F}_2}(Q_i) < 2\Delta$, since Q_i depends on less than 2Δ variables in K .

Let $F_v(y)$ be the restriction of $\langle \alpha, f(x) \rangle$ to $x|_K = y$, given the fixing $x|_{K^c} = v$. Then, $F_v(y) = \sum_{i \in A^*} Q_i(y) + R(y)$, where Q_i are polynomials on disjoint inputs of degrees $2\Delta \leq \deg_{\mathbb{F}_2}(Q_i) \leq \deg_{\mathbb{F}_2}(P) \leq d$, and where $\deg_{\mathbb{F}_2}(R) < 2\Delta$. By Claim 4.5, we obtain that

$$\left| \mathbb{E}_y \left[(-1)^{F_v(y)} \right] \right| \leq \exp(-|A^*|/2^{d+2\Delta}) \leq \exp(-t/12d^24^d),$$

where the last inequality follows by noting that $2\Delta \leq 2 \min(r/4, e/2) < d$. As this holds for any restriction v , we conclude that

$$\left| \mathbb{E}_x \left[(-1)^{\langle \alpha, f(x) \rangle} \right] \right| \leq \mathbb{E}_v \left| \mathbb{E}_y \left[(-1)^{F_v(y)} \right] \right| \leq \exp(-t/12d^24^d),$$

and the lemma follows. \square

4.2 Proof of Claim 4.5

Let $x = (x^1, \dots, x^\tau) \in \{0, 1\}^{n\tau}$, $F(x) = Q_1(x^1) + \dots + Q_\tau(x^\tau) + R(x^1, \dots, x^\tau)$ where Q_1, \dots, Q_τ are nonzero polynomials of degrees $e \leq \deg_{\mathbb{F}_2}(Q_i) \leq d$ and $\deg_{\mathbb{F}_2}(R) < e$. Our goal is to bound $\mathbb{E}_x[(-1)^{F(x)}]$. By the Gowers-Cauchy-Schwarz lemma [33], we can bound this by the e -th Gowers uniformity norm,

$$\left| \mathbb{E}_x \left[(-1)^{F(x)} \right] \right| \leq \|(-1)^F\|_{U^e} = \prod_{i=1}^{\tau} \|(-1)^{Q_i}\|_{U^e}.$$

Here, we used the basic fact that as $\deg_{\mathbb{F}_2}(R) < e$, R is annihilated in the computation of the e -th Gowers uniformity norm. Now, for any polynomial Q of degree $\deg_{\mathbb{F}_2}(Q) \geq e$, we have that $\|(-1)^Q\|_{U^e}^e$ is the bias of a nonzero polynomial of the same degree as Q . By the minimal distance property of polynomials, this bias is at most $1 - 2^{-\deg_{\mathbb{F}_2}(Q)}$. We obtain that

$$\left| \mathbb{E}_x \left[(-1)^{F(x)} \right] \right| \leq \left(1 - 2^{-d} \right)^{\tau/2^e} \leq \exp(-\tau/2^{d+e}).$$

The lemma follows. \square

4.3 Proof of Claim 4.6

Let $A \subset [m]$ be a set of at least t hyperedges. We gradually mark hyperedges and nodes in G . At the end, we let A^* and K be the set of unmarked hyperedges/nodes. We start by marking the hyperedges outside A and the nodes which are not adjacent to A , and proceed with three main steps.

In the first step we locate the set X of nodes with high degree in A and the set of hyperedges A_0 with high degree in X and mark both X and A_0 . Formally, define the set of frequently used nodes

$$X = \left\{ j \in [n] : j \in I_i \text{ for at least } \frac{2d}{t}|A| \text{ elements } i \in A \right\}.$$

As $|I_i| = d$ for all $i \in A$ we have $|X| \leq t/2$. Next, consider the set A_0 of hyperedges with large degree in X ,

$$A_0 = \{i \in A : |I_i \cap X| \geq \Delta - 1\}.$$

We claim that $|A_0| < |X| \leq t/2$. Indeed, if this is not the case then for any subset $S \subset A_0$ of size $|S| = |X| \leq t$, its neighborhood size $|\cup_{i \in S} I_i|$ is upper-bounded by $|X| + |S|(d - \Delta + 1) = |S|(d - \Delta)$, which contradicts the assumption on the expansion of G . Mark X and A_0 , and let $A_1 = A \setminus A_0$ be the set of remaining hyperedges. Note that $|A_1| \geq |A|/2$.

Our next step is to find a large set $A_2 \subset A_1$ of pairwise disjoint hyperedges, and mark all other hyperedges $A_1 \setminus A_2$ and all the nodes that do not participate in A_2 . Formally, for each $i \in A_1$ define $J_i = I_i \setminus X$. (This is the non-marked part of the hyperedges). We have that $|J_i| \geq d - \Delta$ for

all $i \in A_1$. Since the nodes in J_i belong to at most $(2d/t)|A|$ hyperedges in A , each J_i intersects at most $|J_i| \cdot (2d/t)|A|$ of the other sets $\{J_{i'} : i' \in A_1\}$. This implies that we can find $A_2 \subset A_1$, for which the sets $\{J_i : i \in A_2\}$ are pairwise disjoint, of size

$$|A_2| \geq \frac{|A_1|}{(2d^2/t)|A|} \geq \frac{t}{4d^2}.$$

Indeed, such an A_2 can be found via a simple greedy process in which in each step we choose an hyperedge J and remove all the hyperedges J' which intersect with J . Mark all the hyperedges outside A_2 and all the nodes that do not appear in the neighborhood of A_2 . The only non-marked nodes are now the nodes in $T = \cup_{i \in A_2} J_i$ (i.e., the non-marked neighborhood of A_2).

In our final step, we mark the nodes $C \subset T$ which participate in too many marked hyperedges, and keep only the hyperedges in A_2 which touch no more than 2Δ marked nodes. Formally, let $B = \{i \in A \setminus A_2 : |I_i \cap T| \geq 2\Delta\}$ be the set of marked hyperedges which touch at least 2Δ unmarked nodes, and let $C = \cup_{i \in B} (I_i \cap T)$ denote this set of nodes. We claim that

$$|C| \leq \sum_{i \in B} |I_i \cap T| \leq 2\Delta |A_2|.$$

To see this, assume that the claim does not hold. Then, there exists a set of hyperedges $B' \subseteq B$ of size $|B'| \leq |A_2|$ for which $\sum_{i \in B'} |I_i \cap T| \geq 2\Delta |A_2|$ (Such a B' exists since each hyperedge in B contributes at least 2Δ nodes to the sum). This means that the neighborhood of $A_2 \cup B'$ is of size at most $|\sum_{i \in A_2 \cup B'} I_i| \leq |A_2|d + |B'|d - 2\Delta |A_2| \leq (|A_2| + |B'|)(d - \Delta)$ where the last inequality follows from $|B'| \leq |A_2|$. Since $|A_2| + |B'| \leq 2|A_2| \leq t$ this violates expansion.

Mark the nodes in C . We will show that at least a third of the hyperedges $\{J_i : i \in A_2\}$ contain at most 3Δ marked nodes. First observe that the hyperedges $\{J_i : i \in A_2\}$ are disjoint and therefore their marked neighborhood (viewed as a multiset) can be upper-bounded by

$$\sum_{i \in A_2} |J_i \cap C| \leq |C| \leq 2\Delta |A_2|.$$

So on average, $|J_i \cap C| \leq 2\Delta$ for $i \in A_2$. By Markov's inequality, $|J_i \cap C| \leq 3\Delta$ for at least a third of $i \in A_2$. So, we can find $A^* \subset A_2$ of size $|A^*| = |A_2|/3$ such that each hyperedges in A^* touches at most 3Δ marked nodes, that is

$$|J_i \cap C| \leq 3\Delta \quad \forall i \in A^*.$$

Mark all the hyperedges outside A^* and all the nodes of the hyperedges outside A^* . For each of the unmarked hyperedges $i \in A^*$, let $K_i = J_i \setminus C$ denote the unmarked part of each of the unmarked hyperedges $i \in A^*$. Let $K = \cup_{i \in A^*} K_i$ be the set of unmarked nodes. Observe that A^* and K satisfy the required properties. Indeed,

- For each $i \in A^*$, the sets $K_i = I_i \cap K$ are disjoint and of size $|K_i| \geq d - 4\Delta$.
- For each $i \notin A^*$, we have $|I_i \cap K| < 2\Delta$.

The claim follows. □

5 Algebraic Attacks

5.1 Definitions

Algebraic attacks against a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ start with an output $y = f(x)$ and use it to initialize a system of polynomial equations over the hidden input variables $x = (x_1, \dots, x_n)$. Next, the system is *extended* by adding more equations that follow from the original one, and finally the system is *solved*. There are different strategies for implementing the extension step (e.g., by multiplying the polynomials by some low-degree polynomial) and the solution step (e.g., via linearization and Gaussian elimination or by computing a Gröbner basis of the expanded system). We abstract these strategies via the following general definition of algebraic attack which is parameterized with some *scheduling algorithm*.

Definition 4 (Algebraic inversion attack). *An algebraic inversion attack against a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is parameterized with a stateful scheduling algorithm S . Given an input $y \in \{0, 1\}^m$ the attack has the following form:*

1. (Axiom Loading) Initialize a list L of polynomial equations that contains all output equations

$$f_i(x) - y_i = 0, \quad \forall i \in [m],$$

where f_i denotes the Boolean function that computes the i -th output bit of f represented as an \mathbb{F}_2 polynomial in the input variables $x = (x_1, \dots, x_n)$.

2. (Extension) Based on y, f, L and its internal state, the scheduler S either decides to terminate with failure, or does one of the following atomic operations:

- (a) Chooses an equation $Q(x) = 0$ from L and a variable x_i and adds to L the equation

$$x_i \cdot Q(x) = 0.$$

- (b) Chooses a pair of equations $Q(x) = 0$ and $R(x) = 0$ from L and add to L the equation

$$Q(x) + R(x) = 0.$$

3. (Inversion) If, for some $b \in \mathbb{F}_2^n$, the list L contains the equations

$$x_i - b_i = 0 \quad \forall i \in [n],$$

terminate with the assignment b .

4. Go to Step 2.

The above attack attempts to recover x by deducing the equations $x_i - b_i = 0$ from the initial equations (axioms) using simple derivation rules (atomic operations). These derivation rules correspond to the well known Polynomial Calculus (PC) proof system which was introduced by Clegg, Edmonds and Impagliazzo [15].¹⁰

¹⁰Over general (non-binary) fields, the second derivation rule (addition) is generalized to $(Q, R) \models \alpha Q(x) + \beta R(x) = 0$ for any field elements α, β . Also, in order to force a binary solution the axioms $x_i^2 - x_i = 0$ are added for each input variable x_i .

While the PC derivation rules are seemingly elementary, they can be used to efficiently implement a rich family of algorithms including Gaussian elimination, and algorithms for polynomial division, polynomial reduction, polynomial greatest common divisors, and for computing Gröbner basis. As a result, standard algebraic attacks from the cryptanalysis literature (e.g., linearization, XL [22], F4 [26] and F5 [27]) fall into the above framework.

Refutation attacks. While typical algebraic attacks attempt to invert f on $y \in \text{Im}(f)$, one can also consider an *algebraic distinguishing attack* which attempts to break the pseudorandomness of f by distinguishing between $y \in \text{Im}(f)$ and a random $y \leftarrow \{0, 1\}^m$. This is done by trying to recognize the event that y falls out of the image of f , i.e., that the system $f(x) = y$ has no solution. We formalize this as follows.

Definition 5 (Algebraic refutation attack). *An algebraic refutation attack against a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is defined similarly to algebraic inversion attack except that the Inversion step (Step 3) is replaced with the following step:*

(Refutation) If L contains the equation $1 = 0$ (contradiction), terminate with the output “Unsatisfied”.

A refutation attack gives rise to a distinguisher with one-sided error: It outputs “Unsatisfied” only if it certified that $y \notin \text{Im}(f)$. Furthermore, the certificate is given as a proof in polynomial calculus for the unsatisfiability of the $f(x) = y$. We measure the success probability of an algebraic refutation attack \mathcal{A} as the probability that \mathcal{A} outputs “Unsatisfied” (encoded as 0) over a random $y \leftarrow \{0, 1\}^m$. Since \mathcal{A} has one-sided error, the success probability p equals to the standard distinguishing advantage of \mathcal{A} :

$$\left| \Pr_{y \leftarrow \{0, 1\}^m} [\mathcal{A}(y) = 0] - \Pr_{x \leftarrow \{0, 1\}^n} [\mathcal{A}(f(x)) = 0] \right|.$$

5.2 Algebraic Attacks against Predicates with Low Rational Degree

Theorem 5.1. *Assume that $P : \{0, 1\}^d \rightarrow \{0, 1\}$ is a predicate of rational degree e . Then there exists an efficient algebraic refutation algorithm \mathcal{A} such that for every $m \geq \Omega_d(n^e)$ it holds that*

$$\Pr[\mathcal{A}(f, y) \text{ derives a contradiction}] > 1 - \exp(-\Omega(n^e)),$$

where $f \leftarrow \mathcal{F}_{P, n, m}$ and $y \leftarrow \{0, 1\}^m$.

Before describing our attack we will need to make few observations. First, since P has a rational degree e , we are guaranteed to have a non-zero degree e polynomial $Q(z_1, \dots, z_d)$ for which one of the following holds: $P = 0 \Rightarrow Q = 0$ or $P = 1 \Rightarrow Q = 0$. We let β be zero in the former case, and one in the latter case and note that $(P - \beta)Q = Q$.

Second, observe that, due to the completeness of the polynomial calculus proof system [15], there exists an algebraic refutation algorithm that, given an unsatisfiable set of polynomial equations over a set of variables S , generates the equation $1 = 0$ in time which depends only on the size of S . In particular, if S has constant size then the algorithm runs in constant time.¹¹ We refer to this refutation algorithm as the *enumerate* algorithm.

¹¹Indeed, [31, Lemma 5.3] shows that any (multilinear) polynomial equation which logically follows from H can be derived using $\exp(|S|)$ atomic steps. By trying all possible such sequences we get the desired algorithm.

Finally, we will say that a polynomial equation $T(x) = 0$ is spanned by a set of polynomial equations $\mathcal{T} = \{T_i(x) = 0 : 1 \leq i \leq t\}$ if the former equation can be written as a linear combination of the equations in \mathcal{T} . Note that this happens if and only if the polynomial T is spanned by the polynomials T_1, \dots, T_t . Therefore, using Gaussian elimination, we can efficiently check if the equation $T(x) = 0$ is spanned by \mathcal{T} , and, in case it does, we can efficiently find a sequence of at most t atomic steps that derives the equation $T(x) = 0$ from \mathcal{T} .

We can now describe our algebraic refutation algorithm:

1. For every $i \in [m]$ for which $y_i = \beta$ derive the equation $Q(x_{I_i}) = 0$ and put it in the set $\mathcal{L}(G, y)$. Each of the above equations can be derived by multiplying the axiom $P(x_{I_i}) - \beta = 0$ by the degree e polynomial $Q(x_{I_i})$. Therefore the cost per equation is $O(e \cdot d^e) = O(1)$ atomic steps.
2. For every d -tuple $I \in [n]_d$, if the polynomial equation $Q(x_I) = 0$ is spanned by $\mathcal{L}(G, y)$, then derive it from $\mathcal{L}(G, y)$ using at most $|\mathcal{L}(G, y)| \leq m$ atomic steps.
3. Check if there exists a set $A \subset [m]$ of size $2d$ for which the following holds:
 - i The set $S = \bigcup_{i \in A} I_i$ is of size $2d^2$; Namely, the collection of sets $\{I_i : i \in A\}$ is pair-wise disjoint.
 - ii For every d -tuple I with d distinct elements from S (i.e., $I \in S_d$) the polynomial equation $Q(x_I) = 0$ was generated in the previous step.
 - iii The string y_A is balanced, i.e., it contains d ones and d zeroes.
4. If there exists a set A which passes the test, derive a contradiction by applying the “enumerate” algorithm to the polynomial equations $\{Q(x_I) = 0 : I \subset S\} \cup \{P(x_{I_i}) - y_i = 0 : i \in A\}$ which are defined over the constant-size set of variables S . (We will later see that this set of equations is indeed unsatisfiable.)
5. Otherwise, if the test in Step 3 fails, abort with failure.

It is not hard to verify that the algorithm runs in time $\text{poly}(n, m)$. (The most expensive step is Step 3 which can be implemented by trying all $O(m^{2d})$ subsets.) We first prove that Step 4 indeed derives a contradiction.

Claim 5.2. *If A satisfies conditions (i)–(iii) then the set of equations $\{Q(x_I) = 0 : I \subset S\} \cup \{P(x_{I_i}) - y_i = 0 : i \in A\}$ is unsatisfiable.*

Proof. Fix some $z \in \{0, 1\}^d$ for which $Q(z) = 1$. (Such a string z exists since Q is not constant.) We first claim that any assignment for x_S which satisfies the system $\{Q(x_I) = 0 : I \subset S\}$ either: (a) contains less than d ones; or (b) contains less than d zeroes. Indeed, if an assignment ρ to x_S contains both d ones and d zeroes then there exists a tuple $I \subset S$ for which $\rho_I = z$ which means that the equation $Q(x_I) = 0$ is not satisfied.

Next, we show that any assignment for x_S which satisfies $H = \{P(x_{I_i}) - y_i = 0 : i \in A\}$ cannot contain less than d ones. Fix some assignment ρ to x_S which contains less than d ones. Then, for all but $d - 1$ of $i \in A$, it holds that $\rho_{I_i} = 0^d$. Therefore, for at least $d + 1$ of the elements $i \in A$ it holds that $P(\rho_{I_i}) - b = 0$ where $b := P(0^d)$. Since y_A is balanced, this means that there exists $i \in A$ for which

$$P(\rho_{I_i}) - b = 0 \quad \text{but} \quad (P(x_{I_i}) - (1 - b) = 0) \in H,$$

and ρ does not satisfy H . A similar argument shows that assignments for x_S which contain less than d zeroes cannot satisfy H . Overall, we conclude that the system is unsatisfiable. \square

Our next goal is to show that a random (G, y) is likely to pass the test in Step 3. Before that we will need the following notation and claim. For a hypergraph $G = (I_1, \dots, I_m)$ and a string y we let $p(G, y)$ denote

$$\Pr_{I \leftarrow [n]_d} [Q(x_I) \in \text{span}(\mathcal{L}(G, y))],$$

where $\mathcal{L}(G, y)$ is the set of polynomials $\{Q(x_{I_j}) : y_j = \beta\}$.

Claim 5.3. *For every constant $\varepsilon \in [0, 1]$, constant $c > 2/\varepsilon$, and $m \geq cn^e$, it holds that*

$$\Pr_{G \leftarrow \mathcal{G}_{n,m,d}, y \leftarrow \{0,1\}^m} [p(G, y) \geq 1 - \varepsilon] > 1 - \exp(-\Omega(n^e)).$$

Proof. We upper-bound the probability that $p = p(G, y)$ is smaller than $1 - \varepsilon$ as follows. Consider the experiment in which (G, y) is gradually sampled in m steps where at the i -th step the pair (I_i, y_i) is chosen uniformly at random, i.e., $I_i \leftarrow [n]_d$ and $y_i \leftarrow \{0, 1\}$. Let χ_i denote the indicator random variable which takes the value 1 if the rank of $\mathcal{L}(G_{[1:i]}, y_{[1:i]})$ is larger than the rank of $\mathcal{L}(G_{[1:i-1]}, y_{[1:i-1]})$, where $G_{[1:i]} = (I_1, \dots, I_i)$ and $y_{[1:i]} = (y_1, \dots, y_i)$. Then, conditioned on $p(G_{[1:i-1]}, y_{[1:i-1]}) < 1 - \varepsilon$, it holds that

$$\Pr[\chi_i = 1] = \Pr_{I_i} [Q(x_{I_i}) \notin \text{span}(\mathcal{L}(G_{[1:i-1]}, y_{1:i-1}))] \cdot \Pr[y_i = \beta] \geq \varepsilon/2.$$

Also, observe that if $p < 1 - \varepsilon$ then $\text{rank}(\mathcal{L}(G, y)) < n^e$, since \mathcal{L} contains only degree e polynomials, and there exists a degree e polynomial which is not spanned by it. We conclude that

$$\Pr_{G,y} [p(G, y) < 1 - \varepsilon] \leq \Pr[\forall i, \chi_i \text{ takes the value 1 w/p at least } \varepsilon/2 \text{ and } \sum_{i=1}^m \chi_i < n^e].$$

Since $m \geq cn^e$ for some constant $c > 2/\varepsilon$, we can use a Chernoff bound to upper-bound the probability of the latter event by $\exp(-\Omega(m))$. \square

We complete the proof of Theorem 5.1 by proving that a random (G, y) is likely to pass the test in Step 3.

Lemma 5.4. *There exists a constant $c = c(d) > 0$ such that for every $m \geq cn^e$ it holds*

$$\Pr_{G \leftarrow \mathcal{G}_{n,m,d}, y \leftarrow \{0,1\}^m} [(G, y) \text{ passes the test in Step 3}] > 1 - \exp(-\Omega(n^e)).$$

Proof. Let $D = \binom{2d^2}{d} d!$ denote the number of all possible d -tuples with distinct elements taken from a universe S of size $2d^2$. Let $\varepsilon = 1/(2D)$ and assume that $m \geq (4D + 2)n^e$. (We did not attempt to optimize the constants.) We partition $G = (I_1, \dots, I_m)$ and $y = (y_1, \dots, y_m)$ to two parts: a left part

$$G_L = (I_1, \dots, I_{m_L}) \quad \text{and} \quad y_L = (y_1, \dots, y_{m_L}),$$

where $m_L = (4D + 1)n^e$, and a right part

$$G_R = (I_{m_L+1}, \dots, I_m) \quad \text{and} \quad y_R = (y_{m_L+1}, \dots, y_m).$$

From now on, we condition on the event that $p(G_L, y_L) \geq 1 - \varepsilon$ which, by Claim 5.3, happens with probability $1 - \exp(-\Omega(n^\varepsilon))$. We say that a d -tuple I is *good* if the polynomial $Q(x_I)$ is spanned by $\mathcal{L}(G_L, y_L)$. For a $2d^2$ -size set $S \subseteq [n]$, let $T(S)$ denote the number of good d -tuples I with distinct elements from S . We call S *good* if $T(S) = D$, that is, S is good if every d -tuple $I \in S_d$ is good. Let α denote the fraction of good S 's among all $2d^2$ -size subsets of $[n]$. Then,

$$(1 - \varepsilon)D \leq p(G_L, y_L)D = \mathbb{E}_S[T(S)] \leq \alpha D + (1 - \alpha)(D - 1).$$

To see the equality, recall that $p(G_L, y_L)$ measures the fraction of good d -tuples and note that choosing a random d -tuple $I \leftarrow [n]_d$ is equivalent to choosing a random $2d$ -subset $S \subset [n]$ and then selecting a random d -tuple $I \in S_d$. The above equation implies that $\alpha \geq 1 - \varepsilon D = 1/2$, and so at least half of all the S 's are good.

Let us now partition (G_R, y_R) to $k = \lfloor n^\varepsilon/2d^2 \rfloor$ blocks of size $\ell = 2d^2$ each. That is, for $i \in [k]$ define $I^{(i)} = (I_{m_L+(i-1)\ell+1}, \dots, I_{m_L+i\ell})$ and $y^{(i)} = (y_{m_L+(i-1)\ell+1}, \dots, y_{m_L+i\ell})$. Call a block i *good* if:

1. the tuples in $I^{(i)}$ are pair-wise disjoint; and
2. their union $S = \bigcup_{I \in I^{(i)}} I$ is good; and
3. the corresponding string $y^{(i)}$ is balanced.

Observe that $I^{(i)}$ satisfies (1) with probability $1 - o(1)$ and, conditioned on (1), it satisfies (2) with probability $\alpha \geq 1/2$. Also, $y^{(i)}$ satisfies (3) with probability $\Omega(1/\sqrt{|y^{(i)}|}) = \Omega(1/d)$. Overall, each block is good with constant probability. Since the distribution of each block is statistically independent, we conclude that there exists a good block with probability $1 - \exp(-\Omega(k))$. The lemma follows. \square

5.3 High Rational Degree Resists Algebraic Attacks

We will prove lower-bounds against algebraic attacks. Our lower-bounds apply to the total number of monomials which are stored in L which provide a lower-bound on the total running time. Notably we ignore the computational complexity of the scheduling algorithm and so our results apply to the “best algebraic algorithm” which, for every y , chooses the best (cheapest) sequence of extension steps.

We begin with lower-bounds against algebraic refutation attacks.

Theorem 5.5. *Let $m = an^s$ for some arbitrary constants $a, s \geq 1$. Let d be a constant and let $P : \{0, 1\}^d \rightarrow \{0, 1\}$ be a predicate with rational degree of $\ell + 1$ where $\ell > 8s$. Then, with probability at least $1 - o(1)$ over the choice of $f \leftarrow \mathcal{F}_{P, n, m}$, the following holds. For every $y \notin \text{Im}(f)$ the complexity of any algebraic refutation attack against f on an input y is $\exp(\Omega(n^{1-\delta}))$ where $\delta = \frac{16(s-1)}{\ell}$.*

Note that when $m = O(n)$ we derive an exponential lower bound of $\exp(\Omega(n))$.

Proof. Recall that the input-output dependency hypergraph $G = (I_1, \dots, I_m)$ is (r, c) expanding if for every set $A \subseteq [m]$ of cardinality at most r , it holds that $|\bigcup_{i \in A} I_i| > c|A|$. A standard calculation shows that for every $d \geq 3$, $c < d - 2$, and $m/n = o(n^{(d-c-2)/2})$ a random $G \leftarrow \mathcal{G}_{n, m, d}$ is likely to be an (r, c) expander for some $r = \Omega(\frac{n}{(m/n)^{2/(d-c-2)}})$ (cf., [2, Lemma 4.1]). In our case $m = an^s$ and

so we can take c to be some real number $c \in (d - \ell/4, d - 2s)$ and $r = \Omega(n^{1-\varepsilon})$ where $\varepsilon = \frac{2(s-1)}{d-c-2}$. (The reason for the lower-bound on c will be clear later. For now observe that, by the condition $\ell > 8s$, we can choose such a constant c .)

Fix some G that satisfies (r, c) -expansion and some y outside the image of $f = f_{G,P}$. Any algebraic refutation algorithm that terminates on y with the output “unsatisfiable” provides a Polynomial Calculus proof that the system of polynomial equations L constructed in the initialization step is unsatisfiable. The *degree* of such a proof is defined to be the maximal degree of a polynomial Q that appears in the proof.

Alekhovich and Razborov [2, Theorem 3.8] prove that any PC refutation of an unsatisfiable system $L = \{Q_i(x) = 0\}$ must contain a polynomial of degree larger than $D = r(\ell/4 - (d - c))$, provided that the following conditions hold for some r, ℓ, d, s :

1. The underlying dependency hypergraph of the system $L = \{Q_i(x) = 0\}$ is (r, c) expanding;
2. Each polynomial Q_i depends on at most $d < D$ inputs; and
3. For every i , the polynomial Q_i is ℓ -immune in the sense that there is no degree- ℓ non-zero polynomial Q' which satisfies $Q_i(x) = 0 \Rightarrow Q'(x) = 0$.

Observe that all three conditions hold in our case. Indeed, (1) holds by definition. Condition (2) follows by noting that d is constant, while D is super-constant since r is super-constant and $\ell/4 - (d - c)$ is positive (as implied by $c > d - \ell/4$). Finally, to see that the last condition holds, recall that in our system the initial polynomials are $P(x) - y_i$ and therefore, since P has rational degree of $\ell + 1$, all these polynomials are ℓ -immune in the Alekhovich-Razborov sense.

It follows that our unsatisfiable system has no PC proof of degree $D = \Omega(n^{1-\varepsilon})$. Impagliazzo et al. [34, Corollary 6.3] show that any such degree lower bound translates into a size-lower bound of $\exp(\Omega(D^2/n))$, provided that the initial system consists of constant-degree polynomials (which is indeed the case in our situation). This implies that the running time of an algebraic attack on y is lower-bounded by $\exp(\Omega(n^{1-2\varepsilon}))$. Recall that s can be taken to be an arbitrary constant in the interval $(d - \ell/4, d - 2s)$ and that $\varepsilon = \frac{2(s-1)}{d-c-2}$. We can therefore choose s to be $c = d - \ell/4 + \alpha$ for some (possibly tiny) $\alpha < 2$ and derive a lower-bound of $\exp(\Omega(n^{1-\frac{4(s-1)}{\ell/4}}))$ as required. \square

The above lower-bound applies to y which are outside the image of f and so they measure the *algebraic refutation complexity*. We show that when f is “simple” enough, the lower-bound extends to algebraic *inversion* attacks which are applied on $y \in \text{Im}(f)$.

For the following lemma, we measure the *simplicity* of a boolean function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ via the size of the minimal *skew circuit* that computes it, where an arithmetic circuit (over the binary field) is skew if each of its multiplication gates involves at least one argument that is an input variable. It is known that any language computable in log-space (or even non-deterministic log-space) has a polynomial-size skew-circuit [43].

Lemma 5.6. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function and let $f' : \{0, 1\}^n \rightarrow \{0, 1\}^{m-1}$ be the function obtained by omitting from f its last output bit. Suppose that there exists an algebraic inversion attack against f' that has complexity t on some input $y' \in \text{Im}(f')$. Then, there exists an algebraic refutation attack against f that, on some input $y \notin \text{Im}(f)$, has complexity of $t + O(s)$ where s is the size of the skew-circuit of f_m , the last output of f .*

Proof. Let $y' \in \text{Im}(f')$ be the $(m-1)$ -bit string on which f' has t -time algebraic inversion attack, and let $b \in \{0, 1\}^n$ be the output of the attack. Consider the m -bit string $y = (y', 1 + f_m(b))$; namely, y is obtained by concatenating the complement of the last bit of $f(b)$ to the end of the string y' . We will describe an algebraic refutation attack against f that, given y , outputs “unsatisfiable” in time $t + O(s)$. (Since the polynomial-calculus is sound, this also shows that $y \notin \text{Im}(f)$.)

The attack consists of three high-level steps. First, apply the f' -attack that inverts y' and derive the equations $x_i - b_i$ for all $i \in [n]$. Second, use these equations to derive the polynomial equation $f_m(x) - f_m(b) = 0$. Finally, subtract this equation from the last-output equation $f_m(x) - y_m = 0$ and derive the contradiction $1 = 0$.

The second step is implemented in time $O(s)$ by traversing the skew circuit C that computes f_m from the inputs to the outputs while generating, for every internal gate $g(x)$, a polynomial equation of the form $g(x) - g(b) = 0$. For the input gates such equations are already presented in L . For an addition gate $g(x) = g_1(x) + g_2(x)$, we recursively generate the equations $g_1(x) - g_1(b) = 0$ and $g_2(x) - g_2(b) = 0$ and add them together. For a multiplication gate $g(x) = x_i \cdot h(x)$, we recursively generate the equation $h(x) - h(b) = 0$, multiply the latter by the variable x_i to obtain the equation $g(x) - x_i \cdot h(b) = 0$, and, if the constant $h(b)$ is non-zero, add the latter equation to the equation $x_i - b_i = 0$. It is not hard to verify that this yields an equation of the form $g(x) - g(b) = 0$. The lemma follows. \square

By combining Lemma 5.6 with Theorem 5.5 we derive the following corollary.

Corollary 5.7. *Let $m = an^s$ for some arbitrary constants $a, s \geq 1$. Let d be a constant and let $P : \{0, 1\}^d \rightarrow \{0, 1\}$ be a predicate with rational degree of $\ell + 1$ where $\ell > 8s$. Then, with probability at least $1 - o(1)$ over the choice of $f \leftarrow \mathcal{F}_{P,n,m}$, the complexity of any algebraic inversion attack against f on any input $y \in \text{Im}(f)$ is $\exp(\Omega(n^{1-\delta}))$ where $\delta = \frac{16(s-1)}{\ell}$.*

Proof. Call a function t -invertible (resp., t -refutable) if there exists a string in its image (resp., outside the image) which can be inverted (resp., refuted) by some algebraic attack in time t . Recall that, by Theorem 5.5, with probability of $1 - o(1)$, a random function $f \leftarrow \mathcal{F}_{P,n,m+1}$ cannot be algebraically refuted in less than $t = \exp(\Omega(n^{1-\delta}))$ time. By Lemma 5.6, it follows that $1 - o(1)$ fraction of $f' \leftarrow \mathcal{F}_{P,n,m}$ cannot be algebraically inverted in time less than $t' = t - O(2^d) = t - O(1)$. To see this observe that (1) any d -local function has a 2^d skew circuit, and (2) the restriction of a function $f \leftarrow \mathcal{F}_{P,n,m+1}$ to its first m output bits is distributed according to $\mathcal{F}_{P,n,m}$. The corollary follows. \square

6 Refuting the FPV conjecture

Feldman, Perkins and Vempala [30] presented a general model of Constraint Satisfaction Problems with planted assignments, proved strong upper-bounds and lower-bounds against a large class of statistical algorithms, and made a conjecture about the intractability of such planted-CSP's. We will show that their conjecture does not hold. For this it suffices to consider a simpler notion of planted CSPs which is a special case of the model of [30].

The model. Let $P : \{0, 1\}^d \rightarrow \{0, 1\}$ be a d -ary predicate. A P -CSP instance φ over n variables and m constraints is a list of m d -clauses C_1, \dots, C_m where each clause is represented by a d -tuple $I \in [n]_d$ of distinct indices and by a string $z \in \{0, 1\}^d$ that represents the *polarity* of each

variable. (For example, $I = (1, 3, 6)$ and $z = (0, 1, 0)$ represents the clause (x_1, \bar{x}_3, x_6) .) A clause (I, z) is satisfied by an assignment $x \in \{0, 1\}^n$ if $P(x|_I \oplus z) = 1$. For a planted assignment $x \in \{0, 1\}^n$, we let $\Phi_{P,n,m}(x)$ denote the distribution over $\varphi = (C_1, \dots, C_m)$ in which each clause C_i is chosen independently and uniformly at random among all clauses which satisfy x . Equivalently, each clause is sampled by first choosing a random tuple $I_i \in [n]_d$ and then choosing a random $z_i \in \{0, 1\}^d$ subject to $P(x|_I \oplus z) = 1$. We say that an algorithm \mathcal{A} solves $\Phi(P, n, m)$ if

$$\Pr_{x \leftarrow \{0,1\}^n, \varphi \leftarrow \Phi_{P,n,m}(x)} [\mathcal{A}(\varphi) = x] > 1/3.$$

The following conjecture is implied by Conjecture 1 of [30]:

Conjecture 1. *If the predicate P is $(r - 1)$ -resilient and its algebraic degree is larger than $r/2$ then, for any $c < r/2$, there is no efficient algorithm that solves $\Phi(P, n, n^c)$.*

We refute the above conjecture.

Theorem 6.1. *For every constant r there exists an $(r - 1)$ -resilient predicate P with algebraic degree larger than $r/2$ for which $\Phi(P, n, m = \omega(n^3))$ can be solved efficiently.*¹²

Proof. For given r , let $P = \text{XOR-AND}_{r,r}$ be the $2r$ -ary predicate that computes the XOR of the first r inputs and the AND of the last r inputs and outputs the XOR of the results. We assume, without loss of generality, that r is odd (though one can tweak the algorithm to work with an even r as well). In the following we will say that a variable i participates as a positive (resp., negative) XOR variable in a clause $(I = (i_1, \dots, i_{2r}), z)$ if for some $1 \leq j \leq r$ it holds that $i = i_j$ and $z_j = 0$ (resp., $z_j = 1$). Similarly, we say that i participates as a positive (resp., negative) AND variable in (I, z) if for some $r + 1 \leq j \leq 2r$ it holds that $i = i_j$ and $z_j = 0$ (resp., $z_j = 1$).

Let $m = \omega(n^3)$. Our algorithm takes a P -CSP instance $\varphi = (C_1, \dots, C_m)$ and does the following. First, we let A_0 (resp., A_1) denote the set of clauses in which the last variable participates as a positive AND variable (resp., negative AND variable). Next, we derive two candidate assignments x^0 and x^1 as follows. Set the last entry of x^b to the value b (i.e., we guess that the last variable of the planted assignment x^* is set to b) and assign the other values of x^b such that the following system of linear equations is satisfied. For every $j \in A_b$, generate the linear equation (over the formal variables (x_1, \dots, x_{n-1})) induced by the j -th constraint under the assumption that the AND part is set to b . That is, if the j -th clause C_j contains the variables i_1, \dots, i_r as XOR variables with polarities z_1, \dots, z_r then we get the constraint: $(x_{i_1} \oplus z_1) \oplus \dots \oplus (x_{i_r} \oplus z_r) = 1$ or equivalently

$$x_{i_1} \oplus \dots \oplus x_{i_r} = (1 \oplus z_{i_1} \dots \oplus z_{i_r}). \quad (1)$$

Finally, output x^0 if it satisfies all the original constraints, and otherwise output x^1 .

We claim that for any planted assignment x^* , when given a random instance $\varphi \leftarrow \Phi_{P,n,m}(x^*)$, the above algorithm recovers x^* with probability at least $2/3$. In fact, we will show that if x^* sets the last variable to the value b then, with probability $1 - o(1)$, the assignment x^b satisfies all the constraints. This suffices to prove the claim since, by a standard probabilistic argument, whenever $m = \omega(n)$ the only satisfying assignment is the planted one (except with probability $o(1)$).

Fix some x^* and let us assume without loss of generality that $x_n^* = 0$. The main observation is that for each constraint C_i the variable x_n participates as a positive AND variable with probability $\Omega(1/n)$, and, conditioned on being selected, the r -tuple (i_1, \dots, i_r) of the XOR variables is

¹²We did not attempt to optimize m , and we believe that a more careful analysis can reduce it to $\Omega(n^2/\log n)$.

distributed uniformly over $[n-1]_r$ and the corresponding polarities (z_1, \dots, z_r) are uniform subject to Eq. (1). By a Chernoff bound, it follows that A_0 contains, with probability $1 - o(1)$, at least $\omega(n^2)$ clauses, and, conditioned on this, the system that we get contains $\omega(n^2)$ linear equations where each equation contains r random inputs from $[n-1]$ and the RHS is consistent with x^* . To conclude the proof, we will show (via a standard probabilistic argument) that such a system is likely to have a single solution.

First, observe that for any fixed $x \neq x^*$, the probability that x is consistent with x^* with respect to a random r -ary linear constraint is at most $1 - \Omega(1/n)$. Indeed, if the set of indices S on which x and x^* disagree is smaller than $n/2$ then with probability $\Omega(1/n)$ a random constraint touches a single location in S . On the other hand, if S is larger than $n/2$ then with constant probability of 2^{-r} all the r entries of the constraint fall in S and since r is odd this means that x^* is not consistent with the constraint. Overall, x is a valid solution with probability at most $(1 - \Omega(1/n))^{\omega(n^2)} < \exp(\omega(-n))$ and by applying a union-bound over all possible n -bit strings $x \neq x^*$, we conclude that x^* is likely to be the single satisfying assignment. \square

References

- [1] M. Alekhnovich. More on average case vs approximation complexity. In *44th FOCS*, pages 298–307, Cambridge, Massachusetts, USA, Oct. 11–14, 2003. IEEE Computer Society Press.
- [2] M. Alekhnovich and A. A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. In *42nd FOCS*, pages 190–199, Las Vegas, Nevada, USA, Oct. 14–17, 2001. IEEE Computer Society Press.
- [3] B. Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. *SIAM J. Comput.*, 42(5):2008–2037, 2013.
- [4] B. Applebaum. Cryptographic hardness of random local functions - survey. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:27, 2015.
- [5] B. Applebaum, B. Barak, and A. Wigderson. Public-key cryptography from different assumptions. In L. J. Schulman, editor, *42nd ACM STOC*, pages 171–180, Cambridge, Massachusetts, USA, June 5–8, 2010. ACM Press.
- [6] B. Applebaum, A. Bogdanov, and A. Rosen. A dichotomy for local small-bias generators. In R. Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 600–617, Taormina, Sicily, Italy, Mar. 19–21, 2012. Springer, Heidelberg, Germany.
- [7] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in NC^0 . *SIAM J. Comput.*, 36(4):845–888, 2006. Preliminary version in FOCS 2004.
- [8] B. Applebaum, Y. Ishai, and E. Kushilevitz. On pseudorandom generators with linear stretch in NC^0 . *J. of Computational Complexity*, 17(1):38–69, 2008.
- [9] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, May 1998. Prelim version FOCS '92.
- [10] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, Jan. 1998. Prelim version FOCS '92.

- [11] B. Barak, G. Kindler, and D. Steurer. On the optimality of semidefinite relaxations for average-case and generalized constraint satisfaction. In R. D. Kleinberg, editor, *ITCS 2013*, pages 197–214, Berkeley, CA, USA, Jan. 9–12, 2013. ACM.
- [12] A. Bogdanov and Y. Qiao. On the security of Goldreich’s one-way function. *Computational Complexity*, 21(1):83–127, 2012.
- [13] A. Bogdanov and A. Rosen. Input locality and hardness amplification. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 1–18, Providence, RI, USA, Mar. 28–30, 2011. Springer, Heidelberg, Germany.
- [14] C. Carlet. *Boolean models and methods in mathematics, computer science, and engineering*, chapter Boolean functions for cryptography and error-correcting codes, pages 257–397. Cambridge University Press, 2010.
- [15] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *28th ACM STOC*, pages 174–183, Philadelphia, Pennsylvania, USA, May 22–24, 1996. ACM Press.
- [16] J. Cook, O. Etesami, R. Miller, and L. Trevisan. Goldreich’s one-way function candidate and myopic backtracking algorithms. In O. Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 521–538. Springer, Heidelberg, Germany, Mar. 15–17, 2009.
- [17] J. Cook, O. Etesami, R. Miller, and L. Trevisan. On the one-way function candidate proposed by goldreich. *ACM Transactions on Computation Theory*, 6(3):14:1–14:35, 2014.
- [18] S. Cook. The complexity of theorem proving procedures. pages 151–158. ACM Press, 1971.
- [19] N. Courtois. The security of hidden field equations (HFE). In D. Naccache, editor, *CT-RSA 2001*, volume 2020 of *LNCS*, pages 266–281, San Francisco, CA, USA, Apr. 8–12, 2001. Springer, Heidelberg, Germany.
- [20] N. Courtois. Fast algebraic attacks on stream ciphers with linear feedback. In D. Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 176–194, Santa Barbara, CA, USA, Aug. 17–21, 2003. Springer, Heidelberg, Germany.
- [21] N. Courtois. General principles of algebraic attacks and new design criteria for cipher components. In H. Dobbertin, V. Rijmen, and A. Sowa, editors, *Advanced Encryption Standard AES*, volume 3373 of *Lecture Notes in Computer Science*, pages 67–83. Springer Berlin Heidelberg, 2005.
- [22] N. Courtois, A. Klimov, J. Patarin, and A. Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In B. Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 392–407, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.
- [23] N. Courtois and W. Meier. Algebraic attacks on stream ciphers with linear feedback. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 345–359, Warsaw, Poland, May 4–8, 2003. Springer, Heidelberg, Germany.

- [24] M. Cryan and P. B. Miltersen. On pseudorandom generators in NC. In *Proc. of 26th Mathematical Foundations of Computer Science (MFCS)*, pages 272–284, 2001.
- [25] I. Dinur, S. Goldwasser, and H. Lin. The computational benefit of correlated instances. In T. Roughgarden, editor, *ITCS 2015*, pages 219–228, Rehovot, Israel, Jan. 11–13, 2015. ACM.
- [26] J.-C. Faugère. A new efficient algorithm for computing Gröbner bases F_4 . *Journal of Pure and Applied Algebra*, 139(1-3):61–88, 1999.
- [27] J.-C. Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F_5). In T. Mora, editor, *ISSAC'2002*, pages 75–83. ACM Press, 2002.
- [28] U. Feige. Relations between average case complexity and approximation complexity. In *34th ACM STOC*, pages 534–543, Montréal, Québec, Canada, May 19–21, 2002. ACM Press.
- [29] V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. Statistical algorithms and a lower bound for detecting planted cliques. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 655–664, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.
- [30] V. Feldman, W. Perkins, and S. Vempala. On the complexity of random satisfiability problems with planted solutions. In R. A. Servedio and R. Rubinfeld, editors, *47th ACM STOC*, pages 77–86, Portland, OR, USA, June 14–17, 2015. ACM Press.
- [31] Y. Filmus, M. Lauria, M. Miksa, J. Nordström, and M. Vinyals. Towards an understanding of polynomial calculus: New separations and lower bounds - (extended abstract). In F. V. Fomin, R. Freivalds, M. Z. Kwiatkowska, and D. Peleg, editors, *ICALP 2013, Part I*, volume 7965 of *LNCS*, pages 437–448, Riga, Latvia, July 8–12, 2013. Springer, Heidelberg, Germany.
- [32] O. Goldreich. Candidate one-way functions based on expander graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(090), 2000.
- [33] W. T. Gowers. A new proof of szemerédi’s theorem. *Geometric and Functional Analysis*, 11(3):465–588, 2001.
- [34] R. Impagliazzo, P. Pudlák, and J. Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.
- [35] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Cryptography with constant computational overhead. In R. E. Ladner and C. Dwork, editors, *40th ACM STOC*, pages 433–442, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.
- [36] M. J. Kearns. Efficient noise-tolerant learning from statistical queries. *J. ACM*, 45(6):983–1006, 1998.
- [37] L. A. Levin. Universal sequential search problems. *PINFTRANS: Problems of Information Transmission (translated from Problemy Peredachi Informatsii (Russian))*, 9, 1973.
- [38] E. Mossel, A. Shpilka, and L. Trevisan. On e-biased generators in NC0. In *44th FOCS*, pages 136–145, Cambridge, Massachusetts, USA, Oct. 11–14, 2003. IEEE Computer Society Press.

- [39] R. O’Donnell and D. Witmer. Goldreich’s PRG: evidence for near-optimal polynomial stretch. In *Proc. of IEEE 29th Conference on Computational Complexity, CCC*, pages 1–12, 2014.
- [40] J. Patarin. Cryptanalysis of the Matsumoto and Imai public key scheme of eurocrypt’88. In D. Coppersmith, editor, *CRYPTO’95*, volume 963 of *LNCS*, pages 248–261, Santa Barbara, CA, USA, Aug. 27–31, 1995. Springer, Heidelberg, Germany.
- [41] C. E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28-4:656–715, 1949.
- [42] T. Siegenthaler. Correlation-immunity of nonlinear combining functions for cryptographic applications. *IEEE Transactions on Information Theory*, 30(5):776–778, 1984.
- [43] Toda. Classes of arithmetic circuits capturing the complexity of computing the determinant. *TIEICE: IEICE Transactions on Communications/Electronics/Information and Systems*, 1992.