Electronic Colloquium on Computational Complexity, Report No. 208 (2026)

Exponential Correlation Bounds for Polynomials

Eshan Chattopadhyay* Cornell University

Pooya Hatami† Ohio State University

Chin Ho Lee‡ NC State University

Shachar Lovett§ UC San Diego

Avishay Tal¶ UC Berkeley

Emanuele Viola‖ Northeastern University

September 23, 2026

Abstract

We prove that the XOR of k majorities on disjoint blocks of ℓ bits has correlation at most $(2d/\sqrt{\ell})^k$ with every degree-d polynomial over $\mathbb{F}_2$. By known techniques, this implies pseudorandom generators with polylogarithmic seed length for low-degree polynomials over $\mathbb{F}_2$ and for alternating circuits with parity gates.

1 Introduction

Proving *correlation bounds for low-degree polynomials* is a fundamental challenge in theoretical computer science. Such bounds have received a lot of attention, because they are a prerequisite for progress on a number of other long-standing challenges in computational complexity, related to matrix rigidity, multiparty communication complexity, circuit lower bounds, pseudorandomness, and more. For background on correlation bounds, see the survey [Vio22] or the book [Vio]. For two functions $f, g\colon \mathbb{F}_2^n \to \mathbb{F}_2$, we define their correlation by

$$\mathrm{corr}(f,g) := |\mathbb{E}_x[(-1)^{f(x)+g(x)}]|,$$

where x is uniform over $\mathbb{F}_2^n$.

For degree-d polynomials over $\mathbb{F}_2$, the fundamental works of Razborov [Raz87] and Smolensky [Smo87] proved correlation $< cd/\sqrt{n}$ with the Majority function on n bits; this is tight [Vio21]. Remarkably, for polynomials with degree $d \geq \log n$, no explicit function (e.g.

*eshan@cs.cornell.edu. Supported by NSF Award CCF-2514586.

†hatami.2@osu.edu.

‡chinho.lee@ncsu.edu.

§slovett@ucsd.edu. Supported by Simons Investigator Award #929894 and NSF Award CCF-2425349.

¶atal@berkeley.edu. Supported by NSF CAREER award CCF-2145474.

‖mathematicsoftheimpossible@gmail.com. Supported by NSF award 2430026.

ISSN 1433-8092

in NP) was known to have asymptotically smaller correlation. Improving this bound was a prerequisite for each of the other long-standing challenges mentioned above.

Here, we finally solve this challenge and in fact obtain exponential correlation bounds, even for polynomials of degree a power in the input length. Our results apply to any *versatile* function, as discussed below. To keep this exposition focused, we consider polynomials over $\mathbb{F}_2$ and their correlation with the XOR of Majority on k disjoint inputs of ℓ bits, for odd ℓ, denoted $\mathsf{Maj}_\ell^{\oplus k}$. This is a natural candidate that was considered in previous works [Vio06, VW08, CHH+20, Vio21].

Theorem 1.1. *For every polynomial $p\colon \mathbb{F}_2^{k\ell} \to \mathbb{F}_2$ of total degree at most d we have*

$$\operatorname{corr}(\mathsf{Maj}_\ell^{\oplus k}, p) \leq \left(\frac{2d}{\sqrt{\ell}}\right)^k.$$

The analysis actually gives a constant better than 2, namely $\sqrt{2/\pi} \cdot 2 \approx 1.596$. This bound is tight up to the constant. This follows by the construction in [Vio21] applied to each block.

Applications. This result has several applications. Applying the construction by Nisan and Wigderson [Nis91, NW94] based on our hard function $\mathsf{Maj}_\ell^{\oplus k}$ gives seed length $O\big(d^4 \cdot \mathsf{polylog}(n/\varepsilon)\big)$ for degree-d polynomials over $\mathbb{F}_2$ on n variables. Furthermore, the PRG construction of [CGL+21] gives the improved seed length $O(d^2 \cdot \mathsf{polylog}(n/\varepsilon))$. The seed length of previous generators [Bog05, Lov09, Vio09] had an exponential dependence on the degree, and so gave nothing for $d \geq \log n$. We provide details in Appendices A and B.

Combining Razborov's polynomial approximation [Raz87] with Nisan and Wigderson's construction [Nis91, NW94] (see also [Vio]) gives pseudorandom generators (PRGs) with polylogarithmic seed length for polynomial-size, alternating circuits with parity gates (i.e., $\mathsf{AC}^0[\oplus]$). The previous best was seed length $n(1 - o(1))$ [FSUV13]. We provide details in Appendix C.

2 Proof of Main Result

The proof of Theorem 1.1 is an extension of Smolensky's celebrated encoding argument to tuples. We introduce necessary notation and preliminaries in Section 2.1. A new notion of weights for functions is introduced in Section 2.2. We prove the multiplication lemma in Section 2.2 and the main theorem by dimension counting in Section 2.3.

2.1 Preliminaries and Notation

For convenience, set $\ell := 2m + 1$ and write h for the Boolean Majority function on ℓ bits, so $h(z) = 1$ exactly when $|z| \geq m + 1$, where $|z|$ is the Hamming weight of z. Let $n = k\ell$. An n-bit input x is partitioned into k blocks $x^{(1)}, \ldots, x^{(k)}$ of length ℓ. For the remainder of the paper, for ease of notation, we use the shorthand XMaj for $\mathsf{Maj}_\ell^{\oplus k}$. We also let $h_i(x) := h(x^{(i)})$ so that

$$\mathsf{XMaj}(x) = \bigoplus_{i=1}^{k} h_i(x).$$

Interpolation. We first review the basic interpolation in (some expositions of) Smolensky's correlation bound with (one) majority. Let us consider one block ($k = 1$) of ℓ bits. Let V be the space of multilinear polynomials of degree at most m. Its dimension is $\sum_{r=0}^{m} \binom{\ell}{r} = 2^{\ell-1}$. We use the fact that every function on either of the two sets

$$H_0 \coloneqq \{z \in \{0,1\}^\ell : |z| \le m\}, \quad H_1 \coloneqq \{z \in \{0,1\}^\ell : |z| \ge m+1\}$$

agrees with a unique polynomial in V. We call this its *interpolating polynomial*. Consequently, every function $f \colon \mathbb{F}_2^\ell \to \mathbb{F}_2$ can be written uniquely as

$$f = q + hr, \qquad q, r \in V. \tag{1}$$

(This property is also referred to as h being a *versatile* function, see [Kop11]. In fact, this is the only property of Majority that we are using in the proof.) Indeed, let q agree with f on H_0, let s agree with f on H_1, and set $r = q + s$. Then $q + hr$ equals $q = f$ on H_0 and $q + r = s = f$ on H_1. For uniqueness, the values of f on H_0 determine q, and its values on H_1 then determine $q + r$.

Next we apply this interpolation separately in each block. For $\alpha = (\alpha_1, \dots, \alpha_k) \in \{0,1\}^k$, set

$$H_\alpha \coloneqq H_{\alpha_1} \times \cdots \times H_{\alpha_k}$$

and

$$h_\alpha(x) \coloneqq \prod_{i:\alpha_i=1} h_i(x) \prod_{i:\alpha_i=0} (1 - h_i(x)) = \prod_{i=1}^{k} \mathbf{1}[h_i(x) = \alpha_i].$$

Thus H_α consists of the inputs whose block majorities are $\alpha_1, \dots, \alpha_k$, and h_α is its indicator. Write $\mathrm{par}(\alpha) = \bigoplus_{i=1}^{k} \alpha_i$ for the parity of α.

An important space in this work consists of polynomials where the degree *in each block* is bounded.

Definition 1. Let W denote the space of multilinear polynomials in $k\ell$ variables of degree at most m in each block.

We can now state and prove the new interpolation result.

Lemma 2.1. *Every function $f \colon \mathbb{F}_2^{k\ell} \to \mathbb{F}_2$ can be written uniquely as*

$$f = \sum_{\alpha \in \{0,1\}^k} f_\alpha h_\alpha, \qquad f_\alpha \in W. \tag{2}$$

Here f_α is the unique polynomial in W that agrees with f on H_α.

In the following, it is useful to think of the individual h_i as new formal variables.

Proof. Fix α. We interpolate successively in the k blocks to obtain a unique polynomial $f_\alpha \in W$ agreeing with f on H_α. Specifically, this step is done by induction. Fix all the blocks except the first. We can interpolate the resulting function on that block. The coefficients of this interpolation are functions of the other blocks, which can be interpolated by induction.

At every input, exactly one indicator h_α equals 1, so the sum in Equation (2) equals f. Conversely, restricting this identity to H_α determines f_α by uniqueness of interpolation. □

The h-basis. We next specify a convenient basis for functions. A monomial of degree at most m in each of the k blocks is specified by a tuple of subsets $\mathbf{A} = (A_1, \ldots, A_k)$, where $A_i \subseteq [\ell]$ and $|A_i| \leq m$. Let $\mathcal{M}$ be the set of all such tuples, and write

$$x^{\mathbf{A}} \coloneqq \prod_{i=1}^{k} \prod_{j \in A_i} x_j^{(i)}.$$

The degree of this monomial in the i-th block is simply $|A_i|$. Expanding each function f_α as a linear combination of the monomials $x^{\mathbf{A}}$, and each function h_α as a linear combination of $h^S \coloneqq \prod_{i \in S} h_i$, we get that any Boolean function f can be uniquely expressed as

$$f = \sum_{\mathbf{A} \in \mathcal{M}, S \subseteq [k]} c_{\mathbf{A},S}(f) x^{\mathbf{A}} h^S.$$

We call $\{x^{\mathbf{A}} h^S\}_{\mathbf{A},S}$ the *h-basis*.

2.2 Weight

We define a notion of *weight* w for functions, based on the expansion in the h-basis with the useful property that it increases by at most t under multiplication by a degree-t polynomial. For a basis function $x^{\mathbf{A}} h^S$, we define its weight by letting every block $i \in S$ contribute $m+1$, and every block $i \notin S$ contribute $|A_i|$:

$$w(x^{\mathbf{A}} h^S) \coloneqq (m+1)|S| + \sum_{i \notin S} |A_i|.$$

For nonzero f, let $w(f)$ be the largest weight of a basis function with nonzero coefficient in its unique expansion, and set $w(0) = -\infty$. Note that this is not the usual polynomial degree.[1]

Lemma 2.2 (Multiplication Lemma)**.** *Let $f \colon \mathbb{F}_2^{k\ell} \to \mathbb{F}_2$ be a function, and q be a polynomial in $k\ell$ variables of total degree at most t. Then*

$$w(qf) \leq w(f) + t.$$

Proof. Pick a basis function (with nonzero coefficient) for f. Let it be $b \coloneqq x^{\mathbf{A}} h^S$. We prove that after multiplication by q the weight increases by at most t. First we prove it when q is a single variable z. Suppose $z = x_e^{(i)}$ is a variable in block i. There are three cases to consider:

1. If $i \notin S$ and $|A_i| < m$, the weight increases by at most one.
2. Assume $i \notin S$ and $|A_i| = m$. This is the boundary case where the block has already reached its maximum degree and does not have the factor h_i. For this case we use Equation (1) to rewrite the monomial $z \prod_{j \in A_i} x_j^{(i)}$ as $q_i + h_i r_i$, with $q_i, r_i \in V$. Every resulting term contributes at most $m+1$ in block i, again increasing the weight by at most one.

[1] Instead, weight is a capped version of polynomial degree in the following sense: for each monomial in the ordinary multilinear representation of f, we sum its block degrees capped at $m+1$, and then take the maximum over monomials with nonzero coefficients.

3. Finally, suppose $i \in S$. Again using Equation (1) write $z \prod_{j \in A_i} x_j^{(i)} = q_i + h_i r_i$, with $q_i, r_i \in V$. Since $h_i^2 = h_i$, the new factor in block i is $h_i(q_i + h_i r_i) = h_i(q_i + r_i)$. Every nonzero resulting term still contains h_i and therefore contributes $m + 1$ in this block. The weight does not increase in this case.

This verifies the case of one variable. To show the bound for an arbitrary monomial, repeat for each variable. For an arbitrary polynomial q, apply to each monomial, and then add the results. Combining equal terms can cancel coefficients but cannot introduce a term of greater weight. □

2.3 Proof of the Main Theorem by Dimension Counting

We will use the weight defined above and the multiplication lemma (Lemma 2.2). The role of the weight will be to show that a finite rewriting procedure terminates.

Let p be a multilinear polynomial in the $k\ell$ variables of total degree at most d. Define

$$E \coloneqq \{x \in \{0,1\}^{k\ell} : p(x) = \mathsf{XMaj}(x)\}.$$

We need to show that the relative size of E is close to $1/2$. To do so we will precisely follow Smolensky's idea and show that every function on E can be written as (the restriction of) a function from a space of dimension only slightly larger than half the dimension of the full function space. The new space is defined in the h-basis as follows. For each monomial $x^{\mathbf{A}}$ with a low-degree block, we shall pick the first such block j and keep only the 2^{k-1} basis terms $\{x^{\mathbf{A}} h^S\}_{S \subseteq [k] \setminus \{j\}}$ not involving h_j. For monomials $x^{\mathbf{A}}$ whose degree is high in *every block*, we keep all 2^k terms $\{x^{\mathbf{A}} h^S\}_{S \subseteq [k]}$. A simple calculation bounds the dimension of the space spanned by these basis terms, which then bounds $|E|$.

We define the set of exceptional monomial indices

$$\mathcal{B} \coloneqq \{\mathbf{A} \in \mathcal{M} : |A_i| > m - d \text{ for every } i \in [k]\}.$$

For every non-exceptional $\mathbf{A} \notin \mathcal{B}$ choose a low-degree block:

$$j(\mathbf{A}) \coloneqq \min\{i \in [k] : |A_i| \le m - d\}.$$

Define $\mathcal{R}$ to be the $\mathbb{F}_2$-linear span of the following basis functions:

$$\{x^{\mathbf{A}} h^S : \mathbf{A} \in \mathcal{B},\ S \subseteq [k]\} \cup \{x^{\mathbf{A}} h^S : \mathbf{A} \notin \mathcal{B},\ S \subseteq [k],\ j(\mathbf{A}) \notin S\} \tag{3}$$

In words, for a non-exceptional monomial $x^{\mathbf{A}}$ we choose one low-degree block $j(\mathbf{A})$ and forbid all products containing the factor $h_{j(\mathbf{A})}$. For an exceptional monomial we impose no restriction.

Lemma 2.3. *For every function $a\colon E \to \mathbb{F}_2$, there exists a function $r \in \mathcal{R}$ such that $r(x) = a(x)$ for every $x \in E$.*

Proof. As any function $a\colon E \to \mathbb{F}_2$ can be expressed as a linear combination of basis functions $\{x^{\mathbf{A}} h^S\}_{\mathbf{A} \in \mathcal{M}, S \subseteq [k]}$, it suffices to prove the claim for each basis function separately. Let

$$T = x^{\mathbf{A}} h^S.$$

The proof is by induction on the weight of T. If $w(T) = 0$ then T is a constant function which belongs to $\mathcal{R}$. Assume from now that $w(T) \geq 1$. If T is one of the allowed terms in (3) there is nothing to prove. Otherwise we have $\mathbf{A} \notin \mathcal{B}$ and $j = j(\mathbf{A})$ belongs to S. Set

$$Q \coloneqq x^{\mathbf{A}} h^{S \setminus \{j\}}.$$

On E we have

$$p = \mathsf{XMaj} = h_1 + \cdots + h_k,$$

and therefore

$$h_j = p + \sum_{i \neq j} h_i.$$

Multiplying this last equation by Q gives, on E,

$$T = Qp + \sum_{i \neq j} Q h_i. \tag{4}$$

Every term Qh_i on the right is allowed because it does not contain the factor h_j. Repeated factors are simplified using $h_i^2 = h_i$.

It remains to handle Qp. Since h_j contributes $m + 1$ to $w(T)$, whereas after removing h_j the j-th block contributes only $|A_j|$ to $w(Q)$, we have

$$w(Q) = w(T) - \big((m+1) - |A_j|\big).$$

Because $|A_j| \leq m - d$,

$$w(Q) \leq w(T) - (d+1).$$

The multiplication lemma (Lemma 2.2) then gives

$$w(Qp) \leq w(Q) + d \leq w(T) - 1.$$

Expand Qp in the full basis $\{x^{\mathbf{A}} h^S\}_{\mathbf{A},S}$. Every basis term (with non-zero coefficient) in this expansion has weight strictly smaller than $w(T)$. We may therefore apply induction to each of these terms. Equation (4) then expresses T, on E, as a linear combination of allowed generators. □

We now count the dimension of $\mathcal{R}$. Let

$$M_\ell \coloneqq \sum_{r=0}^{m} \binom{\ell}{r} = 2^{\ell-1}, \qquad N_{\ell,d} \coloneqq \sum_{r=\max\{0,m-d+1\}}^{m} \binom{\ell}{r}.$$

There are M_ℓ^k ordinary monomials in $\mathcal{M}$, and exactly $N_{\ell,d}^k$ exceptional monomials in $\mathcal{B}$. For each nonexceptional monomial $x^{\mathbf{A}}$, exactly 2^{k-1} subsets $S \subseteq [k]$ avoid $j(\mathbf{A})$. For each exceptional monomial, all 2^k subsets are allowed. Hence

$$\dim(\mathcal{R}) = 2^{k-1}\big(M_\ell^k - N_{\ell,d}^k\big) + 2^k N_{\ell,d}^k = 2^{k-1} M_\ell^k + 2^{k-1} N_{\ell,d}^k. \tag{5}$$

By Lemma 2.3, the restriction map $\mathcal{R} \to \mathbb{F}_2^E$ is surjective. Since the vector space of functions on E has dimension $|E|$, Equation (5) gives

$$|E| \le 2^{k-1} M_\ell^k + 2^{k-1} N_{\ell,d}^k = 2^{k\ell-1} + 2^{k-1} N_{\ell,d}^k.$$

Consequently,

$$\frac{2|E| - 2^{k\ell}}{2^{k\ell}} \le \left(\frac{N_{\ell,d}}{2^{\ell-1}}\right)^k.$$

Applying the same argument to $p+1$ replaces E by the disagreement set. Therefore

$$\mathrm{corr}(\mathsf{XMaj}, p) \le \left(\frac{N_{\ell,d}}{2^{\ell-1}}\right)^k.$$

Finally, there are at most d summands in $N_{\ell,d}$, each at most $\binom{\ell}{m} \le \sqrt{\frac{2}{\pi}} \frac{2^\ell}{\sqrt{\ell}}$. Thus

$$\mathrm{corr}(\mathsf{XMaj}, p) \le \left(\sqrt{\frac{2}{\pi}} \frac{2d}{\sqrt{\ell}}\right)^k \le \left(\frac{2d}{\sqrt{\ell}}\right)^k. \quad \square$$

Acknowledgment. This work was done while most of the authors were at the Simons Institute for Theory of Computing, attending the Fall 2026 Pseudorandomness & High-Dimensional Expansion program, which provided an ideal working environment.

AI methodology. This work began with several authors prompting AI tools about correlation bounds for the XOR of Majority. One of these interactions eventually produced a new correlation bound, albeit much weaker than the one presented in this paper, and with a much more complicated proof. The authors then undertook a research project to improve the bounds and simplify the argument through repeated interactions with AI tools, interspersed with their own insights and simplifications. Some intermediate proofs were extremely complex; the final proof emerged through many rounds of refinement.

References

[Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *ACM Symp. on the Theory of Computing (STOC)*, pages 21–30, 2005. (cit. on page 2.)

[CGL+21] Eshan Chattopadhyay, Jason Gaitonde, Chin Ho Lee, Shachar Lovett, and Abhishek Shetty. Fractional pseudorandom generators from any Fourier level. In *36th Computational Complexity Conference (CCC 2021)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:24. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021. (cit. on pages 2, 10, and 12.)

[CHH+20] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, Shachar Lovett, and David Zuckerman. Xor lemmas for resilient functions against polynomials. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 234–246, 2020. (cit. on page 2.)

[CHHL19] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. Pseudorandom generators from polarizing random walks. *Theory Comput.*, 15:1–26, 2019. (cit. on page 10.)

[FSUV13] Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. On beating the hybrid argument. *Theory of Computing*, 9:809–843, 2013. (cit. on page 2.)

[HH24] Pooya Hatami and William Hoza. Paradigms for unconditional pseudorandom generators. *Foundations and Trends in Theoretical Computer Science*, 16(1–2):1–210, 2024. (cit. on page 9.)

[Kop11] Swastik Kopparty. On the complexity of powering in finite fields. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 489–498, 2011. (cit. on page 3.)

[Lov09] Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. *Theory of Computing*, 5(1):69–82, 2009. (cit. on page 2.)

[Nis91] Noam Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica. A Journal on Combinatorics and the Theory of Computing*, 11(1):63–70, 1991. (cit. on pages 2 and 9.)

[NW94] Noam Nisan and Avi Wigderson. Hardness vs Randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994. (cit. on pages 2 and 9.)

[Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in Mathematical Notes of the Academy of Sci. of the USSR, 41(4):333-338, 1987. (cit. on pages 1, 2, and 12.)

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *19th ACM Symp. on the Theory of Computing (STOC)*, pages 77–82. ACM, 1987. (cit. on page 1.)

[Vio] Emanuele Viola. *Mathematics of the impossible*. (cit. on pages 1 and 2.)

[Vio06] Emanuele Viola. New correlation bounds for GF(2) polynomials using Gowers uniformity. *Electronic Colloquium on Computational Complexity*, Technical Report TR06-097, 2006. www.eccc.uni-trier.de/. (cit. on page 2.)

[Vio09] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *Computational Complexity*, 18(2):209–217, 2009. (cit. on page 2.)

[Vio21] Emanuele Viola. Fourier conjectures, correlation bounds, and majority. In *Coll. on Automata, Languages and Programming (ICALP)*, 2021. Available at http://www.ccs.neu.edu/home/viola/. (cit. on pages 1 and 2.)

[Vio22] Emanuele Viola. Correlation bounds against polynomials, a survey. 2022. (cit. on page 1.)

[VW08] Emanuele Viola and Avi Wigderson. Norms, xor lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(7):137–168, 2008. (cit. on page 2.)

A PRG from Nisan-Wigderson Construction

We present the details of the PRG consequence stated in the introduction. We use the construction of Nisan and Wigderson [NW94], originating in [Nis91], with a separate seed for each majority block. Replacing the output bits one at a time in the hybrid argument, keeping the XOR across the majority blocks, lets us bound the degree of each restricted output without paying a factor for the number of blocks. Directly applying the construction to the entire hard function would give a slightly worse seed length, as we explain below.

Throughout the appendices, logarithms are base two, U_r denotes the uniform distribution on $\{0,1\}^r$, and $0 < \varepsilon \le 1/2$. A map $G\colon \{0,1\}^r \to \{0,1\}^n$ ε-fools a class of Boolean functions if $|\Pr[f(G(U_r)) = 1] - \Pr[f(U_n) = 1]| \le \varepsilon$ for every function f in the class.

Theorem A.1. *There is an explicit PRG that ε-fools degree-d polynomials over $\mathbb{F}_2$ on n variables with seed length $O\big(d^4 \log^3(n) \log(n/\varepsilon)\big)$.*

Construction. Set $t = \lceil \log(2n) \rceil$. For any $\ell \ge t$, we can construct sets $S_1, \ldots, S_n \subseteq [R]$ in time polynomial in n and ℓ, where

$$|S_i| = \ell, \qquad |S_i \cap S_j| < t \quad (i \ne j), \qquad R = O(\ell^2/t).$$

(See, for example, [HH24, Lemma 4.5].)

Choose an odd $\ell \ge 16 d^2 t^2$ with $\ell = O(d^2 t^2)$, and set $k = \lceil \log(2n/\varepsilon) \rceil$. Using k independent seeds $z^{(1)}, \ldots, z^{(k)} \in \{0,1\}^R$, define

$$G(z^{(1)}, \ldots, z^{(k)})_i = \bigoplus_{a=1}^{k} \mathrm{Maj}(z^{(a)}_{S_i}), \qquad i \in [n].$$

Here $z^{(a)}_{S_i}$ lists the coordinates in increasing order. The seed length is

$$kR = O(k\ell^2/t) = O\big(d^4 \log^3(n) \log(n/\varepsilon)\big).$$

Analysis. Fix a degree-d polynomial p. For $0 \le i \le n$, let Y_i be a random vector in $\{0,1\}^n$ consisting of the first i outputs of G followed by $n - i$ independent uniform bits. Thus Y_0 is uniform and Y_n is the output of G. To compare Y_i and Y_{i-1}, fix the seed bits outside S_i in each of the k seeds, and fix the last $n - i$ output bits. The remaining seed bits form k independent uniform blocks $x^{(1)}, \ldots, x^{(k)} \in \{0,1\}^\ell$. Write $F(x) = \mathsf{XMaj}(x)$ for the i-th output.

For $j < i$, the j-th output is now of the form $q_j(x) = \bigoplus_{a=1}^{k} q_{j,a}(x^{(a)})$, where $q_{j,a}$ depends on fewer than t variables, since $|S_j \cap S_i| < t$. Note that $\deg q_j \le t$, since the XOR across the blocks is addition over $\mathbb{F}_2$, so it does not add their degrees. For $b \in \{0,1\}$, let $Q_b(x)$ be the result of substituting these earlier outputs, the bit b, and the fixed later bits into p. Both Q_0 and Q_1 have total degree at most dt.

Treating Boolean outputs as real numbers, the conditional difference between the two acceptance probabilities is

$$\mathbb{E}_x\left[Q_{F(x)}(x) - \frac{Q_0(x)+Q_1(x)}{2}\right] = \frac{1}{4}\mathbb{E}_x\left[(-1)^{F(x)+Q_1(x)} - (-1)^{F(x)+Q_0(x)}\right].$$

By Theorem 1.1, its absolute value is at most

$$\frac{1}{2}\left(\frac{2dt}{\sqrt{\ell}}\right)^k \le 2^{-k-1} \le \frac{\varepsilon}{4n}.$$

This holds for every fixing. Averaging over the fixed bits and summing over the n hybrids proves Theorem A.1.

Remark. The degree bound in the hybrid argument is the reason for applying NW separately in each majority block. Although q_j may depend on kt variables, it is an XOR of functions of at most t variables, so its degree is at most t. Thus we pay for k independent seeds of length $O(\ell^2/t)$, while the polynomials Q_b still have degree at most dt. If instead we apply NW directly to the entire $k\ell$-bit hard function, using sets of size $k\ell$ with intersections of size at most t, the same correlation estimate requires $\ell = \Theta(d^2t^2)$ and $k = \Theta(\log(n/\varepsilon))$. The design then gives seed length $O((k\ell)^2/t) = O\big(d^4 \log^3(n) \log^2(n/\varepsilon)\big)$. Using the XOR structure thus saves a factor of $k = \Theta(\log(n/\varepsilon))$.

B Fourier Bounds for Polynomials and an alternate PRG

Our correlation bound resolves a conjectured Fourier bound for polynomials [CGL+21, Conjecture 31]. As a consequence, using a result from [CGL+21], based on the frameworks of polarizing random walks [CHHL19], we derive an alternative pseudorandom generator for polynomials.

For $p\colon \mathbb{F}_2^n \to \mathbb{F}_2$, define its sign function on $\{-1,1\}^n$ by $f(x) = (-1)^{p((1-x_1)/2,\ldots,(1-x_n)/2)}$. Write its real Fourier expansion as $f(x) = \sum_{S\subseteq[n]} \widehat{f}(S)x_S$, where $x_S = \prod_{i\in S} x_i$ and $\widehat{f}(S) = \mathbb{E}_x[f(x)x_S]$ for uniform $x \in \{-1,1\}^n$. Define

$$M_k(f) = \max_{x\in\{-1,1\}^n}\left|\sum_{|S|=k} \widehat{f}(S)x_S\right|.$$

Thus the absolute value is taken after summing the level-k coefficients with the signs specified by x. Let $\mathcal{P}_{n,d}$ be the class of degree-d polynomials on n variables over $\mathbb{F}_2$, and set $M_k(\mathcal{P}_{n,d}) = \max_{p\in\mathcal{P}_{n,d}} M_k((-1)^p)$.

Theorem B.1. *For $1 \le k \le n$ and $d \ge 1$,*

$$M_k(\mathcal{P}_{n,d}) \le \frac{k^k}{k!}(2d)^k \le (2ed)^k.$$

Proof. Note that negating input coordinates multiplies $\widehat{f}(S)$ by the corresponding product of signs. Thus it suffices to bound $|\sum_{|S|=k} \widehat{f}(S)|$, where f is the sign function of an arbitrary $p \in \mathcal{P}_{n,d}$. Fix a partition $B_1, \dots, B_k$ of $[n]$, allowing empty parts. Pad each part with fresh variables to the same odd length ℓ, and let f ignore these additional variables. On $\{-1, +1\}$, majority is the function $h_\ell(y) = \operatorname{sgn}(y_1 + \cdots + y_\ell)$. Using the exact constant in the proof of Theorem 1.1 gives

$$\left| \mathbb{E}\left[f(x) \prod_{a=1}^{k} h_\ell(x^{(a)}) \right] \right| \le \left(2\sqrt{\frac{2}{\pi}} \frac{d}{\sqrt{\ell}} \right)^k, \tag{6}$$

where $x^{(a)}$ contains the variables of B_a and its padding.

We recover the sum of coefficients that use one variable from each part by letting ℓ grow. We prove this as follows. Write $\ell = 2m + 1$ and $a_\ell = \widehat{h}_\ell(\{1\}) = 2^{-2m}\binom{2m}{m}$. The even-level Fourier coefficients of h_ℓ vanish, and for odd $|T| = 2j + 1$,

$$\widehat{h}_\ell(T) = (-1)^j a_\ell \frac{\binom{m}{j}}{\binom{2m}{2j}}.$$

Indeed, differentiating in one coordinate of T gives the indicator that the other $2m$ signs sum to zero. On this event, the expected product of any $2j$ coordinates is $(-1)^j\binom{m}{j}/\binom{2m}{2j}$, as follows by expanding $(1+z)^m(1-z)^m$. Thus $\sqrt{\ell}\, a_\ell \to \sqrt{2/\pi}$, and for fixed odd $|T| > 1$ we have $\widehat{h}_\ell(T)/a_\ell \to 0$.

Expanding in the Fourier basis and dividing by a_ℓ^k gives

$$\frac{1}{a_\ell^k}\mathbb{E}\left[f(x) \prod_{a=1}^{k} h_\ell(x^{(a)}) \right] = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{a=1}^{k} \frac{\widehat{h}_\ell(S \cap B_a)}{a_\ell}.$$

If S misses a part, the corresponding factor is zero because majority has mean zero. Otherwise, each factor is 1 when $|S \cap B_a| = 1$, and tends to zero when $|S \cap B_a| > 1$. Since the sum has only 2^n terms, taking $\ell \to \infty$ leaves exactly those sets containing one variable from each part. By Equation (6) and $\sqrt{\ell}\, a_\ell \to \sqrt{2/\pi}$, we obtain

$$\left| \sum_{S:\, |S \cap B_a| = 1 \text{ for every } a} \widehat{f}(S) \right| \le (2d)^k.$$

Now assign each coordinate independently and uniformly to one of the k parts. Every k-element set meets all parts with probability $k!/k^k$. Averaging the last inequality gives

$$\frac{k!}{k^k} \left| \sum_{|S|=k} \widehat{f}(S) \right| \le (2d)^k. \qquad \square$$

PRG consequence. We recall the PRG result that we need from [CGL+21]. If a class $\mathcal{F}$ of sign functions on n variables is closed under restrictions and satisfies $M_k(\mathcal{F}) \le b^k$ for $k \ge 3$ and $1 \le b \le n$, then it has an explicit PRG of error ε with seed length

$$O\left(b^2 k \log(n) \log(n/\varepsilon) \left(\frac{b^2 \log(n/\varepsilon)}{\varepsilon}\right)^{2/(k-2)}\right).$$

Using the Fourier bounds proved above, taking $b = 2ed$ and $k = O\left(\log\left(\frac{b^2 \log(n/\varepsilon)}{\varepsilon}\right)\right)$, we obtain the following result.

Corollary B.2. *There is an explicit PRG that ε-fools degree-d polynomials over $\mathbb{F}_2$ on n variables with seed length*

$$O\left(d^2 \log(n) \log(n/\varepsilon) \log\left(\frac{2d^2 \log(n/\varepsilon)}{\varepsilon}\right)\right).$$

C PRG for $\mathsf{AC}^0[\oplus]$

Using Razborov's polynomial approximation of $\mathsf{AC}^0[\oplus]$, we easily derive the following PRGs for $\mathsf{AC}^0[\oplus]$ using our results on fooling polynomials.

Theorem C.1. *Fix a depth $\Delta \ge 1$. There are explicit PRGs that ε-fool size-s, depth-Δ $\mathsf{AC}^0[\oplus]$ circuits on n inputs with the following seed lengths. The Nisan–Wigderson construction gives*

$$O\big(\log^{4\Delta-1}(s) \log^4(1/\varepsilon) \log(s/\varepsilon)\big),$$

and the construction from Fourier bounds gives

$$O\left(\log^{2\Delta-1}(s) \log^2(1/\varepsilon) \log(s/\varepsilon) \log\left(\frac{\log(s)}{\varepsilon}\right)\right).$$

Proof. Fix such a circuit C. Razborov's polynomial approximation method [Raz87] gives a distribution of polynomials of degree $d = O(\log^{\Delta-1}(s) \log(1/\varepsilon))$ such that, for every input x, a sampled polynomial p satisfies $\Pr_p[p(x) \ne C(x)] \le \varepsilon/3$.

Let G be a PRG that $(\varepsilon/3)$-fools degree-d polynomials. The pointwise approximation guarantee holds under both the uniform distribution and the generator. Consequently,

$$\begin{aligned}
&|\Pr[C(G(U_r)) = 1] - \Pr[C(U_n) = 1]| \\
&\quad \le \mathbb{E}_p \Pr[p(G(U_r)) \ne C(G(U_r))] + \mathbb{E}_p \left|\Pr[p(G(U_r)) = 1] - \Pr[p(U_n) = 1]\right| \\
&\qquad + \mathbb{E}_p \Pr[p(U_n) \ne C(U_n)] \le \varepsilon.
\end{aligned}$$

Thus both the PRGs constructed in previous sections also fool $\mathsf{AC}^0[\oplus]$ circuits. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il