New correlation bounds for $GF(2)$ polynomials using Gowers uniformity

Emanuele Viola*

August 6, 2006

Abstract

We study the correlation between low-degree $GF(2)$ polynomials p and explicit functions. Our main results are the following:

I We prove that the Mod_m function on n bits has correlation at most $\exp\left(-\Omega\left(n/4^d\right)\right)$ with any $GF(2)$ polynomial of degree d, for any fixed odd integer m. This improves on the previous $\exp\left(-\Omega\left(n/8^d\right)\right)$ bound by Bourgain (C. R. Acad. Sci. Paris, 2005) and Green et al. (C. R. Acad. Sci. Paris, 2005).

II We exhibit a polynomial-time computable function on n bits that has correlation at most $\exp\left(-\Omega\left(n/2^d\right)\right)$ with any $GF(2)$ polynomial of degree d. Previous to our work the best correlation bound for an explicit function was $\exp\left(-\Omega\left(n/\left(d \cdot 2^d\right)\right)\right)$, which follows from (Chung and Tetali; SIAM J. Discrete Math., 1993).

III We derive an 'XOR Lemma' for low-degree $GF(2)$ polynomials: We show that if a function f has correlation at most $1 - 4^{-d}$ with any $GF(2)$ polynomial of degree d (and $\Pr_x[f(x) = 1] \approx 1/2$) then the XOR of m independent copies of f has correlation at most $\exp\left(-\Omega\left(m/4^d\right)\right)$ with any $GF(2)$ polynomial of degree d.

Our results rely on a measure of the 'complexity' of a function due to Gowers (Geom. Funct. Anal., 1998 & 2001).

*`viola@eecs.harvard.edu`. Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138. Research supported by NSF grant CCR-0133096, US-Israel BSF grant 2002246, ONR grant N-00014-04-1-0478.

1 Introduction

In this work we study the *correlation* between low-degree $GF(2)$ polynomials p and explicit functions $f : \{0,1\}^n \to \{+1,-1\}$, which is defined as

$$\text{Correlation}(p, f) := \left| \Pr_{x:f(x)=1}[p(x) = 1] - \Pr_{x:f(x)=-1}[p(x) = 1] \right|.$$

The study of correlation bounds is motivated, in large part, by their applications to proving *lower bounds* on the size of important classes of circuits with parity and majority gates. We refer the reader to, e.g., [AB, Gre] for a discussion of these applications. Here we content ourselves with recalling that exhibiting an explicit function which has negligible correlation with every $GF(2)$ polynomial of polylogarithmic degree would solve the famous open problem of establishing a superpolynomial lower bound on the size of constant-depth circuits with parity gates and one majority gate.[1] An additional motivation for studying correlation bounds is that balanced functions with negligible correlation with low-degree $GF(2)$ polynomials can be used to construct *pseudorandom generators* that fool polynomial-size $GF(2)$ polynomials and, more generally, polynomial-size constant-depth circuits with few parity gates. This follows from the techniques in [Vio1, Han].

Besides the above two motivations, the amount of research devoted to proving correlation bounds (e.g., [Raz, Smo, AB, Gre, Bou, GRS]), as well as the innovative techniques developed for this purpose, make the study of correlation bounds interesting in its own right. Finally, our ability to prove correlation bounds is a fundamental benchmark for our understanding of complexity theory: Currently no explicit function on n bits is known to have negligible correlation with $GF(2)$ polynomials of degree $\log_2 n$.

Our results. In this work we obtain the following results.

I We prove that the Mod_m function on n bits has correlation at most $\exp\left(-\Omega\left(n/4^d\right)\right)$ with any $GF(2)$ polynomial of degree d, for any fixed odd integer m. This improves on a result by Bourgain [Bou] and Green et al. [GRS] that shows that the same correlation is at most $\exp\left(-\Omega\left(n/8^d\right)\right)$.[2]

II We exhibit a polynomial-time computable function on n bits that has correlation at most $\exp\left(-\Omega\left(n/2^d\right)\right)$ with any $GF(2)$ polynomial of degree d. Previous to our work the best correlation bound for an explicit function was $\exp\left(-\Omega\left(n/\left(d \cdot 2^d\right)\right)\right)$, which follows from the multiparty communication complexity results by Chung and Tetali [CT] and the fact that any function computed by a polynomial of degree d can be computed by a multiparty communication complexity protocol (with $d+1$ parties) exchanging few bits [HG, Proof of Lemma 4].

[1] This follows by the so-called "ϵ-discriminator lemma" in [HMP$^+$] and the fact that any function computed by a polynomial-size constant-depth circuit with parity gates can be approximated by a $GF(2)$ polynomial of polylogarithmic degree [Raz, Smo].

[2] Bourgain's proof [Bou] contains all the main ideas but is slightly incorrect. A correct proof is given by Green et al. [GRS].

III We derive an 'XOR Lemma' for low-degree $GF(2)$ polynomials: We show that if a function $f : \{0,1\}^n \to \{+1,-1\}$ has correlation at most $1 - 4^{-d}$ with any degree d $GF(2)$ polynomial (and $\Pr_x[f(x) = 1] \approx 1/2$) then the function $g : \{0,1\}^{n \cdot m} \to \{+1,-1\}$ defined as $g(x_1, x_2, \ldots, x_m) := \prod_{i \le m} f(x_i)$ has correlation at most $\exp\left(-\Omega\left(m/4^d\right)\right)$ with any degree d $GF(2)$ polynomial. Previous to our work such a result was not even known for polynomials of degree $d = 2$.

As we mentioned at the beginning of this introduction, correlation bounds imply circuit lower bounds. To briefly illustrate some consequences of our results in this direction, we remark that by combining the "ϵ-discriminator lemma" by Hajnal et al. [HMP+] with our result (I), one obtains the following: Computing the Mod_3 function via the majority of s $GF(2)$ polynomials of degree d requires at least $s \ge \exp\left(\Omega\left(n/4^d\right)\right)$ polynomials. (Equivalently, using the notation in [AB], computing the Mod_3 function via $Maj \circ Parity \circ And_d$ circuits requires at least $\exp\left(\Omega\left(n/4^d\right)\right)$ parity gates.) The previous best lower bound was $\exp\left(\Omega\left(n/8^d\right)\right)$ [GRS]. Similar considerations hold for our result (II).

Techniques. Our results rely on a measure of the 'complexity' of a function that has become known as *Gowers uniformity*. This uniformity was introduced by Gowers [Gow1, Gow2] and has been studied and applied in computer science in the area of property testing by Alon et al. [AKK+] (who actually defined Gowers uniformity independently) and Samorodnitsky and Trevisan [Sam, ST] (who also applied it to probabilistically checkable proofs). The Gowers uniformity of a function f is parameterized by an integer k, and is denoted here by $U_k(f)$. When the parameter k is not immediately relevant, we refer to $U_k(f)$ as the uniformity of f.

Our results rely on the following inequality:

$$\text{Correlation}(f, p) \le U_{d+1}(f)^{1/2^{d+1}}, \tag{1}$$

for every function $f : \{0,1\}^n \to \{+1,-1\}$ and every $GF(2)$ polynomial of degree d. Actually, this is true only when the function f satisfies $\Pr_x[f(x) = 1] \approx 1/2$, which is the case in our results (II) and (III) but not in our result (I). For the result (I) we work with a complex-valued version of the Mod_m function for which Inequality (1) also holds. We ignore this issue in this introduction and proceed with the intuition behind our results.

For our result (I), we give an exact calculation of the uniformity of the Mod_m function: $U_{d+1}(Mod_m) = \exp\left(-\Theta\left(n/2^d\right)\right)$. Given some elementary facts about Gowers uniformity, our proof seems simpler and more modular than the proofs in [Bou, GRS]. However, the techniques in [Bou, GRS] generalize to polynomials modulo q for arbitrary q relatively prime to m, as opposed to $q = 2$ in this work. It is not clear to us how to generalize the techniques in this work to any $q \neq 2$.

For our result (II), we note that a random function $F : \{0,1\}^n \to \{+1,-1\}$ satisfies $U_d(F) \le \exp(-\Omega(n))$ with high probability. We derandomize this probabilistic construction by showing that the same holds when the truth-table of F (of length 2^n) is selected at random from a *small-bias space* [NN, AGHP]. Such a sample space F_s can be generated using only

$O(n)$ random bits s, which we can include as part of the input to our function. Thus, we obtain that the function $f(s,x) := F_s(x)$ has correlation at most $\exp\left(-\Omega\left(n/2^d\right)\right)$ with any $GF(2)$ polynomial of degree d. In particular, using a construction in [AGHP], we obtain that this correlation bound holds for the function $(\alpha, \beta, x) \mapsto \langle \alpha^x, \beta \rangle$, where α is an element of $GF(2^n)$ and $\langle \cdot, \cdot \rangle$ denotes inner product modulo 2.

For our result (III), we use a result by Alon et al. [AKK+] that shows that if a function $f : \{0,1\}^n \to \{+1,-1\}$ has correlation at most $1 - 4^{-d}$ with every $GF(2)$ polynomial of degree d (and $\Pr_x[f(x) = 1] \approx 1/2$) then $U_{d+1}(f) \leq 1 - \Omega\left(2^{-d}\right)$. We then observe that the uniformity of the product of functions multiplies.

2 Gowers uniformity

In this section we discuss Gowers uniformity. For our results, we need to work with both real-valued and complex-valued functions. We denote the complex conjugate of a complex number $a+ib$ by $\overline{a+ib} := a - ib$, and its norm by $|a+ib| := \sqrt{a^2+b^2}$. Although the Gowers uniformity of a function is syntactically defined as the expectation of a complex-valued random variable, it is always a non-negative real number (cf. [ST]).

Definition 1 (Gowers uniformity). *Let $f : \{0,1\}^n \to \mathbb{C}$ be a function and $k \geq 1$ an integer. The* k-uniformity *of f is defined as*

$$U_k(f) := E_{y_1, y_2, \ldots, y_k, x \in \{0,1\}^n}\left[\prod_{S \subseteq [k], |S| \text{ even}} f\left(x + \sum_{j \in S} y_j\right) \cdot \prod_{S \subseteq [k], |S| \text{ odd}} \overline{f\left(x + \sum_{j \in S} y_j\right)}\right],$$

where '+' denotes bit-wise XOR.

The following crucial lemma essentially lets us upper bound the correlation of f with *any* low-degree polynomial by the uniformity of f.

Lemma 2 ([GT]). *For every function $f : \{0,1\}^n \to \mathbb{C}$ and every $GF(2)$ polynomial $p : \{0,1\}^n \to \{0,1\}$ of degree at most d,*

$$\left|E_{x \in \{0,1\}^n}\left[f(x) \cdot (-1)^{p(x)}\right]\right| \leq U_{d+1}(f)^{1/2^{d+1}}.$$

Proof sketch of Lemma 2. The lemma follows readily from the following facts, which hold for every function $f : \{0,1\}^n \to \mathbb{C}$:

1 $\left|E_{x \in \{0,1\}^n}[f(x)]\right| = \sqrt{U_1(f)}$,

2 for every k, $U_k(f) \leq \sqrt{U_{k+1}(f)}$,

3 for every $GF(2)$ polynomial p of degree at most d, $U_{d+1}(f \cdot (-1)^p) = U_{d+1}(f)$.

Above, (1) and (2) follow easily from the definition. (3) follows from the fact that for every $GF(2)$ polynomial $p(x)$ of degree d and every fixed $y \in \{0,1\}^n$, the polynomial $q(x) := p(x) + p(x+y)$ has degree $d-1$. The same three facts above are stated in [GT, Equations 1.1, 1.2, and 2.1], for a definition of uniformity that corresponds to our $U_d(f)$ raised to the power of $1/2^d$. □

We make use of the fact, formally stated next, that the uniformity of the product of two functions defined on disjoint input bits is the product of the uniformity of the two functions. The proof of this fact is immediate from the definition.

Fact 3. *For functions $f, f' : \{0,1\}^n \to \mathbb{C}$, define the function $(f \cdot f') : \{0,1\}^{2n} \to \mathbb{C}$ by $(f \cdot f')(x,y) := f(x) \cdot f'(y)$. Then*

$$U_k(f \cdot f') = U_k(f) \cdot U_k(f').$$

3 The correlation of the Mod_m function with GF(2) polynomials

In this section we study the correlation of low-degree $GF(2)$ polynomials with the function $Mod_m : \{0,1\}^n \to \{0,1\}$, for odd $m \geq 3$, where $Mod_m(x_1, x_2, \ldots, x_n)$ equals 1 if and only if $\sum_i x_i$ is divisible by m.

Theorem 4. *For any odd m, the correlation between the Mod_m function on n bits and any $GF(2)$ polynomial of degree d is at most $\exp\left(-\alpha \cdot n/4^d\right)$, where $\alpha > 0$ is a constant that depends on m only.*

Proof. To model the Mod_m function, define $f : \{0,1\}^n \to \mathbb{C}$ as $f(x_1, \ldots, x_n) := e_m\left(\sum_i x_i\right) = \prod_i e_m(x_i)$, where $e_m(x) := e^{2\pi \cdot i \cdot x/m}$. As shown in [Bou], the correlation between a $GF(2)$ polynomial $p(x)$ of degree d and the Mod_m function can be bound from above by the maximum over $a \in \{1, \ldots, m-1\}$ of

$$\left|E_{x \in \{0,1\}^n}\left[f(x)^a \cdot (-1)^{p(x)}\right]\right|, \tag{2}$$

up to a factor $O(m)$ and a term $2^{-\epsilon \cdot n}$ for a constant $\epsilon > 0$ which depends only on m. To bound the above norm (2), we use Lemma 2 to relate it to the $(d+1)$-uniformity of f, and then we use the fact that the uniformity of the product of functions on disjoint input bits multiplies (Fact 3). Formally, letting $k := d+1$, we obtain:

$$\left|E_{x \in \{0,1\}^n}\left[f(x)^a \cdot (-1)^{p(x)}\right]\right| \leq U_k(f^a)^{1/2^k} = U_k(e_m^a)^{n/2^k}.$$

Thus, we are left with the task of bounding

$$U_k(e_m^a) = E_{y_1, \ldots, y_k, x \in \{0,1\}}\left[e_m\left(a \cdot \sum_{S \subseteq [k]} (-1)^{|S|} \cdot \left(x \oplus \left(\bigoplus_{j \in S} y_i\right)\right)\right)\right].$$

To bound $U_k\left(e_m^a\right)$, note that whenever $y_1 = y_2 = \ldots = y_k = 1$, we have that

$$\begin{aligned}
& E_{x\in\{0,1\}}\left[e_m\left(a\cdot\sum_{S\subseteq[k]}(-1)^{|S|}\cdot\left(x\oplus\left(\bigoplus_{j\in S}y_i\right)\right)\right)\right] \\
&= E_{x\in\{0,1\}}\left[e_m\left(a\cdot\sum_{S\subseteq[k]}(-1)^{|S|}\cdot\left(x\oplus\left(\bigoplus_{j\in S}1\right)\right)\right)\right] \\
&= \frac{e_m\left(a\cdot 2^{k-1}\right)+e_m\left(-a\cdot 2^{k-1}\right)}{2} \\
&= \Re\left(e_m\left(a\cdot 2^{k-1}\right)\right) < 1,
\end{aligned}$$

where $\Re(\cdot)$ denotes the real part, and the last inequality holds because m is odd and $a \in \{1,\ldots,m-1\}$. It is also easy to see that the expectation is 0 whenever $y_j = 0$ for some j (though we do not need this for the upper bound). Since it is the case that $y_1 = y_2 = \ldots = y_k = 1$ with probability 2^{-k}, we have, letting $\delta := \Re\left(e_m\left(a\cdot 2^{k-1}\right)\right)$:

$$U_k\left(e_m^a\right) = \left(\delta\cdot 2^{-k} + 1 - 2^{-k}\right).$$

Putting everything together, we obtain

$$\left|E_{x\in\{0,1\}^n}\left[f(x)^a\cdot(-1)^{p(x)}\right]\right| \le \left(1-\frac{1-\delta}{2^k}\right)^{n/2^k} < e^{-(1-\delta)n/2^{2k}},$$

which concludes our proof. (Recall that $\delta < 1$ and that $k = d+1$.) □

4 A function with correlation $\exp\left(-\left(n/2^d\right)\right)$

In this section we exhibit a polynomial-time computable function on n bits that has correlation $\exp\left(-\Omega\left(n/2^d\right)\right)$ with any $GF(2)$ polynomial of degree d. Rather than working directly with the correlation, in the next theorem we exhibit a function $f : \{0,1\}^n \to \{+1,-1\}$ that satisfies $E_x\left[f(x)\cdot(-1)^{p(x)}\right] \le \exp\left(-\Omega\left(n/2^d\right)\right)$ for every $GF(2)$ polynomial p of degree d. This in particular implies that $|\Pr_x[f(x)=1]-0.5| \le \exp\left(-\Omega\left(n/2^d\right)\right)$ (e.g. by considering $p = 0$ and $p = 1$), from which the bound on the correlation follows.

Theorem 5. *There is a polynomial-time computable function $f : \{0,1\}^n \to \{+1,-1\}$ such that for every $d < n/2$ and every $GF(2)$ polynomial $p : \{0,1\}^n \to \{0,1\}$ of degree d we have:*

$$E_x\left[f(x)\cdot(-1)^{p(x)}\right] \le \exp\left(-\alpha\cdot n/2^d\right),$$

where $\alpha > 0$ is a universal constant.

Proof. It is sufficient and more convenient to prove the theorem for a function with input length $O(n)$ rather than n. We prove that the theorem holds for the function that on input (σ, x) equals the x-th output bit of a small-bias generator on seed σ. The following lemma summarizes the definition and the existence of small-bias generators.

Lemma 6 ([NN, AGHP][3]). *There is a polynomial-time computable function $f : \{0,1\}^{O(n)} \times \{0,1\}^n \to \{+1,-1\}$ such that for every $T \subseteq \{0,1\}^n$, we have:*

$$E_\sigma\left[\prod_{x\in T} f(\sigma,x)\right] \le 2^{-n}.$$

Let f be the function in Lemma 6 and write f_σ for the function that maps x to $f(\sigma,x)$. We now show that, over the choice of σ, we expect f_σ to have small uniformity.

Claim 7. $E_\sigma\left[U_k\left(f_\sigma\right)\right] \le 2^{-\alpha\cdot n}$*, for every $k \le n/2$, where $\alpha > 0$ is a universal constant.*

Proof. Let D be the event that for every $S, S' \subseteq [k]$ we have $\sum_{j\in S} y_j \ne \sum_{j\in S'} y_j$ (over the choice of $y_1,\ldots,y_k$). We have:

$$\begin{aligned} E_\sigma\left[U_k\left(f_\sigma\right)\right] &= E_{x,y_1,\ldots,y_k}\left[E_\sigma\left[\prod_{S\subseteq[k]} f_\sigma\left(x+\sum_{j\in S} y_j\right)\right]\right] \\ &\le E_{x,y_1,\ldots,y_k}\left[E_\sigma\left[\prod_{S\subseteq[k]} f_\sigma\left(x+\sum_{j\in S} y_j\right)\right]\middle| D\right] + \Pr[\neg D] \le 2^{-\alpha\cdot n}, \end{aligned}$$

where in the last inequality the first term is at most 2^{-n} by the property of f in Lemma 6, and the second term is

$$\Pr[\neg D] = 1-\left(1-2^{-n}\right)\left(1-2^{-n+1}\right)\cdots\left(1-2^{-n+k-1}\right) \le 1-\left(1-2^{-n+k-1}\right)^{k-1} \le 2^{-\alpha\cdot n}$$

for a universal constant $\alpha > 0$, using that $k \le n/2$. $\square$

To conclude the proof of the theorem, let $p : \{0,1\}^n \to \{0,1\}$ be any $GF(2)$ polynomial of degree d, and notice that

$$\begin{aligned} E_{\sigma,x}\left[f(\sigma,x)\cdot(-1)^{p(\sigma,x)}\right] &= E_\sigma\left[E_x\left[f_\sigma(x)\cdot(-1)^{p(\sigma,x)}\right]\right] \\ &\le E_\sigma\left[U_{d+1}\left(f_\sigma\right)^{1/2^{d+1}}\right] \le E_\sigma\left[U_{d+1}\left(f_\sigma\right)\right]^{1/2^{d+1}} \le 2^{-\alpha\cdot n/2^d}, \end{aligned}$$

where $\alpha > 0$ is a universal constant, the first inequality holds by Lemma 2, the second is Jensen's inequality, and the last holds by Claim 7. $\square$

Remark 8 (On the tightness of Theorem 5). *It is natural to ask whether the* $\exp\left(-\Omega\left(n/2^d\right)\right)$ *correlation bound is tight for the particular function f given by Theorem 5, which recall computes the x-th bit of a small-bias generator, given the seed and x. We observe that this bound is tight in the sense that, for some small-bias generator, the associated function f has correlation $1-o(1)$ with some $GF(2)$ polynomial of degree $d = \log^{O(1)} n$. This follows*

[3] Our presentation is syntactically different from the one in [AGHP], which is in terms of sample spaces. The lemma stated here follows from the results in [AGHP] by considering a small bias sample space over $\{0,1\}^N$, where $N := 2^n$, and defining $f(\alpha,x)$ to be the x-th bit of the sample that corresponds to α.

from the fact that, for some small-bias generator, the associated function f is computable by polynomial-size constant-depth circuits with parity gates [GV, Hea][4] and the well-known fact that any such function has correlation at least $1-o(1)$ with some $GF(2)$ polynomial of degree $\log^{O(1)} n$ [Raz, Smo].

5 An XOR Lemma for low-degree $GF(2)$ polynomials

Yao's XOR lemma [GNW] states that if a function $x \mapsto f(x) \in \{+1,-1\}$ is somewhat hard to compute by small circuits, then the function $(x_1, x_2, \ldots, x_k) \mapsto \prod_{i \leq k} f(x) \in \{+1,-1\}$ is very hard to compute by small circuits, where hardness is measured (up to normalization) by the minimum, over any small circuit, of the fraction of inputs on which the circuit fails to compute the function correctly. Although many proofs of the XOR lemma have been obtained (see, e.g., [GNW]), none of them can be applied to the computational models for which we actually can establish the existence of hard functions (i.e. prove lower bounds), such as constant-depth circuits or low-degree $GF(2)$ polynomials (cf. [Vio2, Chapter 6]). This naturally raises the question of whether one can prove XOR lemmas for these computational models.

In this section we show how Gowers uniformity can be used to obtain the following XOR Lemma for low-degree $GF(2)$ polynomials.

Theorem 9. *Let $f : \{0,1\}^n \to \{+1,-1\}$ be a function that satisfies $E_x\left[f(x) \cdot (-1)^{p(x)}\right] \leq 1 - 4^{-d}$ for every $GF(2)$ polynomial p of degree d. Consider the function $g : \{0,1\}^{n \cdot m} \to \{+1,-1\}$ defined as $g(x_1, x_2, \ldots, x_m) := \prod_{i \leq m} f(x_i)$. Then, for every $GF(2)$ polynomial p of degree d, $E_x\left[g(x) \cdot (-1)^{p(x)}\right] \leq \exp\left(-\Omega\left(m/4^d\right)\right)$.*

To prove Theorem 9 we need to bound the uniformity of a function f such that, for every low-degree polynomial p, we have $E_x\left[f(x) \cdot (-1)^{p(x)}\right] \leq 1 - 4^{-d}$. Such a bound arose from the study of testing of low-degree polynomials. Specifically, Alon et al. [AKK$^+$] define, for a given function $f : \{0,1\}^n \to \{+1,-1\}$, a probabilistic procedure and essentially show that if the function satisfies $E_x\left[f(x) \cdot (-1)^{p(x)}\right] \leq 1 - 4^{-d}$ for every degree-d polynomial p then their procedure rejects with probability $\Omega\left(2^{-d}\right)$. As noted in [Sam], the rejection probability of their procedure is $\left(1 - U_{d+1}(f)\right)/2$. Thus we have the following lemma (stated in [JPRZ, Theorem 4.1] but essentially proved in [AKK$^+$]).

Lemma 10 ([AKK$^+$, JPRZ]). *Let $f : \{0,1\}^n \to \{+1,-1\}$ be a function such that, for every $GF(2)$ polynomial p of degree d, we have $E_x\left[f(x) \cdot (-1)^{p(x)}\right] \leq 1 - 4^{-d}$. Then $U_{d+1}(f) \leq 1 - \Omega\left(2^{-d}\right)$.*

Proof of Theorem 9. We have

$$E_x\left[g(x) \cdot (-1)^{p(x)}\right] \leq U_{d+1}(g)^{1/2^{d+1}} = U_{d+1}(f)^{m/2^{d+1}} \leq \left(1 - \Omega\left(2^{-d}\right)\right)^{m/2^{d+1}} \leq 2^{-\Omega(m/4^d)},$$

[4] There exist simpler non-uniform constructions which would suffice for the point made here, but we do not have a reference for that.

where the first inequality holds by Lemma 2, the next equality by Fact 3, and the next inequality by Lemma 10. $\square$

Theorem 9 can be generalized to the setting where we start with a function that satisfies, for every $GF(2)$ polynomial of degree d, $E_x\left[f(x)\cdot(-1)^{p(x)}\right] \le 1-\epsilon$ for some arbitrary ϵ (as opposed to $\epsilon = 4^{-d}$ in the formulation above). For small $\epsilon \le 4^{-d}$, one can replace the conclusion of Theorem 9 with $E_x\left[g(x)\cdot(-1)^{p(x)}\right] \le \exp\left(-\Omega\left(m\cdot\epsilon\right)\right)$. For big $\epsilon \ge 4^{-d}$, one does not get a bound on $E_x\left[g(x)\cdot(-1)^{p(x)}\right]$ that is better than what stated in Theorem 9 (i.e., $\exp\left(-\Omega\left(m/4^d\right)\right)$), see [JPRZ, Theorem 4.1].

Acknowledgments. We thank Salil Vadhan for his helpful reading of this paper.

References

[AB] N. Alon and R. Beigel. Lower bounds for approximations by low degree polynomials over Z_m. In *Proceedings of the Sixteenth Annual Conference on Computational Complexity*, pages 184–187. IEEE, June 18–21 2001.

[AGHP] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

[AKK$^+$] N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron. Testing low-degree polynomials over GF(2). In *Approximation, randomization, and combinatorial optimization*, volume 2764 of *Lecture Notes in Comput. Sci.*, pages 188–199. Springer, Berlin, 2003.

[Bou] J. Bourgain. Estimation of certain exponential sums arising in complexity theory. *C. R. Math. Acad. Sci. Paris*, 340(9):627–631, 2005.

[CT] F. R. K. Chung and P. Tetali. Communication complexity and quasi randomness. *SIAM J. Discrete Math.*, 6(1):110–123, 1993.

[GNW] O. Goldreich, N. Nisan, and A. Wigderson. On Yao's XOR lemma. Technical Report TR95–050, Electronic Colloquium on Computational Complexity, March 1995. `http://www.eccc.uni-trier.de/eccc`.

[Gow1] W. T. Gowers. A new proof of Szemerédi's theorem for arithmetic progressions of length four. *Geom. Funct. Anal.*, 8(3):529–551, 1998.

[Gow2] W. T. Gowers. A new proof of Szemerédi's theorem. *Geom. Funct. Anal.*, 11(3):465–588, 2001.

[GT] B. Green and T. Tao. An inverse theorem for the Gowers U^3 norm, 2005. arXiv.org:math/0503014.

[Gre] F. Green. The correlation between parity and quadratic polynomials mod 3. *J. Comput. System Sci.*, 69(1):28–44, 2004.

[GRS] F. Green, A. Roy, and H. Straubing. Bounds on an exponential sum arising in Boolean circuit complexity. *C. R. Math. Acad. Sci. Paris*, 341(5):279–282, 2005.

[GV] D. Gutfreund and E. Viola. Fooling Parity Tests with Parity Gates. In *Proceedings of the Eight International Workshop on Randomization and Computation (RANDOM)*, Lecture Notes in Computer Science, Volume 3122, pages 381–392. Springer-Verlag, August 22–24 2004.

[HMP+] A. Hajnal, W. Maass, P. Pudlák, M. Szegedy, and G. Turán. Threshold circuits of bounded depth. *J. Comput. System Sci.*, 46(2):129–154, 1993.

[Han] K. A. Hansen. Lower Bounds for Circuits with Few Modular Gates using Exponential Sums. *Electronic Colloquium on Computational Complexity*, Technical Report TR06-079, 2006. `http://www.eccc.uni-trier.de/eccc`.

[HG] J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Comput. Complexity*, 1(2):113–129, 1991.

[Hea] A. Healy. Randomness-Efficient Sampling within NC^1. In *Proceedings of the 10th International Workshop on Randomization and Computation (RANDOM)*, Lecture Notes in Computer Science., 2006.

[JPRZ] C. S. Jutla, A. C. Patthak, A. Rudra, and D. Zuckerman. Testing Low-Degree Polynomials over Prime Fields. *focs*, 00:423–432, 2004.

[NN] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. In *Proceedings of the 22nd Annual ACM Symposium on the Theory of Computing*, pages 213–223, 1990.

[Raz] A. A. Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Mat. Zametki*, 41(4):598–607, 623, 1987.

[Sam] A. Samorodnitsky. Low degree tests at large distances, 2006. Manuscript.

[ST] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influence of variables, and PCPs. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing 2006, Seattle, WA, USA, May 21–23, 2006*, pages 11–20, 2006.

[Smo] R. Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, pages 77–82, New York City, 25–27 May 1987.

[Vio1] E. Viola. Pseudorandom Bits for Constant-Depth Circuits with Few Arbitrary Symmetric Gates. In *Proceedings of the Twentieth Annual Conference on Computational Complexity*, pages 198–209. IEEE, June 12–15 2005. To appear in *SIAM Journal of Computing*.

[Vio2] E. Viola. *The Complexity of Hardness Amplification and Derandomization*. PhD thesis, Harvard University, 2006.