



A Hypergraph Dictatorship Test with Perfect Completeness

Victor Chen *

November 12, 2008

Abstract

A hypergraph dictatorship test is first introduced by Samorodnitsky and Trevisan in [21] and serves as a key component in their unique games based PCP construction. Such a test has oracle access to a collection of functions and determines whether all the functions are the same dictatorship, or all their low degree influences are $o(1)$. The test in [21] makes $q \geq 3$ queries and has amortized query complexity $1 + O\left(\frac{\log q}{q}\right)$ but has an inherent loss of perfect completeness. In this note, we give another hypergraph dictatorship test that achieves both perfect completeness and amortized query complexity $1 + O\left(\frac{\log q}{q}\right)$.

1 Introduction

Linearity and dictatorship testing have been studied in the past decade both for their combinatorial interest and connection to complexity theory. These tests distinguish functions which are linear/dictator from those which are far from being a linear/dictator function. The tests do so by making queries to a function at certain points and receiving the function's values at these points. The parameters of interest are the number of queries a test makes and the completeness and soundness of a test.

In this paper we shall work with boolean functions of the form $f : \{0, 1\}^n \rightarrow \{-1, 1\}$. We say a function f is *linear* if $f = (-1)^{\sum_{i \in S} x_i}$ for some subset $S \subset [n]$. A *dictator* function is simply a linear function where $|S| = 1$, i.e., $f(x) = (-1)^{x_i}$ for some i . A dictator function is often called a *long code*, and it is first used in [4] for the constructions of probabilistic checkable proofs (PCPs), see e.g., [2, 1]. Since then, it has become standard to design a PCP verifier as the composition of two parts, an outer verifier and an inner verifier. In such case, a PCP verifier expects the proof to be written in such a way so that the outer verifier, typically based on the verifier obtained from Raz's Parallel Repetition Theorem [17], selects some tables of the proof at random and then passes the control to the inner verifier. The inner verifier, with oracle access to these tables, makes queries into these tables and ensures that the tables are the encoding of some error-correcting codes and satisfy some joint constraint. Since these table typically have constant size, efficiency is not an issue. The long code encoding is usually employed in these proof constructions, and the inner verifier simply tests whether the collection of tables (functions) are long codes satisfying some constraints. In this way, constructing a PCP with certain parameters reduces to the problem of designing a long code test with similar parameters.

One question of interest is the tradeoff between the soundness and query complexity of a tester. If a tester queries the functions at every single value, then trivially the verifier can determine all the functions. One would like to construct a dictatorship test that has the lowest possible soundness while making as few

*MIT CSAIL. victor@csail.mit.edu. Research supported in part by an NSF graduate fellowship and NSF Award CCR-0514915.

queries as possible. One way to measure this tradeoff between the soundness s and the number of queries q is *amortized query complexity*, defined as $\frac{q}{\log s - 1}$. This line of research, initiated in [25], has since spurred a long sequence of works [22, 20, 12, 7]. All these tests essentially find a way to reuse the queries by running dependent iterations of a basic linear/dictatorship test. The techniques used are Fourier analytic, and the best amortized query complexity from this sequence of works has the form $1 + O\left(\frac{1}{\sqrt{q}}\right)$.

More progress is made in [19] when the notion of a *relaxed* linearity test is introduced along with new ideas from additive combinatorics. In property testing, the goal is to distinguish objects that are very structured from those that are pseudorandom. In the case of linearity/dictatorship testing, the structured objects are the linear/dictator functions, and functions that are far from being linear/dictator are interpreted as pseudorandom. The recent paradigm in additive combinatorics is to find the right framework of structure and pseudorandomness and analyze combinatorial objects by dividing them into structured and pseudorandom components, see e.g., [24] for a survey. One success is the notion of Gowers norm [8], which has been fruitful in attacking many problems in additive combinatorics and computer science. In [19], the notion of pseudorandomness for linearity testing is relaxed; instead of designating the functions that are far from being linear as pseudorandom, the functions having small low degree Gowers norm are considered to be pseudorandom. By doing so, an optimal tradeoff between soundness and query complexity is obtained for the problem of relaxed linearity testing. (Here the tradeoff is stronger than the tradeoff for the traditional problem of linearity testing.)

In a similar fashion, in the PCP literature since [10], the pseudorandom objects in dictatorship tests are not functions that are far from being a dictator. The pseudorandom functions are typically defined to be either functions that are far from all “juntas” or functions whose “low-degree influences” are $o(1)$. Both considerations of a dictatorship test are sufficient to compose the test in a PCP construction. In [21], building on the analysis of the relaxed linearity test in [19], a dictatorship test (taking the view that functions with arbitrary small “low-degree influences are pseudorandom”) is constructed with amortized query complexity $1 + O\left(\frac{\log q}{q}\right)$. Furthermore, the test is used as the inner verifier in a conditional PCP construction (based on unique games [13]) with the same parameters. However, their dictatorship test suffers from an inherent loss of perfect completeness. Ideally one would like testers with one-sided errors. One, for aesthetic reasons, testers should always accept valid inputs. Two, for some hardness of approximation applications, in particular coloring problems (see e.g., [11], [6]), it is important to construct PCP systems with one-sided errors.

In this paper, we prove the following theorem:

Theorem 1.1 *For every $q \geq 3$, there exists an (adaptive) dictatorship test that makes q queries, has completeness 1, and soundness $\frac{O(q^2)}{2^q}$; in particular it has amortized query complexity $1 + O\left(\frac{\log q}{q}\right)$.*

Our tester is a variant of the one given in [21]. Our tester is adaptive in the sense that it makes its queries in two stages. It first makes roughly $\log q$ nonadaptive queries into the function. Based on the values of these queries, the tester then selects the rest of the query points nonadaptively. Our analysis is based on techniques developed in [12, 21, 11, 9].

1.1 Future Direction

Unfortunately, the adaptivity of our test is a drawback. The correspondence between PCP constructions and hardness of approximation needs the test to be fully nonadaptive. However, a more pressing issue is that our hypergraph dictatorship test does not immediately imply a new PCP characterization of NP. The reason is that a dictatorship test without “consistency checks” is most easily composed with the unique label

cover defined in [13] as the outer verifier in a PCP reduction. As the conjectured NP-hardness of the unique label cover cannot have perfect completeness, the obvious approach in combining our test with the unique games-based outer verifier does not imply a new PCP result. However, there are variants of the unique label cover (e.g., Khot’s 2 to 1 Conjecture) [13] that do have conjectured perfect completeness, and these variants are used to derive hardness of coloring problems in [6]. We hope that our result combined with similar techniques used in [6] may obtain a new conditional PCP construction and will motivate more progress on the unique games conjecture.

1.2 Related Works

The problem of linearity testing was first introduced in [5]. The framework of property testing was formally set up in [18]. The PCP Theorems were first proved in [2, 1]; dictatorship tests first appeared in the PCP context in [4], and many dictatorship tests and variants appeared throughout the PCP literature. Dictatorship test was also investigated in a general, combinatorial property testing setting in [16]. As mentioned, designing testers and PCPs focusing on amortized query complexity was first investigated in [25], and a long sequence of works [22, 20, 12, 7] followed. The first tester/PCP system focusing on this tradeoff while obtaining perfect completeness was achieved in [11].

The orthogonal question of designing testers or PCPs with as few queries as possible was also considered. In a highly influential paper [10], Håstad constructed a PCP system making only three queries. Many variants also followed. In particular PCP systems with perfect completeness making three queries were also achieved in [9, 14], and similar to our approach, O’Donnell and Wu designed an optimal three bit dictatorship test with perfect completeness in [15] as a step toward constructing a conditional PCP system.

2 Preliminaries

We fix some notation and provide the necessary background in this section. We let $[n]$ denote the set $\{1, 2, \dots, n\}$. For a vector $v \in \{0, 1\}^n$, we write $|v| = \sum_{i \in [n]} v_i$. We let \wedge denote the boolean AND, where $a \wedge b = 1$ iff $a = b = 1$. For vectors $v, w \in \{0, 1\}^n$, we write $v \wedge w$ to denote the vector obtained by applying AND to v and w component-wise. We abuse notation and sometimes interpret a vector $v \in \{0, 1\}^n$ as a subset $v \subset [n]$ where $i \in v$ iff $v_i = 1$. For a boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we make the convenient notational change from $\{0, 1\}$ to $\{-1, 1\}$ and write $f : \{0, 1\}^n \rightarrow \{-1, 1\}$.

2.1 Fourier Analysis

Definition 2.1 (*Fourier transform*) For a real-valued function $f : \{0, 1\}^n \rightarrow \mathbb{R}$, we define its Fourier transform $\hat{f} : \{0, 1\}^n \rightarrow \mathbb{R}$ to be

$$\hat{f}(\alpha) = \mathbb{E}_{x \in \{0, 1\}^n} f(x) \chi_\alpha(x),$$

where $\chi_\alpha(x) = (-1)^{\sum_{i \in [n]} \alpha_i x_i}$. We say $\hat{f}(\alpha)$ is the Fourier coefficient of f at α , and the characters of $\{0, 1\}^n$ are the functions $\{\chi_\alpha\}_{\alpha \in \{0, 1\}^n}$.

It is easy to see that for $\alpha, \beta \in \{0, 1\}^n$, $\mathbb{E} \chi_\alpha \cdot \chi_\beta$ is 1 if $\alpha = \beta$ and 0 otherwise. Since there are 2^n characters, they form an orthonormal basis for $\{0, 1\}^n$, and we have the Fourier inversion formula

$$f(x) = \sum_{\alpha \in \{0, 1\}^n} \hat{f}(\alpha) \chi_\alpha(x)$$

and Parseval's Identity

$$\sum_{\alpha \in \{0,1\}^n} \hat{f}(\alpha)^2 = \mathbb{E}_x[f(x)^2].$$

2.2 Influence of Variables

For a boolean function $f : \{0,1\}^n \rightarrow \{-1,1\}$, the *influence* of the i -variable, $I_i(f)$, is defined to be $\Pr_{x \in \{0,1\}^n}[f(x) \neq f(x + e_i)]$, where e_i is a vector in $\{0,1\}^n$ with 1 on the i -th coordinate 0 everywhere else. This corresponds to our intuitive notion of influence: how likely the outcome of f changes when the i -th variable on a random input is flipped. For the rest of this paper, it will be convenient to work with the Fourier analytic definition of $I_i(f)$ instead, and we leave it to the readers to verify that the two definitions are equivalent when f is a boolean function.

Definition 2.2 Let $f : \{0,1\}^n \rightarrow \mathbb{R}$. We define the influence of the i -th variable of f to be

$$I_i(f) = \sum_{\alpha \in \{0,1\}^n: \alpha_i=1} \hat{f}(\alpha)^2.$$

We shall need the following technical lemma that gives an upper bound on the influence of a product of functions.

Lemma 2.3 (Lemma 4 from [21]) Let $f_1, \dots, f_k : \{0,1\}^n \rightarrow [-1,1]$ be a collection of k bounded real-valued functions, and define $f(x) = \prod_{i=1}^k f_i(x)$ to be the product of these k functions. Then for each $i \in [n]$,

$$I_i(f) \leq k \cdot \sum_{j=1}^k I_i(f_j).$$

When $\{f_i\}$ are boolean functions, it is easy to see that $I_i(f) \leq \sum_{j=1}^k I_i(f_j)$ by the union bound.

We now define the notion of *low-degree influence*.

Definition 2.4 Let w be an integer between 0 and n . We define the w -th degree influence of the i -th variable of a function $f : \{0,1\}^n \rightarrow \mathbb{R}$ to be

$$I_i^{\leq w}(f) = \sum_{\alpha \in \{0,1\}^n: \alpha_i=1, |\alpha| \leq w} \hat{f}(\alpha)^2.$$

While the definition of low-degree influence is standard in the literature, we shall make a few remarks since this definition does not have a clean combinatorial interpretation or an immediate justification. Dictatorship tests (those based on influences) classify functions in the NO instances to be those whose low-degree influences are $o(1)$ for two reasons. One is that large parity functions, which have many variables with influence 1 but no variables with low-degree influence, must be rejected by the test. The second is that if w is fixed, then a bounded function has only a finite number of variables with large w -th degree influence. This easy fact, though we won't need it here, is often needed to lift a dictatorship test to a PCP construction. Both such considerations fail if we substitute the low-degree influence requirement by just influence, thus the need for a thresholded version of influence.

2.3 Gowers norm

In [8], Gowers uses analytic techniques to give a new proof of Szemerédi’s Theorem [23] and in particular, initiates the study of a new norm of a function as a measure of pseudorandomness. Subsequently this norm is termed the *Gowers uniformity norm* and has been intensively studied and applied in additive combinatorics, see e.g. [24] for a survey. The use of the Gowers norm in computer science is initiated in [19, 21].

Definition 2.5 Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$. We define the d -th dimension Gowers uniformity norm of f to be

$$\|f\|_{U_d} = \left(\mathbb{E}_{x, x_1, \dots, x_d} \left[\prod_{S \subset [d]} f \left(x + \sum_{i \in S} x_i \right) \right] \right)^{1/2^d}.$$

For a collection of 2^d functions $f_S : \{0, 1\}^n \rightarrow \mathbb{R}, S \subset [d]$, we define the d -th dimension Gowers inner product of $\{f_S\}_{S \subset [d]}$ to be

$$\langle \{f_S\}_{S \subset [d]} \rangle_{U_d} = \mathbb{E}_{x, x_1, \dots, x_d} \left[\prod_{S \subset [d]} f_S \left(x + \sum_{i \in S} x_i \right) \right].$$

When f is a boolean function, one can interpret the Gowers norm as simply the expected number of “affine parallelepiped” of dimension d . While this expression may look cumbersome at first glance, the use of the Gowers norm is in some sense to control expectations over some other expressions. For instance, to count the number of $d+1$ -term progressions of the form $x, x+y, \dots, x+d \cdot y$ in a subset, one may be interested in approximating expressions of the form $\mathbb{E}_{x,y}[f_1(x)f_2(x+y) \cdots f_d(x+d \cdot y)]$, where f_1, \dots, f_d are some bounded functions over some appropriate domain. In fact, as shown by Gowers, these expectations are upper bounded by the Gowers inner product of f_i , which is also upper bounded by $\min_{i \in [d]} \|f_i\|_{U_d}^{2^d}$. Thus, in a rough sense, questions regarding progressions are then reduced to questions regarding the Gowers norms, which are more amenable to analytic techniques.

The proof showing that $\mathbb{E}_{x,y}[f_1(x)f_2(x+y) \cdots f_d(x+d \cdot y)]$ is upper bounded by the minimum Gowers norm of all the functions f_i is not difficult; it proceeds by repeated applications of the Cauchy-Schwarz inequality and substitution of variables. Collectively, statements saying that certain expressions are governed by the Gowers norm are coined *von-Neumann type theorems* in the literature.

For the analysis of hypergraph-based dictatorship test, we shall encounter the following expression.

Definition 2.6 Let $\{f_S\}_{S \subset [d]}$ be a collection of functions where $f_S : \{0, 1\}^n \rightarrow \mathbb{R}$. We define the d -th dimension Gowers linear inner product of $\{f_S\}$ to be

$$\langle \{f_S\}_{S \subset [d]} \rangle_{LU_d} = \mathbb{E}_{x_1, \dots, x_d} \left[\prod_{S \subset [d]} f_S \left(\sum_{i \in S} x_i \right) \right].$$

This definition is a variant of the Gowers inner product and is in fact upper bounded by the square root of the Gowers inner product as shown in [21]. Furthermore they showed that if a collection of functions has large Gowers inner product, then two functions must share an influential variable. Thus, one can infer the weaker statement that large linear Gowers inner product implies two functions have an influential variable. For our purposes, we can encapsulate all the prior discussion into the following statement, which is Lemma 16 from [21]. This is the only fact on the Gowers norm that we explicitly need.

Lemma 2.7 (Lemma 16 from [21]) Let $\{f_S\}_{S \subset [d]}$ be a collection of bounded functions of the form $f_S : \{0, 1\}^n \rightarrow [-1, 1]$. Suppose $LU_d(\{f_S\}_{S \subset [d]}) \geq \epsilon$ and $\mathbb{E} f_{[d]} = 0$. Then there exists some variable i , some subsets $S \neq T \subset [d]$ such that the influences of the i -th variable in both f_S and f_T are at least $\frac{\epsilon^4}{2^{O(d)}}$.

3 Dictatorship Test

Definition 3.1 (*dictatorship*) For $i \in [n]$, the i -th dictator is the function $f(x) = (-1)^{x_i}$.

Alternatively, it is often referred to as the long code encoding of i , $\langle (-1)^{x_i} \rangle_{x \in \{0,1\}^n}$, which is simply the evaluation of the i -th dictator function at all points. However, for most applications of long code, n is a large constant.

Now let us define a t -function dictatorship test. Suppose we are given oracle access to a collection of boolean functions f_1, \dots, f_t . We want to make as few queries as possible into these functions to decide if all the functions are the same dictatorship, or no two functions have some common structure. More precisely, we have the following definition:

Definition 3.2 We say that a test $T = T^{f_1, \dots, f_t}$ is a t -function dictatorship test with completeness c and soundness s if T is given oracle access to a family of t functions $f_1, \dots, f_t : \{0, 1\}^n \rightarrow \{-1, 1\}$, such that

- if there exists some variable $i \in [n]$ such that for all $a \in [t]$, $f_a(x) = (-1)^{x_i}$, then T accepts with probability at least c , and
- for every $\epsilon > 0$, there exist a positive constant $\tau > 0$ and a fixed positive integer w such that if T accepts with probability at least $s + \epsilon$, then there exist two functions f_a, f_b where $a, b \in [t]$, $a \neq b$ and some variable $i \in [n]$ such that $I_i^{\leq w}(f_a), I_i^{\leq w}(f_b) \geq \tau$.

A q -function dictatorship test making q queries, with soundness $\frac{q+1}{2^q}$ was proved in [21], but the test suffers from imperfect completeness. We obtain a $q - O(\log q)$ -dictatorship test that makes q queries, has completeness 1, soundness $\frac{O(q^2)}{2^q}$, and in particular has amortized query complexity $1 + O\left(\frac{\log q}{q}\right)$, the same as the test in [21]. By a simple change of variable, we can more precisely state the following:

Theorem 3.3 (main theorem restated) For infinitely many t , there exists an adaptive t -function dictatorship test that makes $t + \log(t + 1)$ queries, has completeness 1, and soundness $\frac{t+1}{2^t}$.

Our test is adaptive and selects queries in two passes. During the first pass, it picks an arbitrary subset of $\log(t + 1)$ functions out of the t functions. For each function selected, our test picks a random entry y and queries the function at entry y . Then based on the values of these $\log(t + 1)$ queries, during the second pass, the test selects t positions nonadaptively from each function, then queries all t positions at once. The adaptivity is necessary in our analysis, and it is unclear if one can prove an analogous result with only one pass.

3.1 Folding

As introduced in [3], we shall assume that the functions are folded. We do so by requiring our dictatorship test to make queries in a special manner. Suppose the test wants to query f at the point $x \in \{0, 1\}^n$. If $x_1 = 1$, then the test queries $f(x)$ as usual. If $x_1 = 0$, then the test queries f at the point $\vec{1} + x = (1, 1 + x_2, \dots, 1 + x_n)$ and negates the value it receives. It is instructive to note that folding ensures $f(\vec{1} + x) = -f(x)$ and $f(\vec{0}) = 0$.

3.2 Basic Test

For ease of exposition, we first consider the following simplistic scenario. Suppose we have oracle access to just one boolean function. Furthermore we ignore the tradeoff between soundness and query complexity. We simply want a dictatorship test that always accepts dictator functions and rejects all functions with small low-degree influence with probability at least $\frac{1}{2} + \epsilon, \epsilon > 0$. There are many such tests in the literature; however, we need a suitable one which our hypergraph dictatorship test can base on. Our basic test below is a close variant of the one considered in [9].

BASIC TEST T : with oracle access to f ,

1. Pick x_i, x_j, y, z uniformly at random from $\{0, 1\}^n$.
2. Query $f(y)$.
3. Let $w = \frac{1-f(y)}{2}$. Accept iff

$$f(x_i)f(x_j) = f(x_i + x_j + ((w\vec{1} + y) \wedge z)).$$

Lemma 3.4 *The test T is a dictatorship test with completeness 1.*

Proof: Suppose f is the ℓ -th dictator, i.e., $f(x) = (-1)^{x_\ell}$. First note that

$$w + y_\ell = \frac{1 - (-1)^{y_\ell}}{2} + y_\ell,$$

which evaluates to 0. Thus by linearity of f

$$\begin{aligned} f(x_i + x_j + (w\vec{1} + y) \wedge z) &= f(x_i)f(x_j)f((w\vec{1} + y) \wedge z) \\ &= f(x_i)f(x_j)(-1)^{(w+y_\ell)\wedge z_\ell} \\ &= f(x_i)f(x_j) \end{aligned}$$

and the test always accepts. ■

To analyze the soundness of the test T , we need to derive a Fourier analytic expression for the acceptance probability of T .

Proposition 3.5 *Let p be the acceptance probability of T . Then*

$$p = \frac{1}{2} + \frac{1}{2} \sum_{\alpha \in \{0,1\}^n} \widehat{f}(\alpha)^3 2^{-|\alpha|} \left(1 + \sum_{\beta \leq \alpha} \widehat{f}(\beta) \right).$$

For sanity check, let us interpret the expression for p . Suppose $f = \chi_\alpha$ for some $\alpha \neq \vec{0} \in \{0, 1\}^n$, i.e., $\widehat{f}(\alpha) = 1$ and all other Fourier coefficients of f are 0. Then clearly $p = \frac{1}{2} + 2^{-|\alpha|}$, which equals 1 whenever f is a dictator function as we have just shown. If $|\alpha|$ is large, then T accepts with probability close to $\frac{1}{2}$. Furthermore, as f is folded, $\widehat{f}(\vec{0}) = 0$. So if f is a constant function, then $p = \frac{1}{2}$. We shall first analyze the soundness and then derive this analytic expression for p .

Lemma 3.6 *The test T is a dictatorship test with soundness $\frac{1}{2}$.*

Proof: Suppose the test T passes with probability at least $\frac{1}{2} + \epsilon$, for some $\epsilon > 0$. By Proposition 3.5 and applying Parseval's Identity and Cauchy-Schwarz, we obtain

$$\begin{aligned} \epsilon &\leq \frac{1}{2} \sum_{\alpha \in \{0,1\}^n} \widehat{f}(\alpha)^3 2^{-|\alpha|} \left(1 + \sum_{\beta \leq \alpha} \widehat{f}(\beta) \right) \\ &\leq \sum_{\alpha \in \{0,1\}^n} \widehat{f}(\alpha)^3 2^{-\frac{|\alpha|}{2}}. \end{aligned}$$

Pick a positive integer w such that $2^{-\frac{w}{2}} \leq \frac{\epsilon}{2}$. Then by Parseval's again,

$$\begin{aligned} \frac{\epsilon}{2} &\leq \sum_{\alpha \in \{0,1\}^n: |\alpha| \leq w} \widehat{f}(\alpha)^3 \\ &\leq \max_{\alpha \in \{0,1\}^n: |\alpha| \leq w} |\widehat{f}(\alpha)|. \end{aligned}$$

So there exists some $\beta \in \{0,1\}^n$, $|\beta| \leq w$ such that $\frac{\epsilon}{2} \leq |\widehat{f}(\beta)|$. With f being folded, $\beta \neq \vec{0}$. Thus, there exists an $i \in [n]$ such that $\beta_i = 1$ and

$$\frac{\epsilon^2}{4} \leq |\widehat{f}(\beta)|^2 \leq \sum_{\alpha \in \{0,1\}^n: \alpha_i=1, |\alpha| \leq w} \widehat{f}(\alpha)^2.$$

■

Now we give the straightforward Fourier analytic calculation for p .

Proof (of Proposition 3.5): As usual, we first arithmetize p . We write

$$\begin{aligned} p &= \mathbb{E}_{x_i, x_j, y, z} \left(\frac{1 + f(y)}{2} \right) \left(\frac{1 + \text{Acc}(x_i, x_j, y, z)}{2} \right) + \\ &\quad \mathbb{E}_{x_i, x_j, y, z} \left(\frac{1 - f(y)}{2} \right) \left(\frac{1 + \text{Acc}(x_i, x_j, \vec{1} + y, z)}{2} \right), \end{aligned}$$

where

$$\text{Acc}(x_i, x_j, y, z) = f(x_i)f(x_j)f(x_i + x_j + (y \wedge z)).$$

Since f is folded, $f(\vec{1} + y) = -f(y)$. As y and $\vec{1} + y$ are both identically distributed in $\{0,1\}^n$, we have

$$p = 2 \mathbb{E}_{x_i, x_j, y, z} \left(\frac{1 + f(y)}{2} \right) \left(\frac{1 + \text{Acc}(x_i, x_j, y, z)}{2} \right).$$

Since $\mathbb{E} f = 0$, we can further simplify the above expression to be

$$p = \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x_i, x_j, y, z} [(1 + f(y)) \text{Acc}(x_i, x_j, y, z)].$$

It suffices to expand out the terms $\mathbb{E}_{x_i, x_j, y, z} [\text{Acc}(x_i, x_j, y, z)]$ and $\mathbb{E}_{x_i, x_j, y, z} [f(y) \text{Acc}(x_i, x_j, y, z)]$.

For the first term, it is not hard to show that

$$\mathbb{E}_{x_i, x_j, y, z} [\text{Acc}(x_i, x_j, y, z)] = \sum_{\alpha \in \{0,1\}^n} \widehat{f}(\alpha)^3 2^{-|\alpha|},$$

by applying the Fourier inversion formula on f and averaging over x_i and x_j and then averaging over y and z over the AND operator.

Now we compute the second term. Applying the Fourier inversion formula to the last three occurrences of f and averaging over x_i and x_j , we obtain

$$\mathbb{E}_{x_i, x_j, y, z} [f(y) \text{Acc}(x_i, x_j, y, z)] = \sum_{\alpha \in \{0,1\}^n} \widehat{f}(\alpha)^3 \mathbb{E}_{y, z} [f(y) \chi_\alpha(y \wedge z)].$$

It suffices to expand out $\mathbb{E}_{y, z} [f(y) \chi_\alpha(y \wedge z)]$. By grouping the z 's according to their intersection with different possible subsets β of α , we have

$$\begin{aligned} & \mathbb{E}_{y, z} [f(y) \chi_\alpha(y \wedge z)] \\ &= \sum_{\beta \subset \alpha} \Pr_{z \in \{0,1\}^n} [z \cap \alpha = \beta] \mathbb{E}_y \left[f(y) \prod_{i: \alpha_i=1} (-1)^{y_i \wedge z_i} \right] \\ &= \sum_{\beta \subset \alpha} 2^{-|\alpha|} \mathbb{E}_y \left[f(y) \prod_{i: \beta_i=1} (-1)^{y_i} \right] \\ &= 2^{-|\alpha|} \sum_{\beta \subset \alpha} \widehat{f}(\beta). \end{aligned}$$

Putting everything together, it is easy to see that we have the Fourier analytic expression for p as stated in the lemma. ■

3.3 Hypergraph Dictatorship Test

We prove the main theorem in this section. The basis of our hypergraph dictatorship test will be very similar to the test in the previous section. We remark that we did not choose to present the exact same basic test for hopefully a clearer exposition.

We now address the tradeoff between query complexity and soundness. If we simply repeat the basic test a number of iterations independently, the error is reduced, but the query complexity increases. In other words, the amortized query complexity does not change if we simply run the basic test for many independent iterations. Following [25], all the dictatorship tests that save query complexity do so by reusing queries made in previous iterations of the basic test. To give an illustration, the test T might query f at the points $x_1 + (y \wedge z_1)$, $x_2 + (y \wedge z_2)$, $x_1 + x_2 + (y \wedge z_{12})$ to make a decision. For the second iteration, we let T query f at the points $x_3 + (y \wedge z_3)$ and $x_1 + x_3 + (y \wedge z_{13})$ and reuse the value $f(x_1 + (y \wedge z_1))$ queried during the first run of T . T then uses the three values to make a second decision. In total T makes five queries to run two iterations.

We may think of the first run of T as parameterized by the points x_1 and x_2 and the second run of T by x_1 and x_3 . In general, we may have k points x_1, \dots, x_k and a graph on $[k]$ vertices, such that each edge e of the graph corresponds to an iteration of T parameterized by the points $\{x_i\}_{i \in e}$. We shall use a complete hypergraph on k vertices to save on query complexity, and we will argue that the soundness of the algorithm decreases exponentially with respect to the number of iterations.

Formally, consider a hypergraph $H = ([k], E)$. Let $\{f_a\}_{a \in [k] \cup E}$ be a collection of boolean functions of the form $f_a : \{0, 1\}^n \rightarrow \{-1, 1\}$. We assume all the functions are folded, and so in particular, $\mathbb{E} f = 0$. Consider the following test:

HYPERGRAPH H -TEST: with oracle access to $\{f_a\}_{a \in [k] \cup E}$,

1. Pick $x_1, \dots, x_k, y_1, \dots, y_k$, and $\{z_a\}_{a \in [k] \cup E}$ independently and uniformly at random from $\{0, 1\}^n$.
2. For each $i \in [k]$, query $f_i(y_i)$.
3. Let $w_i = \frac{1 - f_i(y_i)}{2}$.
Accept iff for every $e \in E$,

$$\prod_{i \in e} \left[f_i(x_i + (w_i \vec{1} + y_i) \wedge z_i) \right] = f_e \left(\sum_{i \in e} x_i + \left(\sum_{i \in e} (w_i \vec{1} + y_i) \right) \wedge z_e \right).$$

We make some remarks regarding the design of H -Test. The hypergraph test in [21] accepts iff for every $e \in E$, $\prod_{i \in e} f_i(x_i + \eta_i)$ equals $f_e(\sum_{i \in e} x_i + \eta_e)$, where the bits in each vector η_a are chosen independently to be 1 with some small constant, say 0.01. The noise vectors η_a rules out the possibility that linear functions with large support can be accepted. To obtain a test with perfect completeness, we use ideas from [9, 16, 11] to simulate the effect of the noise perturbation.

Note that for y, z chosen uniformly at random from $\{0, 1\}^n$, the vector $y \wedge z$ is a $\frac{1}{4}$ -noisy vector. As observed in [16] the test $f(y \wedge z) = f(y) \wedge f(z)$ distinguishes between dictators and linear functions with large support. As shown in [11], one can also combine linearity and dictatorship testing into a single test that checks if $f(x_1 + x_2 + y \wedge z)(f(y) \wedge f(z))$ equals $f(x_1)f(x_2)$. However, iterating such a test increases the number of queries made. In fact, the graph test in [11] has completeness 1 and reads $k^2 + 4k$ bits to obtain a soundness of 2^{-k^2} . The authors in [11] only achieve a test that queries the same number of bits as in [20] ($k^2 + 2k$ bits) when they allow the test to be adaptive. While both the nonadaptive and adaptive tests in [11] have the same amortized query complexity, extending the nonadaptive test in [11] to the hypergraph setting is too costly for us. So to achieve the same amortized query complexity as the hypergraph test in [21], we also exploit adaptivity in our test.

Similar to the tests in [9, 11], our test has two stages, where queries made during the second stage depends on the answers received during the first stage. Essentially the queries from the first stage determine whether the vector y_i or its complement $\vec{1} + y_i$ should be used during the second stage. Of interest is the number of queries made during the first stage. While a single random vector y and the corresponding adaptive query suffices for the analysis of a single edge test, it is unclear how to prove this works for a general hypergraph. To circumvent this we select k vectors $\{y_i\}$ to match them up with the vectors $\{x_i\}$ to facilitate the analysis of the Gowers norm.

Theorem 3.7 (main theorem restated) *For infinitely many t , there exists an adaptive t -function dictatorship test with $t + \log(t + 1)$ queries, completeness 1, and soundness $\frac{t+1}{2^t}$.*

Proof: Take a complete hypergraph on k vertices, where $k = \log(t + 1)$. The statement follows by applying Lemmas 3.8 and 3.9. ■

Lemma 3.8 *The H-Test is a $(k + |E|)$ -function dictatorship test that makes $|E| + 2k$ queries and has completeness 1.*

Proof: The test makes k queries $f_i(y_i)$ in the first pass, and based on the answers to these k queries, the test then makes one query into each function f_a , for each $a \in [k] \cup E$. So the total number of queries is $|E| + 2k$. Now suppose all the functions are the ℓ -th dictator for some $\ell \in [n]$, i.e., for all $a \in [k] \cup E$, $f_a = f$, where $f(x) = (-1)^{x_\ell}$. Note that for each $i \in [k]$,

$$w_i + y_i(\ell) = \frac{1 - (-1)^{y_i(\ell)}}{2} + y_i(\ell),$$

which evaluates to 0. Thus for each $e \in E$,

$$\begin{aligned} \prod_{i \in e} f_i(x_i + (w_i \vec{1} + y_i) \wedge z_i) &= f\left(\sum_{i \in e} x_i\right) \cdot \prod_{i \in e} f((w_i \vec{1} + y_i) \wedge z_i) \\ &= f\left(\sum_{i \in e} x_i\right) \cdot \prod_{i \in e} (-1)^{(w_i + y_i(\ell)) \wedge z_i(\ell)} \\ &= f\left(\sum_{i \in e} x_i\right), \end{aligned}$$

and similarly,

$$f_e\left(\sum_{i \in e} x_i + \left(\sum_{i \in e} (w_i \vec{1} + y_i)\right) \wedge z_e\right) = f\left(\sum_{i \in e} x_i\right).$$

Hence the test always accepts. ■

Lemma 3.9 *The H-Test has soundness $2^{-|E|}$.*

Before proving Lemma 3.9 we first prove a proposition relating the Fourier transform of a function perturbed by noise to the function's transform itself.

Proposition 3.10 *Let $f : \{0, 1\}^n \rightarrow \{-1, 1\}$. Define $g : \{0, 1\}^{2n} \rightarrow [-1, 1]$ such that*

$$g(x; y) = \mathbb{E}_{z \in \{0, 1\}^n} f(c' + x + (c + y) \wedge z),$$

where c, c' are some fixed vectors in $\{0, 1\}^n$. Then

$$\widehat{g}(\alpha; \beta)^2 = \widehat{f}(\alpha)^2 \mathbf{1}_{\{\beta \subset \alpha\}} 4^{-|\alpha|}.$$

Proof: This is a straightforward Fourier analytic calculation. By definition,

$$\widehat{g}(\alpha; \beta)^2 = \left(\mathbb{E}_{x, y, z \in \{0, 1\}^n} f(c' + x + (c + y) \wedge z) \chi_\alpha(x) \chi_\beta(y) \right)^2.$$

By averaging over x it is easy to see that

$$\widehat{g}(\alpha; \beta)^2 = \widehat{f}(\alpha)^2 \left(\mathbb{E}_{y, z \in \{0, 1\}^n} \chi_\alpha((c + y) \wedge z) \chi_\beta(y) \right)^2.$$

Since the bits of y are chosen independently and uniformly at random, if $\beta \setminus \alpha$ is nonempty, the above expression is zero. So we can write

$$\widehat{g}(\alpha; \beta)^2 = \widehat{f}(\alpha)^2 1_{\{\beta \subset \alpha\}} \left(\prod_{i \in \alpha \setminus \beta} \mathbb{E}_{y_i, z_i} (-1)^{(c_i + y_i) \wedge z_i} \cdot \prod_{i \in \beta} \mathbb{E}_{y_i, z_i} (-1)^{(c_i + y_i) \wedge z_i + y_i} \right)^2.$$

It is easy to see that the term $\mathbb{E}_{y_i, z_i} (-1)^{(c_i + y_i) \wedge z_i}$ evaluates to $\frac{1}{2}$ and the term $\mathbb{E}_{y_i, z_i} (-1)^{(c_i + y_i) \wedge z_i + y_i}$ evaluates to $(-1)^{c_i} \frac{1}{2}$. Thus

$$\widehat{g}(\alpha; \beta)^2 = \widehat{f}(\alpha)^2 1_{\{\beta \subset \alpha\}} 4^{-|\alpha|}$$

as claimed. \blacksquare

Now we prove Lemma 3.9.

Proof (of Lemma 3.9): Let p be the acceptance probability of H -test. Suppose that $2^{-|E|} + \epsilon \leq p$. We want to show that there are two functions f_a and f_b such that for some $i \in [n]$, some fixed positive integer w , some constant $\epsilon' > 0$, it is the case that $I_i^{\leq w}(f_a), I_i^{\leq w}(f_b) \geq \epsilon'$. As usual we first arithmetize p . We write

$$p = \sum_{v \in \{0,1\}^k} \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \prod_{i \in [k]} \frac{1 + (-1)^{v_i} f_i(y_i)}{2} \prod_{e \in E} \frac{1 + \text{Acc}(\{x_i, y_i, v_i, z_i\}_{i \in e}, z_e)}{2},$$

where

$$\begin{aligned} \text{Acc}(\{x_i, y_i, v_i, z_i\}_{i \in e}, z_e) &= \prod_{i \in e} [f_i(x_i + (v_i \vec{1} + y_i) \wedge z_i)] \\ &\cdot f_e \left(\sum_{i \in e} x_i + \left(\sum_{i \in e} (v_i \vec{1} + y_i) \right) \wedge z_e \right). \end{aligned}$$

For each $i \in [k]$, f_i is folded, so $(-1)^{v_i} f_i(y_i) = f_i(v_i \vec{1} + y_i)$. Since the vectors $\{y_i\}_{i \in [k]}$ are uniformly and independently chosen from $\{0, 1\}^n$, for a fixed $v \in \{0, 1\}^k$, the vectors $\{v_i \vec{1} + y_i\}_{i \in [k]}$ are also uniformly and independently chosen from $\{0, 1\}^n$. So we can simplify the expression for p and write

$$p = \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \left[\prod_{i \in [k]} (1 + f_i(y_i)) \prod_{e \in E} \frac{1 + (\text{Acc}\{x_i, y_i, \vec{0}, z_i\}_{i \in e}, z_e)}{2} \right].$$

Instead of writing $\text{Acc}(\{x_i, y_i, \vec{0}, z_i\}_{i \in e}, z_e)$, for convenience we shall write $\text{Acc}(e)$ to be a notational shorthand.

Note that the product of sums $\prod_{e \in E} \frac{1 + \text{Acc}(e)}{2}$ expands into a sum of products of the form

$$2^{-|E|} \left(1 + \sum_{\emptyset \neq E' \subset E} \prod_{e \in E'} \text{Acc}(e) \right).$$

f_i being folded implies that $\mathbb{E} f_i = 0$, so we have

$$\epsilon \leq \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \left[\prod_{i \in [k]} (1 + f_i(y_i)) 2^{-|E|} \sum_{\emptyset \neq E' \subset E} \prod_{e \in E'} \text{Acc}(e) \right].$$

By averaging, there must exist some nonempty subset $E' \subset E$ such that

$$\epsilon \leq \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \left[\prod_{i \in [k]} (1 + f_i(y_i)) \prod_{e \in E'} \text{Acc}(e) \right].$$

Since $1 + f_i(y_i)$ is either 0 or 2, we have

$$\frac{\epsilon}{2^k} \leq \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \left[\prod_{e \in E'} \text{Acc}(e) \right].$$

Let Odd consists of the vertices in $[k]$ with odd degree in E' . Expanding out the definition of $\text{Acc}(e)$, we can conclude

$$\frac{\epsilon}{2^k} \leq \mathbb{E}_{\{x_i\}, \{y_i\}, \{z_a\}} \left[\prod_{i \in \text{Odd}} f_i(x_i + y_i \wedge z_i) \cdot \prod_{e \in E'} f_e \left(\sum_{i \in e} x_i + \left(\sum_{i \in e} y_i \right) \wedge z_e \right) \right].$$

We now define a family of functions that represent the “noisy versions” of f_a . For $a \in [k] \cup E$, define $g'_a : \{0, 1\}^{2n} \rightarrow [-1, 1]$ to be

$$g'_a(x; y) = \mathbb{E}_{z \in \{0, 1\}^n} f_a(x + y \wedge z).$$

Thus we have

$$\frac{\epsilon}{2^k} \leq \mathbb{E}_{\{x_i\}, \{y_i\}} \left[\prod_{i \in \text{Odd}} g'_i(x_i; y_i) \cdot \prod_{e \in E'} g'_e \left(\sum_{i \in e} x_i; \sum_{i \in e} y_i \right) \right].$$

Following the approach of [12, 21], we are going to reduce the analysis of the iterated test to one hyperedge. Let d be the maximum size of an edge in E' , and without loss of generality, let $(1, \dots, d)$ be a maximum edge in E' . Now, fix the values of x_{d+1}, \dots, x_k and y_{d+1}, \dots, y_k so that the following inequality holds:

$$\frac{\epsilon}{2^k} \leq \mathbb{E}_{x_1, y_1, \dots, x_d, y_d} \left[\prod_{i \in \text{Odd}} g'_i(x_i; y_i) \cdot \prod_{e \in E'} g'_e \left(\sum_{i \in e} x_i; \sum_{i \in e} y_i \right) \right]. \quad (3.1)$$

We group the edges in E' based on their intersection with $(1, \dots, d)$. We rewrite Inequality 3.1 as

$$\frac{\epsilon}{2^k} \leq \mathbb{E}_{(x_1, y_1), \dots, (x_d, y_d) \in \{0, 1\}^{2n}} \left[\prod_{S \subset [d]} \prod_{a \in \text{Odd} \cup E': a \cap [d] = S} g_a \left(\sum_{i \in S} x_i; \sum_{i \in S} y_i \right) \right], \quad (3.2)$$

where for each $a \in [k] \cup E$, $g_a(x; y) = g'_a(c'_a + x; c_a + y)$, with $c'_a = \sum_{i \in a \setminus [d]} x_i$ and $c_a = \sum_{i \in a \setminus [d]} y_i$ fixed vectors in $\{0, 1\}^n$.

By grouping the edges based on their intersection with $[d]$, we can rewrite Inequality 3.2 as

$$\begin{aligned} \frac{\epsilon}{2^k} &\leq \mathbb{E}_{(x_1, y_1), \dots, (x_d, y_d) \in \{0, 1\}^{2n}} \left[\prod_{S \subset [d]} G_S \left(\sum_{i \in S} (x_i; y_i) \right) \right] \\ &= \langle \{G_S\}_{S \subset [d]} \rangle_{LU_d}, \end{aligned}$$

where G_S is simply the product of all the functions g_a such that $a \in \text{Odd} \cup E'$ and $a \cap [d] = S$.

Since $(1, \dots, d)$ is maximum, all the other edges in E' do not contain $(1, \dots, d)$ as a subset. Thus $G_{[d]} = g_{[d]}$ and $\mathbb{E} G_{[d]} = 0$. By Lemma 2.7, the linear Gowers inner product of a family of functions $\{G_S\}$ is positive implies that two functions from the family must share a variable with positive influence. More precisely, there exist $S \neq T \subset [d]$, $i \in [2n]$, $\tau > 0$, such that $I_i(G_S), I_i(G_T) \geq \tau$, where $\tau = \frac{\epsilon^4}{2^{\mathcal{O}(d)}}$.

Note that G_\emptyset is the product of all g'_a that depends on the fixed variables with indices outside of $[d]$. So G_\emptyset is a constant function and all its variables clearly have influence 0. Thus neither S nor T is empty. Since G_S and G_T are products of at most 2^k functions, by Lemma 2.3 there must exist some $a \neq b \in [d] \cup E'$ such that $I_i(g_a), I_i(g_b) \geq \frac{\tau}{2^{2k}}$. Recall that we have defined $g_a(x; y)$ to be $\mathbb{E}_z f_a(c'_a + x + (c_a + y) \wedge z)$. Thus we can apply Proposition 3.10 to obtain

$$\begin{aligned} I_i(g_a) &= \sum_{(\alpha, \beta) \in \{0,1\}^{2n}; i \in (\alpha, \beta)} \widehat{g}_a(\alpha; \beta)^2 \\ &= \sum_{\alpha \in \{0,1\}^n; i \in \alpha} \sum_{\beta \subset \alpha} \widehat{f}_a(\alpha)^2 4^{-|\alpha|} \\ &= \sum_{\alpha \in \{0,1\}^n; i \in \alpha} \widehat{f}_a(\alpha)^2 2^{-|\alpha|}. \end{aligned}$$

Let w be the least positive integer such that $2^{-w} \leq \frac{\tau}{2^{2k+1}}$. Then it is easy to see that $I_i^{\leq w}(f_a) \geq \frac{\tau}{2^{2k+1}}$. Similarly, $I_i^{\leq w}(f_b) \geq \frac{\tau}{2^{2k+1}}$ as well. Hence this completes the proof. ■

4 Acknowledgments

I am grateful to Alex Samorodnitsky for many useful discussions and his help with the Gowers norm. I also thank Madhu Sudan for his advice and support and Swastik Kopparty for an encouraging discussion during the initial stage of this research.

References

- [1] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [2] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [3] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. Linearity testing in characteristic two. In *IEEE Symposium on Foundations of Computer Science*, pages 432–441, 1995.
- [4] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM J. Comput.*, 27(3):804–915, 1998.
- [5] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993.

- [6] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 344–353, New York, NY, USA, 2006. ACM.
- [7] Lars Engebretsen and Jonas Holmerin. More efficient queries in PCPs for NP and improved approximation hardness of maximum csp. In *STACS*, pages 194–205, 2005.
- [8] W. T. Gowers. A new proof of Szemerédi’s theorem. *Geom. Funct. Anal.*, 11(3):465–588, 2001.
- [9] Venkatesan Guruswami, Daniel Lewin, Madhu Sudan, and Luca Trevisan. A tight characterization of NP with 3 query PCPs. In *FOCS '98: Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, page 8, Washington, DC, USA, 1998. IEEE Computer Society.
- [10] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [11] Johan Håstad and Subhash Khot. Query efficient PCPs with perfect completeness. *Theory of Computing*, 1(7):119–148, 2005.
- [12] Johan Håstad and Avi Wigderson. Simple analysis of graph tests for linearity and PCP. *Random Struct. Algorithms*, 22(2):139–160, 2003.
- [13] Subhash Khot. On the power of unique 2-prover 1-round games. In *STOC '02: Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 767–775, New York, NY, USA, 2002. ACM.
- [14] Subhash Khot and Rishi Saket. A 3-query non-adaptive PCP with perfect completeness. In *CCC '06: Proceedings of the 21st Annual IEEE Conference on Computational Complexity*, pages 159–169, Washington, DC, USA, 2006. IEEE Computer Society.
- [15] Ryan O’Donnell and Yi Wu. 3-bit dictator testing: 1 vs. 5/8. In *SODA '09: Proceedings of the 20th Annual Symposium on Discrete Algorithms*, To Appear.
- [16] Michal Parnas, Dana Ron, and Alex Samorodnitsky. Testing basic boolean formulae. *SIAM Journal on Discrete Mathematics*, 16(1):20–46, 2002.
- [17] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [18] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [19] Alex Samorodnitsky. Low-degree tests at large distances. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 506–515, New York, NY, USA, 2007. ACM.
- [20] Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 191–199, New York, NY, USA, 2000. ACM.
- [21] Alex Samorodnitsky and Luca Trevisan. Gowers uniformity, influence of variables, and PCPs. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 11–20, New York, NY, USA, 2006. ACM.
- [22] M. Sudan and L. Trevisan. Probabilistically checkable proofs with low amortized query complexity. In *FOCS '98: Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, page 18, Washington, DC, USA, 1998. IEEE Computer Society.

- [23] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975.
- [24] Terence Tao. Structure and randomness in combinatorics. In *FOCS '07: Proceedings of the forty-eighth annual ACM symposium on Foundations of computer science*, pages 3–15, New York, NY, USA, 2007. ACM.
- [25] Luca Trevisan. Recycling queries in PCPs and in linearity tests (extended abstract). In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 299–308, New York, NY, USA, 1998. ACM.