Linear Consistency Testing*

Yonatan Aumann† Johan Håstad‡ Michael O. Rabin§ Madhu Sudan¶

February 18, 2002

*A preliminary version of this paper appeared in the Proceedings of the 3rd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM '99), Berkeley, California.

†Department of Mathematics and Computer Science, Bar-Ilan University, Ramat-Gan, 52900, Israel. Email: aumann@cs.biu.ac.il.

‡Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, Sweden. Email: johanh@nada.kth.se.

§DEAS, Harvard University, Cambridge, MA 02138, USA and Institute of Computer Science, Hebrew University, Jerusalem, Israel. Email: rabin@deas.harvard.edu. Research supported, in part, by NSF Grant NSF-CCR-97-00365.

¶Department of Electrical Engineering and Computer Science, MIT, 545 Technology Square, Cambridge, MA 02139, USA. Email: madhu@mit.edu. Research supported in part by a Sloan Foundation Fellowship, an MIT-NEC Research Initiation Grant and NSF Career Award CCR-9875511.

Abstract

We extend the notion of linearity testing to the task of checking linear-consistency of multiple functions. Informally, functions are "linear" if their graphs form straight lines on the plane. Two such functions are "consistent" if the lines have the same slope. We propose a variant of a test of Blum, Luby and Rubinfeld [8] to check the linear-consistency of three functions f_1, f_2, f_3 mapping a finite Abelian group G to an Abelian group H: Pick $x, y \in G$ uniformly and independently at random and check if $f_1(x) + f_2(y) = f_3(x + y)$. We analyze this test for two cases: (1) G and H are arbitrary Abelian groups and (2) $G = \mathbb{F}_2^n$ and $H = \mathbb{F}_2$.

Questions bearing close relationship to linear-consistency testing seem to have been implicitly considered in recent work on the construction of PCPs and in particular in the work of Håstad [9]. It is abstracted explicitly for the first time here. As an application of our results we give yet another new and tight characterization of NP, namely $\forall \epsilon > 0$, $\mathrm{NP} = \mathrm{MIP}_{1-\epsilon,\frac{1}{2}}[O(\log n), 3, 1]$. I.e., every language in NP has 3-prover 1-round proof systems in which the verifier tosses $O(\log n)$ coins and asks each of the three provers one question each. The provers respond with one bit each such that the verifier accepts instance of the language with probability $1 - \epsilon$ and rejects non-instances with probability at least $\frac{1}{2}$. Such a result is of some interest in the study of probabilistically checkable proofs.

1 Introduction

The study of linearity testing was initiated by Blum, Luby and Rubinfeld in [8]. A function f mapping a finite Abelian group G to an Abelian group H is "linear" (or more conventionally, a homomorphism) if for every $x, y \in G$, $f(x) + f(y) = f(x + y)$. Blum, Luby and Rubinfeld showed that if a function f satisfies the identity above for a large fraction of pairs $x, y \in G$, then f is close to being linear. This seminal result played a catalytic role in the study of program checking/self-testing [7, 8]. It is also a crucial element in the development of efficient PCP characterizations of NP and in particular occupies a central role in the results of [1, 6, 5].

In this paper we extend this study to testing the consistency of multiple functions. Given a triple of functions $f_1, f_2, f_3 : G \to H$, we say that they are "linear-consistent" if they satisfy: $\forall x, y, \in G,\ f_1(x) + f_2(y) = f_3(x + y)$. [1] At first glance this definition does not seem to enforce any structural property in f_1, f_2 or f_3. We show, however, that if f_1, f_2, f_3 are linear-consistent, then they are: (1) Affine: I.e., there exists $a_1, a_2, a_3 \in H$ such that for every $i \in \{1, 2, 3\}$ and $\forall x, y \in G,\ f_i(x) + f_i(y) = f_i(x + y) + a_i$; and (2) Consistent: I.e., $a_1 + a_2 = a_3$ and for every $i, j \in \{1, 2, 3\}$ and $\forall x \in G, f_i(x) - a_i = f_j(x) - a_j$.

We go on to study triples of functions f_1, f_2, f_3 that do not satisfy the identity $f_1(x) + f_2(y) = f_3(x + y)$ everywhere, but do satisfy this identity with high probability over a random choice of x and y. We provide two analyses for this case. The first is a variant of the analysis due to Coppersmith described in [8] for linearity testing over arbitrary Abelian groups. We obtain the following result:

If $f_1, f_2, f_3 : G \to H$ satisfy $\delta \stackrel{\triangle}{=} \Pr_{x,y \in G}[f_1(x) + f_2(y) \neq f_3(x+y)] < \frac{2}{9}$, then there exists

[1] A slightly more symmetric equivalent definition would be to use: $\forall x, y, z \in G$ such that $x + y + z = 0$, $f_1(x) + f_2(y) + f_3'(z) = 0$. To see that this is equivalent set $f_3'(z) = -f_3(-z)$.

a triple of linear-consistent functions $\tilde{f}_1, \tilde{f}_2, \tilde{f}_3 : G \to H$ such that for every $i \in \{1,2,3\}$, $\Pr_{x \in G}[f_i(x) \neq \tilde{f}_i(x)] \leq \delta$.

The second variant we study is when $G = \mathbb{F}_2^n$ and $H = \mathbb{F}_2$, where $\mathbb{F}_2$ is the finite field of two elements. This special case is of interest due to its applicability in the construction of efficient "probabilistically checkable proofs" and has been extensively studied due to this reason — see the work of Bellare et al. [4] and the references therein. Bellare et al. [4] give a nearly tight analysis of the linearity test in this case and show, among other things, that if a function f fails the linearity test with probability at most δ then it is within a distance of δ from some linear function. We extend their analysis to the case of linear-consistency testing and show an analogous result for this test:

If $f_1, f_2, f_3 : \mathbb{F}_2^n \to \mathbb{F}_2$ and $\gamma > 0$, satisfy $\Pr_{x,y \in \mathbb{F}_2^n}[f_1(x) + f_2(y) \neq f_3(x+y)] = \frac{1}{2} - \gamma < \frac{1}{2}$, then there exists a triple of linear-consistent functions $\tilde{f}_1, \tilde{f}_2, \tilde{f}_3 : \mathbb{F}_2^n \to \mathbb{F}_2$ such that for every $i \in \{1,2,3\}$, $\Pr_{x \in \mathbb{F}_2}[f_i(x) \neq \tilde{f}_i(x)] \leq \frac{1}{2} - \frac{2\gamma}{3}$.

Motivation: We believe that the linear-consistency test is a natural variant of the linearity test and will potentially find similar applications in general. Our original motivation came from the analysis of a variant of a protocol for deniable encryption proposed by Aumann and Rabin [3]. However, at this point we do not have any concrete applications to this case.

One scenario where the linear-consistency test does seem to appear naturally is the case of probabilistically checkable proofs or variants thereof. Tasks similar to linear-consistency testing were implicit in the works of Håstad (e.g. in [9]), where probabilisitic checks to check "validity" and "consistency" of two functions A and B are often used. The notion of validity used in [9] is a more stringent one than that of linearity, however the analysis techniques are similar. In this paper

we derive an application to the construction of "multiple prover proof systems for NP". Another situation where linear-consistency testing plays a small role is in a recent result of Håstad and Wigderson [10]. We describe these applications in the paragraphs below.

Multiple-prover Interactive Proofs: An (r, p, a)-restricted MIP verifier V (for a p-prover one-round proof system) is one that acts as follows: On input $x \in \{0,1\}^n$, V tosses $r(n)$ random coins and generates one question each for each of the p provers. The provers respond with a bits each. The response of the ith prover is allowed to be an arbitrary function of x and the query to the i prover, but is independent of the queries to the other provers. The verifier then outputs a verdict "accept/reject" based on the input x, its random coins and the answers of the p-provers. V is said to verify membership of a language L with completeness c and soundness s, if for every $x \in L$, there exist p-provers that are accepted by V with probability at least c; and for every $x \notin L$, for every p-provers, the verifier accepts with probability at most s. The class of all languages with p-prover one-round proof systems, in which the provers respond with a bits and the verifier is $r(\cdot)$ restricted and has completeness c and soundness s is denoted $\mathrm{MIP}_{c,s}[r, p, a]$.

Multiple prover interactive proof systems (MIPs) are a special case of the more familiar case of probabilistically checkable proof systems (PCPs). The difference is that in a PCP, all questions are sent to one "oracle-prover". The two main parameters of interest are the "randomness-parameter" (same as in MIP) and the "query-parameter", which counts the total number of bits of response from the oracle-prover. Thus the following containment is obtained easily $\mathrm{MIP}_{c,s}[r, p, a] \subseteq \mathrm{PCP}_{c,s}[r, p \cdot a]$ (where the second parameter is the number of queries). However, a converse of the form $\mathrm{PCP}_{c,s}[r, q] \subseteq \mathrm{MIP}_{c,s}[r, q, 1]$ is not known to be true and is a subject of some interest. Most strong PCP constructions today are obtained from some strong MIP construction. It is generally believed that MIP is a more restrictive model, but no results are known separating p-prover

1-bit MIPs from p-query PCPs. In view of the recent tight analysis of 3-query proof systems by Håstad [9] showing NP $= \mathrm{PCP}_{1-\epsilon,\frac{1}{2}}[\log, 3]$, it was conceivable that one could separate 3-query PCPs from 3-prover 1-bit proof systems. However, our analysis of the linear-consistency tests leads us to an equally tight characterization of NP with MIPs. We show:

$$\forall \epsilon > 0, \mathrm{NP} = \mathrm{MIP}_{1-\epsilon,\frac{1}{2}}[O(\log n), 3, 1].$$

In fact in view of our analysis we believe that there may be no separation between p-prover 1-bit MIPs and p-query PCPs for any constant p.

Graph-based linearity tests. Graph-based linearity tests were introduced by Trevisan [14], as a means to study a variety of "linearity tests" that are more complicated that the BLR test, but are more efficient in some senses. Nearly-optimal analyses of graph-based linearity tests were given by Samorodnitsky and Trevisan [12]. A recent result of Håstad and Wigderson [10] shows how this analysis could be simplified significantly. Linear-consistency testing plays a small but arguably crucial role in this simplified analysis. The analysis of [10] reexpresses any graph-based linearity test as a linear-consistency test on three related functions. Their analysis abstracts away the complications arising from the definition of the test into the complex relations satisfied by the functions. The analysis then ignores the relations satisfed by these functions and instead just applies the analysis of linear consistency testing to this triple. This yields that these functions are close to some linear-consistent triple, which in their case immediately implies that the function being tested is close to being linear. While their proofs can be (and are) described without mention of linear-consistency testing, the concept seems to play an important role in their analysis.

Outline of this paper. In Section 2 we present some basic definitions of linear-consistency. In Section 3 we provide the analysis of linear-consistency tests over arbitrary Abelian groups. In Section 4 we consider the special case where the groups are vector spaces over $\mathbb{F}_2$. In Section 5 we give the MIP construction.

2 Definitions

For groups G, H, let $\text{HOM}_{G \to H}$ denote the set of homomorphisms from G to H. I.e.,

$$\text{HOM}_{G \to H} \stackrel{\triangle}{=} \{\phi : G \to H | \forall x, y \in G, \phi(x) + \phi(y) = \phi(x + y)\}.$$

For groups G, H, let $\text{AFF}_{G \to H}$ denote the set of *affine* functions from G to H. I.e.,

$$\text{AFF}_{G \to H} \stackrel{\triangle}{=} \{\psi : G \to H | \exists a \in H, \phi \in \text{HOM}_{G \to H} \text{ s.t. } \forall x \in G, \psi(x) = \phi(x) + a\}.$$

A triple of functions (f_1, f_2, f_3) is defined to be *linear-consistent* if there exists a homomorphism $\phi \in \text{HOM}_{G \to H}$ and $a_1, a_2, a_3 \in H$ such that $a_1 + a_2 = a_3$ and for every $i \in \{1, 2, 3\}$ and $x \in G$, $f_i(x) = \phi(x) + a_i$.

The following proposition gives an equivalent characterization of linear-consistent functions.

Proposition 1 *Functions $f_1, f_2, f_3 : G \to H$ are linear-consistent if and only if for every $x, y \in G$, $f_1(x) + f_2(y) = f_3(x + y)$.*

Proof: Let f_1, f_2, f_3 be linear-consistent, and let $\phi \in \text{HOM}_{G \to H}$ and $a_1, a_2, a_3 \in H$ be as guaranteed to exist by the definition of linear-consistency. Then, for every $x, y \in G$, $f_1(x) + f_2(y) -$

$f_3(x+y) = \phi(x) + \phi(y) - \phi(x+y) + a_1 + a_2 - a_3 = 0$ as required. This gives one direction of the proposition.

Now suppose f_1, f_2, f_3 satisfy $\forall x, y, f_1(x) + f_2(y) = f_3(x+y)$. Using $x = y = 0$, we get

$$f_1(0) + f_2(0) = f_3(0) \tag{1}$$

Next we notice that $f_1(x)+f_2(0) = f_3(x)$ (using $y = 0$). Subtracting $f_1(0)+f_2(0) = f_3(0)$ from both sides we get $f_1(x) - f_1(0) = f_3(x) - f_3(0)$. Similarly we get $f_2(x) - f_2(0) = f_3(x) - f_3(0)$. Thus we may define $\phi(x) = f_1(x) - f_1(0) = f_2(x) - f_2(0) = f_3(x) - f_3(0)$. We now verify that $\phi \in \text{HOM}_{G \to H}$. For arbitrary $x, y \in G$, $\phi(x) + \phi(y) - \phi(x+y) = f_1(x) - f_1(0) + f_2(y) - f_2(0) - (f_3(x+y) - f_3(0)) = (f_1(x) + f_2(y) - f_3(x+y)) - (f_1(0) + f_2(0) - f_3(0)) = 0$. Thus for $a_i = f_i(0)$ and ϕ as above, we see that f_1, f_2, f_3 satisfy the definition of linear-consistency. ∎

For $x, y \in G$, the *linear-consistency test* through x and y is the procedure which accepts iff $f_1(x) + f_2(y) = f_3(x+y)$. Our goal in the remaining sections is to derive relationships between the probability with which a triple f_1, f_2, f_3 is rejected by the linear-consistency tests when x and y are chosen at random, and the proximity of f_1, f_2 and f_3 to linear-consistent functions.

3 Linear-consistency over arbitrary Abelian groups

In this section we consider the case of G and H being arbitrary finite Abelian groups. We extend the argument due to Coppersmith that appears in [8] to this case. We show that if the test rejects with probability $\delta < \frac{2}{9}$, then by changing the value of each of the f_i's on at most δ fraction on the inputs, we get a triple of linear-consistent functions. In what follows, we use $d(f, g)$ to denote the distance of f from g, i.e., $\Pr_{x \in G}[f(x) \neq g(x)]$.

Theorem 2 *Let G, H be finite Abelian groups and let $f_1, f_2, f_3 : G \to H$. If*

$$\delta \stackrel{\triangle}{=} \Pr_{x,y \in G}[f_1(x) + f_2(y) \neq f_3(x+y)] < \frac{2}{9},$$

then there exists a triple of linear-consistent functions g_1, g_2, g_3 such that for every $i \in \{1,2,3\}$, $\epsilon_i \stackrel{\triangle}{=} d(f_i, g_i) \leq \delta$. Furthermore, $\epsilon \stackrel{\triangle}{=} \frac{\epsilon_1+\epsilon_2+\epsilon_3}{3}$ satisfies $3\epsilon(1-2\epsilon) \leq \delta$.

Remark 3

1. *If $f_1 = f_2 = f_3$, then we recover the linearity testing theorem of [8] (see also [4]).*
2. *The proof actually shows that $\epsilon_1 + \epsilon_2 + \epsilon_3 - 2(\epsilon_1\epsilon_2 + \epsilon_2\epsilon_3 + \epsilon_3\epsilon_1) \leq \delta$. Tightness of this and other aspect of the theorem are discussed in Section 3.1.*

Proof: For $f : G \to H$, define $\textsc{Corr}_y^f(x)$ to be $f(x+y) - f(y)$. Define

$$\tilde{f}(x) = \textsc{Plurality}_{i \in \{1,2,3\}, y \in G}\{\textsc{Corr}_y^{f_i}(x)\},$$

where $\textsc{Plurality}(S)$ for a multiset S is the most commonly occurring element in S, with ties being broken arbitrarily. Note that if f_1, f_2, f_3 are linear consistent then $\textsc{Corr}_y^{f_i}(x) = \phi(x)$ for any i and y and the hope in general is that $\tilde{f}$ should equal the sought after ϕ.

For $i \in \{1,2,3\}$ and $x \in G$, let $\gamma_i(x) \stackrel{\triangle}{=} \Pr_{y \in G}\left[\tilde{f}(x) \neq \textsc{Corr}_y^{f_i}(x)\right]$. Let $\gamma_i = \mathbf{E}_x[\gamma_i(x)]$. Let $\gamma(x) = \frac{1}{3}[\gamma_1(x) + \gamma_2(x) + \gamma_3(x)]$ and let $\gamma = \mathbf{E}_x[\gamma(x)]$. Note that, by the definitions, $\gamma = \frac{\gamma_1+\gamma_2+\gamma_3}{3}$.

Our plan is to show that the $\gamma_i(x)$'s are all small and then to use this in two ways: First we use it to show that $\tilde{f}$ is a homomorphism. Then we show that the functions f_i's are within a distance of γ_i from affine functions that are in the orbit of $\tilde{f}$.

Claim 4 *For every $x \in G$, and $i \neq j \in \{1,2,3\}$,*

$$\Pr_{y_1,y_2}\left[\textsc{Corr}_{y_1}^{f_i}(x) \neq \textsc{Corr}_{y_2}^{f_j}(x)\right] \leq 2\delta.$$

Proof: We prove the claim only for the case $i = 1, j = 2$. Other cases are proved similarly.

Over the choice of y_1 and y_2, consider two possible "bad" events: (A) $f_1(x + y_1) + f_2(y_2) \neq f_3(x + y_1 + y_2)$ and (B) $f_1(y_1) + f_2(x + y_2) \neq f_3(x + y_1 + y_2)$. Observe first that if neither of the bad events listed above occur, then we have

$$\begin{aligned}
&\textsc{Corr}_{y_1}^{f_1}(x) \\
&= f_1(x + y_1) - f_1(y_1) \\
&= (f_3(x + y_1 + y_2) - f_2(y_2)) - f_1(y_1) && \text{((A) does not occur)} \\
&= (f_3(x + y_1 + y_2) - f_2(y_2)) - (f_3(x + y_1 + y_2) - f_2(x + y_2)) && \text{((B) does not occur)} \\
&= f_2(x + y_2) - f_2(y_2) \\
&= \textsc{Corr}_{y_2}^{f_2}(x).
\end{aligned}$$

Now notice that the event listed in (A) has probability exactly δ (in particular, this event is independent of x). Similarly probability of the event in (B) is also δ. Thus the probability that (A) or (B) occurs may be bounded from above by 2δ. The claim follows. ∎

The claim above allows us to prove upper bounds on the quantities $\gamma_i(x)$ for every x. This implies, in particular, that the function $\tilde{f}$ is defined at every point x by an overwhelming majority; a fact that is critical in proving that $\tilde{f}$ is a homomorphism.

Claim 5 *For every $x \in G$, and $i \in \{1,2,3\}$ and $j \neq i \in \{1,2,3\}$, the following hold:*

1. $\gamma_i(x) \leq 2\delta$.

2. $\gamma_i(x) + \gamma_j(x) - 2\gamma_i(x)\gamma_j(x) \leq 2\delta$.

3. $\gamma(x) < \frac{1}{3}$.

Proof: Fix x and for $\alpha \in H$, let $p_\alpha = \Pr_{y \in G}[\textsc{Corr}_y^{f_i}(x) = \alpha]$ and $q_\alpha = \Pr_{y \in G}[\textsc{Corr}_y^{f_j}(x) = \alpha]$. The hope is that the same value of α maximizes both p_α and q_α and this value has then to be $\tilde{f}(x)$.

We start by showing that $\max_{\alpha \in H}\{p_\alpha\}$ is very large. Observe that

$$\Pr_{y_1,y_2}[\textsc{Corr}_{y_1}^{f_i}(x) = \textsc{Corr}_{y_2}^{f_j}(x)] = \sum_{\alpha \in H} p_\alpha q_\alpha \leq \max_{\alpha \in H}\{p_\alpha\} \cdot \sum_{\alpha \in H} q_\alpha = \max_{\alpha \in H}\{p_\alpha\}.$$

Using Claim 4 the left-hand side of the inequality above is at least $1 - 2\delta$. Thus we establish that $\max_\alpha\{p_\alpha\} \geq 1 - 2\delta > 5/9$. Similarly we can show that $\max_\alpha\{q_\alpha\} > 5/9$.

Next we show that these maxima occur for the same value of $\alpha \in H$. Assume otherwise. Let $\tilde{p} = \max_\alpha\{p_\alpha\}$ and $\tilde{q} = \max_\alpha\{q_\alpha\}$. By the above $\tilde{p}, \tilde{q} > 5/9 > 1/2$. Since the maxima occur for distinct values of α, we may upper bound the quantity $\Pr_{y_1,y_2}[\textsc{Corr}_{y_1}^{f_i}(x) = \textsc{Corr}_{y_2}^{f_j}(x)]$ by $\tilde{p}(1-\tilde{q})+(1-\tilde{p})\tilde{q}$. With some manipulation, the latter quantity is seen to be equal to $\frac{1}{2} - 2(\tilde{p} - \frac{1}{2})(\tilde{q} - \frac{1}{2}) < \frac{1}{2}$, which contradicts Claim 4.

Thus we find that $\textsc{Plurality}_y\{\textsc{Corr}_y^{f_i}(x)\}$ points to the same value for every $i \in \{1,2,3\}$; and this value is $\tilde{f}(x)$. Thus we conclude $\gamma_i(x) = 1 - \max_\alpha\{p_\alpha\} \leq 2\delta$, yielding Part (1) of the claim. Part (2) follows by observing that

$$\Pr_{y_1,y_2}[\textsc{Corr}_{y_1}^{f_i}(x) = \textsc{Corr}_{y_2}^{f_j}(x)] \;\leq\; (1 - \gamma_i(x))(1 - \gamma_j(x)) + \gamma_i(x)\gamma_j(x)$$

and then using Claim 4 to lower bound the left-hand side by $1 - 2\delta$.

Adding the inequalities given by Part (2) for the three different choices of i, j gives

$$2(\gamma_1(x) + \gamma_2(x) + \gamma_3(x)) - 2(\gamma_1(x)\gamma_2(x) + \gamma_2(x)\gamma_3(x) + \gamma_3(x)\gamma_1(x)) \leq 6\delta.$$

Notice that for any a, b, c we have

$$(a+b+c)^2 = a^2 + b^2 + c^2 + 2(ab+bc+ca) \geq (a+b+c)^2/3 + 2(ab+bc+ca)$$

and hence

$$ab + bc + ca \leq (a+b+c)^2/3. \tag{2}$$

Using this inequality for $a = \gamma_1(x), b = \gamma_2(x), c = \gamma_3(x)$ and using the fact that $\gamma(x) = \frac{1}{3}(\gamma_1(x) + \gamma_2(x) + \gamma_3(x))$, we get

$$6\gamma(x) - 6\gamma(x)^2 \leq 6\delta.$$

Using the fact that $\delta < 2/9$, this yields that either $\gamma(x) < \frac{1}{3}$ or $\gamma(x) > \frac{2}{3}$. Since $\gamma_1(x), \gamma_2(x), \gamma_3(x) < \frac{4}{9}$ (by Part (1)) and $\gamma(x) = \frac{1}{3}(\gamma_1(x) + \gamma_2(x) + \gamma_3(x))$, we rule out the latter possibility. This yields Part (3) of the claim. ∎

The following claim now follows by a convexity argument.

Claim 6 *For every distinct $i, j \in \{1, 2, 3\}$, $\gamma_i + \gamma_j - 2\gamma_i\gamma_j \leq 2\delta$.*

Proof: By Part (2) of Claim 5 we know that for every $x \in G$, $\gamma_i(x) + \gamma_j(x) - 2\gamma_i(x)\gamma_j(x) \leq 2\delta$. Rewriting, we get that for every $x \in G$ $(\frac{1}{2} - \gamma_i(x))(\frac{1}{2} - \gamma_j(x)) \geq \frac{1}{4} - \delta$. By Part (1) of Claim 5, we also know that $\gamma_i(x), \gamma_j(x) \leq 2\delta < \frac{1}{2}$. The set $\{(\alpha, \beta) \in \mathbb{R}^2 | \alpha, \beta > 0,\ \alpha\beta > \frac{1}{4} - \delta\}$ is convex and

since the average of a set of points that belong to a convex set belongs to the same convex set, we find that $\gamma_i = \mathbf{E}_x[\gamma_i(x)]$ and $\gamma_j = \mathbf{E}_x[\gamma_j(x)]$ also satisfies the inequality $(\frac{1}{2} - \gamma_i)(\frac{1}{2} - \gamma_j) \geq \frac{1}{4} - \delta$. The claim follows immediately. ∎

Claim 7 *$\tilde{f}$ is a homomorphism. I.e., $\forall\ x, y \in G$, $\tilde{f}(x) + \tilde{f}(y) = \tilde{f}(x+y)$.*

Proof: Fix $x, y \in G$. We will show that there exist $i \in \{1,2,3\}$ and $u \in G$ (by picking them at random) such that none of the following bad events occur. (A) $\tilde{f}(x) \neq f_i(x+u) - f_i(u)$; (B) $\tilde{f}(y) \neq f_i(u) - f_i(u-y)$; and (C) $\tilde{f}(x+y) \neq f_i(x+u) - f_i(u-y)$.

It is immediate that if none of the events (A)-(C) occur, then

$$\tilde{f}(x) + \tilde{f}(y) - \tilde{f}(x+y) = (f_i(x+u) - f_i(u)) + (f_i(u) - f_i(u-y)) - (f_i(x+u) - f_i(u-y)) = 0.$$

The probability that (A) occurs is, by definition, $\gamma(x)$ and similarly the probabilities of (B) and (C) occurring are given by $\gamma(y)$ and $\gamma(x+y)$, respectively. By the union bound, the probability that (A) or (B) or (C) occurs is, using Claim 5, Part (3), strictly less than 1. Thus such a pair (i,u) does exist. ∎

Claim 8 *For every $i \in \{1,2,3\}$, there exists $\alpha_i \in H$ such that*

$$\Pr_{x \in G}[f_i(x) \neq \tilde{f}(x) + \alpha_i] \leq \gamma_i.$$

Furthermore $\alpha_1 + \alpha_2 = \alpha_3$.

Proof: Fix $i \in \{1,2,3\}$. By definition of $\gamma_i(x)$, we have for every x, $\Pr_{a \in G}[\tilde{f}(x) \neq f_i(x+a) - f_i(a)] \leq \gamma_i(x)$. Thus, we get $\Pr_{x,a \in G}[\tilde{f}(x) \neq f_i(x+a) - f_i(a)] \leq \gamma_i$. In particular, there exists $a_0 \in G$

such that $\Pr_{x\in G}[\tilde{f}(x) \neq f_i(x+a_0) - f_i(a_0)] \leq \gamma_i$ or equivalently $\Pr_{x\in G}[\tilde{f}(x-a_0) \neq f_i(x) - f_i(a_0)] \leq \gamma_i$. But $\tilde{f}$ is a homomorphism, and thus we have $\tilde{f}(x-a_0) = \tilde{f}(x) - \tilde{f}(a_0)$. Thus we find that for this choice of a_0, $\Pr_{x\in G}[f_i(x) \neq \tilde{f}(x) + f_i(a_0) - \tilde{f}(a_0)] \leq \gamma_i$. The first part of the claim follows by setting $\alpha_i = f_i(a_0) - \tilde{f}(a_0)$.

To prove the second part assume for contradiction that $\alpha_1 + \alpha_2 \neq \alpha_3$. Say that x is *i-good* if $f_i(x) = \tilde{f}(x) + \alpha_i$. The probability that x is 1-good, y is 2-good and $(x+y)$ is 3-good is at least

$$(1-\gamma_1)(1-(\gamma_2+\gamma_3)).$$

This follows since the probability that x is 1-good is least $1-\gamma_1$ and both the event that y is not 2-good and the event that $x+y$ is not 3-good is independent of x being 1-good. Hence, by the assumption $\alpha_1 + \alpha_2 \neq \alpha_3$, we conclude that

$$(1-\gamma_1)(1-(\gamma_2+\gamma_3)) \leq \delta.$$

Using the symmetric arguments and adding the 3 inequalities we get

$$3 - 3(\gamma_1+\gamma_2+\gamma_3) + 2(\gamma_1\gamma_2 + \gamma_1\gamma_3 + \gamma_2\gamma_3) \leq 3\delta. \tag{3}$$

Using Claim 6 (for all distinct pairs i, j) we get (after some rearrangement) that

$$2(\gamma_1+\gamma_2+\gamma_3) - 6\delta \leq 2(\gamma_1\gamma_2 + \gamma_2\gamma_3 + \gamma_3\gamma_1). \tag{4}$$

Adding Equations (3) and (4) and using $\gamma_1 + \gamma_2 + \gamma_3 = 3\gamma$ we get

$$3 - 3\gamma - 6\delta \leq 3\delta.$$

We conclude that

$$3\delta + \gamma \geq 1,$$

which contradicts $\delta < 2/9$ and $\gamma < 1/3$. ∎

We are almost done with the proof of Theorem 2. The final claim, sharpens the bounds on the proximity of the functions $f_i(x)$ to the functions $\tilde{f}(x) + \alpha_i$.

Claim 9 *The following inequalities hold:*

1. $\gamma_1 + \gamma_2 + \gamma_3 - 2(\gamma_1\gamma_2 + \gamma_2\gamma_3 + \gamma_3\gamma_1) \leq \delta$.
2. $3\gamma - 6\gamma^2 \leq \delta$.
3. $\gamma_1, \gamma_2, \gamma_3 \leq \delta$.

Proof: We proceed as in the proof of Part (2) of Claim 8. Recall that x is i-good if $f_i(x) = \tilde{f}(x) + \alpha_i$. Pick x, y at random and consider the events (A) x is not 1-good, (B) y is not 2-good (3) $x + y$ is not 3-good. Using the pairwise independence of the events, we can lower bound the probability that exactly one of the events (A), (B) or (C) occurs by

$$\gamma_1(1 - (\gamma_2 + \gamma_3)) + \gamma_2(1 - (\gamma_3 + \gamma_1)) + \gamma_3(1 - (\gamma_1 + \gamma_2)).$$

To see this note that the probability that, (A) occurs and (B) and (C) do not occur is at least $\gamma_1(1 - (\gamma_2 + \gamma_3))$ and the other terms follow similarly. But whenever exactly one of (A)-(C) occurs,

then the test rejects. Thus, the quantity above is at most δ and this yields Part (1) of the claim.

Part (2) follows by using $\gamma_1+\gamma_2+\gamma_3 = 3\gamma$ and using $\gamma_1\gamma_2+\gamma_2\gamma_3+\gamma_3\gamma_1 \leq 3\gamma^2$. The latter inequality is just a special case of (2).

For Part (3), we first use Part (2) to improve the bound on γ. Notice that by Part (2) of Claim 5, we know $\gamma < \frac{1}{3}$. Using Part (2) of this claim, we notice that we can improve upon this bound to $\gamma < \frac{1}{6}$ (no value γ in the interval $[\frac{1}{6}, \frac{1}{3}]$ satisfies $3\gamma - 6\gamma^2 < \frac{2}{9}$). Now assume for contradiction that $\gamma_1 > \delta$. Then rearranging the inequality from Part (1), we get

$$\gamma_1(1 - 2(\gamma_2 + \gamma_3)) + \gamma_2 + \gamma_3 - 2\gamma_2\gamma_3 \leq \delta.$$

Since $\gamma_2 + \gamma_3 \leq 3\gamma < \frac{1}{2}$, we notice that $1 - 2(\gamma_2 + \gamma_3) > 0$ and we can use $\gamma_1 > \delta$ to obtain:

$$\delta(1 - 2(\gamma_2 + \gamma_3)) + \gamma_2 + \gamma_3 - 2\gamma_2\gamma_3 < \delta.$$

$$\Rightarrow (\gamma_2 + \gamma_3)(1 - 2\delta) - 2\gamma_2\gamma_3 < 0.$$

$$\Rightarrow (\gamma_2 + \gamma_3)(1 - 2\delta) - \frac{1}{2}(\gamma_2 + \gamma_3)^2 < 0.$$

$$\Rightarrow (\gamma_2 + \gamma_3)(1 - 2\delta - \frac{1}{2}(\gamma_2 + \gamma_3)) < 0.$$

But the last inequality contradicts the fact that $\delta < 2/9$ and $\gamma_2 + \gamma_3 < \frac{1}{2}$. ∎

The theorem now follows from the above claims as follows. Set $g_i(x) = \tilde{f}(x) + \alpha_i$, where α_i's are as given by Claim 8. It follows from Claims 7 and 8 that g_1, g_2, g_3 are linear-consistent. It follows from Claim 8 that f_i is within a distance of γ_i from g_i; and the bounds on γ_i from Claim 9 bound these distances. ∎

3.1 Tightness of Theorem 2

Theorem 2 is tight in that one cannot improve the bound $\delta < \frac{2}{9}$ without significantly weakening the bound on the proximity of the nearest linear-consistent functions to f_1, f_2 and f_3. This tightness is inherited from the tightness of the linearity testing theorem of Blum, Luby and Rubinfeld, whose analysis also imposes the same upper bound on δ. For the sake of completeness, we recall the example, due to Coppersmith, here.

Let $G = H = \mathbb{Z}_{3n}$ for some large n, and let $f = f_1 = f_2 = f_3$ be the function

$$f(x) = \begin{cases} 3n-1 & \text{if } x = -1 \bmod 3 \\ 0 & \text{if } x = 0 \bmod 3 \\ 1 & \text{if } x = 1 \bmod 3 \end{cases}$$

Then the probability that the linearity test rejects is $\frac{2}{9}$, while (for large enough n), the nearest affine functions to f are the constant functions, which disagree from f in at least $\frac{2}{3}$ of the inputs.

As we increase $\delta > 2/9$, the bounds on the proximity of the nearest linear(-consistent) functions become worse, approaching 0 as $\delta \to 1/4$ as demonstrated by the following example. For positive integers m, n let $f : \mathbb{Z}_{(2m+1)n} \to \mathbb{Z}_{(2m+1)n}$ be the function $f(x) = x \bmod (2m+1)$ if $x \bmod (2m+1) \in \{0, \ldots, m\}$ and $f(x) = (x \bmod (2m+1)) + n - 2m - 1$ otherwise. It may be verified that the closest affine functions to f are the constant functions which are at a distance of at least $1 - \frac{1}{2m+1}$ from f. On the other hand the linearity test (and the hence the linear-consistency test on $f_1 = f_2 = f_3 = f$) accepts with probability at least $\frac{3}{4}$.

Thus for $\delta \geq \frac{1}{4}$ the linearity tests can not guarantee any non-trivial proximity with a linear function. In the range $\delta = [2/9, 1/4]$ we do not seem to have tight bounds. For $\delta < \frac{2}{9}$, the bounds given on ϵ_i can not be improved either, as shown in the following proposition.

Proposition 10 *For every $\epsilon_1, \epsilon_2, \epsilon_3 < \frac{1}{4}$, there exist a family of triples of functions $f_1^{(n)}, f_2^{(n)}, f_3^{(n)} : \mathbb{F}_2^n \to \mathbb{F}_2$ such that the distance of $f_i^{(n)}$ to the space of affine functions converges to ϵ_i and the probability that the linear-consistency test rejects is at most $\epsilon_1 + \epsilon_2 + \epsilon_3 - 2(\epsilon_1\epsilon_2 + \epsilon_2\epsilon_3 + \epsilon_3\epsilon_1)$.*

Proof: Let S_i be any subset of $\lfloor \epsilon_i 2^n \rfloor$ vectors from $\mathbb{F}_2^n$ with first coordinate being 1. Let $f_i^{(n)}(x) = 1 \Leftrightarrow x \in S_i$. Then, since $\epsilon_i < \frac{1}{4}$, the nearest affine function is the zero function, thus establishing the claim on distance. By the nature of the S_i's it is not possible that $x \in S_1$, $y \in S_2$ and $x+y \in S_3$. Therefore, the linear-consistency test rejects if and only if exactly one of $x, y, x+y$ fall in S_1, S_2, S_3 respectively. If we let ρ_i denote $2^{-n}|S_i|$, then the probability of this event is easily shown to be (exactly) $\rho_1+\rho_2+\rho_3-2(\rho_1\rho_2+\rho_2\rho_3+\rho_3\rho_1)$ which in turn is at most $\epsilon_1+\epsilon_2+\epsilon_3-2(\epsilon_1\epsilon_2+\epsilon_2\epsilon_3+\epsilon_3\epsilon_1)$. ∎

4 Linear-consistency tests over $\mathbb{F}_2$

In this section we consider the collection of affine functions and homomorphisms from $\mathbb{F}_2^n$ to $\mathbb{F}_2$. The results obtained are stronger in that it shows that any triple of functions that are accepted by the linear-consistency tests with non-trivial probability[2] are non-trivially close to a triple of linear-consistent functions.

For the purposes of this section it is better to think of the elements of $\mathbb{F}_2$ as $\{+1, -1\}$ and we denote a typical element of $\mathbb{F}_2^n$ by $\vec{x} \stackrel{\triangle}{=} (x_1, x_2 \ldots x_n)$ where $x_i \in \mathbb{F}_2$. Multiplication (over the reals) replaces addition modulo two in this representation. The set of homomorphisms $\textsc{Hom}_n$ mapping $\{+1, -1\}^n \to \{+1, -1\}$ is given by $\textsc{Hom}_n = \{\ell_\alpha | \alpha \subseteq [n]\}$, where $\ell_\alpha(\vec{x}) = \prod_{i \in \alpha} x_i$. The set of affine

[2]Since a triple of random functions would pass the linear-consistency tests with probability $\frac{1}{2}$, we consider the passing probability to be non-trivial if it is strictly larger than $\frac{1}{2}$.

functions is given by $\textsc{Aff}_n = \{\ell_\alpha | \alpha \subseteq [n]\} \cup \{-\ell_\alpha | \alpha \subseteq [n]\}$. The homomorphisms now satisfy $\ell_\alpha(\vec{x})\ell_\alpha(\vec{y}) = \ell_\alpha(\vec{x} \cdot \vec{y})$, where $\vec{x} \cdot \vec{y}$ represents the coordinate-wise product of the two vectors.

Let $\langle f, g \rangle$, the inner product between $f, g : \{+1, -1\}^n \to \{+1, -1\}$, be given by

$$\langle f, g \rangle = \frac{1}{2^n} \sum_{\vec{x} \in \{+1,-1\}^n} f(\vec{x}) g(\vec{x}).$$

Then $\langle \ell_\alpha, \ell_\alpha \rangle = 1$ and $\langle \ell_\alpha, \ell_\beta \rangle = 0$ if $\alpha \neq \beta$. The homomorphisms form a orthonormal basis over the reals for the set of functions from $\{+1, -1\}^n \to \mathbb{R}$. I.e. every function $f : \{+1, -1\}^n \to \mathbb{R}$ is given by $f(\vec{x}) = \sum_{\alpha \subseteq [n]} \hat{f}_\alpha \ell_\alpha(\vec{x})$, where $\hat{f}_\alpha = \langle f, \ell_\alpha \rangle$ is the α-th Fourier coefficient of f. It is easily verified that the following (Parseval's identity) holds: $\langle f, f \rangle = \sum_{\alpha \subseteq [n]} \hat{f}_\alpha^2$. For functions $f : \{+1, -1\}^n \to \{+1, -1\}$, $\langle f, f \rangle = 1$. The Fourier coefficients are of interest due to the following easily verified fact.

Proposition 11 *For every function $f : \{+1, -1\}^n \to \{+1, -1\}$:*

- $\epsilon_{\textsc{Hom}}(f) \stackrel{\triangle}{=} \min_{\alpha \subseteq [n]} \{d(f, \ell_\alpha)\} = \min_{\alpha \subseteq [n]} \{\frac{1 - \hat{f}_\alpha}{2}\}$.
- $\epsilon_{\textsc{Aff}}(f) \stackrel{\triangle}{=} \min_{g \in \textsc{Aff}_n} \{d(f, g)\} = \min_{\alpha \subseteq [n]} \{\frac{1 - |\hat{f}_\alpha|}{2}\}$.

Our result is the following:

Theorem 12 *Given functions $f_i : \{+1, -1\}^n \to \{+1, -1\}$, for $i \in \{1, 2, 3\}$, such that*

$$\Pr_{\vec{x}, \vec{y}}[f_1(\vec{x}) f_2(\vec{y}) \neq f_3(\vec{x} \cdot \vec{y})] = \delta,$$

for every $i \in \{1, 2, 3\}$, $\epsilon_{\textsc{Aff}}(f_i) \leq \delta$. Furthermore, there exists a triple of linear-consistent functions g_1, g_2, g_3 such that for every $i \in \{1, 2, 3\}$, $d(f_i, g_i) \leq \frac{1}{2} - \frac{2\gamma}{3}$, where $\gamma = \frac{1}{2} - \delta$.

Remark 13 *Notice that even when $G = \mathbb{F}_2^n$ and $H = \mathbb{F}_2$, Theorem 12 does not subsume Theorem 2. In particular the error bounds given by Theorem 2 are stronger, when $\delta < 2/9$. However for $\delta > 2/9$, and in particular for $\delta \to \frac{1}{2}$, Theorem 12 is much stronger.*

Proof: Let $\hat{f}_{i,\alpha}$ be the Fourier coefficient corresponding to the character ℓ_α of f_i. For the first part it suffices, by Proposition 11, to show that for every $i \in \{1,2,3\}$ $\max_\alpha\{|\hat{f}_{i,\alpha}|\} \geq 1-2\delta$. For the second part notice that the linear consistent functions g_1, g_2, g_3 are given by some homomorphism ℓ_α and $b_1, b_2, b_3 \in \{+1,-1\}$ satisfying $b_1 b_2 b_3 = 1$ such that $g_i(\vec{x}) = b_i \ell_\alpha(\vec{x})$. Thus our task may be rephrased as saying that we wish to show there exists an α such that $\{\min\{|\hat{f}_{1,\alpha}|, |\hat{f}_{2,\alpha}|, |\hat{f}_{3,\alpha}|\}\} \geq \frac{2(1-2\delta)}{3}$ (which captures the distance property) and $\hat{f}_{1,\alpha}.\hat{f}_{2,\alpha}.\hat{f}_{3,\alpha} \geq 0$ (which captures the property that $b_1 b_2 b_3 = 1$).

We proceed as in [4]. We first express the event that the test rejects algebraically. Let $I_{\vec{x},\vec{y}}$ be 1 if $f_1(\vec{x}) f_2(\vec{y}) \neq f_3(\vec{x} \cdot \vec{y})$ and 0 otherwise. Then

$$I_{\vec{x},\vec{y}} = \frac{1}{2}\left(1 - f_1(\vec{x}) f_2(\vec{y}) f_3(\vec{x} \cdot \vec{y})\right).$$

Since the rejection probability of the linear-consistency test is simply the expected value of $I_{\vec{x},\vec{y}}$, we get:

$$\delta = \mathbf{E}_{\vec{x},\vec{y} \in_R \{+1,-1\}^n}\left[\frac{1}{2}\left(1 - f_1(\vec{x}) f_2(\vec{y}) f_3(\vec{x} \cdot \vec{y})\right)\right].$$

Expressing the f_i's in terms of their Fourier basis we simplify the inner expression above.

$$\begin{aligned} 1 - 2\delta &= \mathbf{E}_{\vec{x},\vec{y} \in_R \{+1,-1\}^n}\left[f_1(\vec{x}) f_2(\vec{y}) f_3(\vec{x} \cdot \vec{y})\right] \\ &= \mathbf{E}_{\vec{x},\vec{y} \in_R \{+1,-1\}^n}\left[\sum_{\alpha \subseteq [n]} \hat{f}_{1,\alpha} \ell_\alpha(\vec{x}) \sum_{\beta \subseteq [n]} \hat{f}_{2,\beta} \ell_\beta(\vec{y}) \sum_{\gamma \subseteq [n]} \hat{f}_{3,\gamma} \ell_\gamma(\vec{x} \cdot \vec{y})\right] \end{aligned}$$

$$
\begin{aligned}
&= \sum_{\alpha,\beta,\gamma\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\beta}\hat{f}_{3,\gamma}\mathbf{E}_{\vec{x},\vec{y}\in_R\{+1,-1\}^n}\left[\ell_\alpha(\vec{x})\ell_\beta(\vec{y})\ell_\gamma(\vec{x}\cdot\vec{y})\right] \\
&= \sum_{\alpha,\beta,\gamma\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\beta}\hat{f}_{3,\gamma}\left(\mathbf{E}_{\vec{x}\in_R\{+1,-1\}^n}\left[\ell_\alpha(\vec{x})\ell_\gamma(\vec{x})\right]\mathbf{E}_{\vec{y}\in_R\{+1,-1\}^n}\left[\ell_\beta(\vec{y})\ell_\gamma(\vec{y})\right]\right) \\
&= \sum_{\alpha,\beta,\gamma\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\beta}\hat{f}_{3,\gamma}\left(\langle\ell_\alpha,\ell_\gamma\rangle\langle\ell_\beta,\ell_\gamma\rangle\right) \\
&= \sum_{\alpha\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha},
\end{aligned}
$$

where the last equality is obtained by recalling that $\langle\ell_\alpha,\ell_\gamma\rangle = 0$ if $\alpha \neq \gamma$ and 1 otherwise.

For the first part, assume for contradiction that $\max_\alpha\{\hat{f}_{1,\alpha}\} < 1-2\delta$. Then we get:

$$
\begin{aligned}
1-2\delta &= \sum_{\alpha\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} \\
&\leq \sum_{\alpha\subseteq[n]} |\hat{f}_{1,\alpha}||\hat{f}_{2,\alpha}||\hat{f}_{3,\alpha}| \\
&\leq \max_\alpha\{|\hat{f}_{1,\alpha}|\}\sum_{\alpha\subseteq[n]} |\hat{f}_{2,\alpha}||\hat{f}_{3,\alpha}| \\
&< (1-2\delta)\sum_{\alpha\subseteq[n]} |\hat{f}_{2,\alpha}||\hat{f}_{3,\alpha}| \\
&\leq (1-2\delta)\sum_{\alpha\subseteq[n]} \frac{\hat{f}_{2,\alpha}^2+\hat{f}_{3,\alpha}^2}{2} \\
&= 1-2\delta. \qquad \text{(Using Parseval's Identity)}
\end{aligned}
$$

The next to last inequality follows from the fact that the geometric mean is smaller than the arithmetic mean. ¿From the above contradiction the first part of the theorem follows.

Now to see the second part, assume for contradiction that for every α, either $\hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} < 0$ or there exists an i, s.t. $|\hat{f}_{i,\alpha}| < \frac{2(1-2\delta)}{3}$.

Let $S_0 = \{\alpha|\hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} < 0\}$ and for $i \in \{1,2,3\}$, let $S_i = \{\alpha||\hat{f}_{i,\alpha}| < \frac{2(1-\delta)}{3}, \alpha \notin S_j \text{ for any } j < i\}$. By definition the sets S_i are disjoint. Furthermore, by assumption, $S_0 \cup S_1 \cup S_2 \cup S_3 = 2^{[n]}$. Thus,

the following sequence of inequalities leads to a contradiction.

$$\begin{aligned}
1-2\delta &= \sum_{\alpha\subseteq[n]} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} \\
&= \sum_{\alpha\subseteq S_0} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} + \sum_{\alpha\subseteq S_1} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} + \sum_{\alpha\subseteq S_2} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} + \sum_{\alpha\subseteq S_3} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha} \\
&< 0 + \frac{2(1-2\delta)}{3}\left(\sum_{\alpha\subseteq S_1} |\hat{f}_{2,\alpha}\hat{f}_{3,\alpha}| + \sum_{\alpha\subseteq S_2} |\hat{f}_{1,\alpha}\hat{f}_{3,\alpha}| + \sum_{\alpha\subseteq S_3} |\hat{f}_{1,\alpha}\hat{f}_{2,\alpha}|\right) \\
&\leq \frac{2(1-2\delta)}{3}\left(\sum_{\alpha\subseteq S_1} \frac{\hat{f}^2_{2,\alpha}+\hat{f}^2_{3,\alpha}}{2} + \sum_{\alpha\subseteq S_2} \frac{\hat{f}^2_{1,\alpha}+\hat{f}^2_{3,\alpha}}{2} + \sum_{\alpha\subseteq S_3} \frac{\hat{f}^2_{1,\alpha}+\hat{f}^2_{2,\alpha}}{2}\right) \\
&\leq \frac{(1-2\delta)}{3}\left(\sum_{\alpha} \hat{f}^2_{1,\alpha}+\hat{f}^2_{2,\alpha}+\hat{f}^2_{3,\alpha}\right) \\
&\leq 1-2\delta.
\end{aligned}$$

This contradiction completes the proof of the second part. ∎

5 3-prover 1-bit proof systems

We first recall the definition of an MIP proof system. For integers p, a and function $r : \mathbb{Z}^+ \to \mathbb{Z}^+$, an MIP verifier V is (r, p, a) restricted if on input $x \in \{0,1\}^n$, V tosses $r(n)$ coins and issues p queries $q_1, \ldots, q_p$ to p-provers $P_1, \ldots, P_p$ and receives a bit responses $a_1, \ldots, a_p$ from the p provers. The prover P_i is thus a function mapping q_i to some a bit string a_i. The verifier then outputs a Boolean verdict accept/reject based on x, its random coins and the responses $a_1, \ldots, a_p$. An (r, p, a)-restricted MIP verifier V achieves completeness c and soundness s for a language L if for every $x \in L$ there exists a collection of p provers that force the V to accept with probability at least c, while for $x \notin L$ no tuple of p provers can make V accept with probability greater than s. $\mathrm{MIP}_{c,s}[r, p, a]$ is the collection of all languages L that have (r, p, a) restricted MIP verifiers achieving

completeness c and soundness s.

We prove the following containment for NP.

Theorem 14 *For every $\epsilon > 0$, $NP = MIP_{1-\epsilon,\frac{1}{2}+\epsilon}[O(\log n), 3, 1]$.*

Remark 15

1. *To obtain the equality $NP = MIP_{1-\epsilon,\frac{1}{2}}[O(\log n), 3, 1]$ as stated in the introduction we apply Theorem 14 with the parameter $\epsilon/3$, and then change the verifier to reject with probability $2\epsilon/3$ without looking at the proof. This gives a proof system with completeness at least $(1 - 2\epsilon/3)(1 - \epsilon/3) \geq 1 - \epsilon$ and soundness at most $(1 - 2\epsilon/3)(1/2 + \epsilon/3) \leq 1/2$.*

2. *Zwick [15] proved that for non-adaptive PCPs reading three bits, if $c/s > 2$ only languages in P can be accepted. The result extends to the case of adaptive PCPs using an earlier reduction of Trevisan [13] from adaptive to non-adaptive PCPs. Since a PCP proof system is more powerful than an MIP proof system (for the same choice of parameters), the same lower bound also applies in our situation showing that our result is essentially tight.*

Our verifier and analysis are simple variants of the verifier and analysis of Håstad [9]. We use here the formalism 'inner verifier" of Trevisan [14].

Definition 16 *A $(r, 3, 1)$-good MIP inner-verifier system consists of an $(r, 3, 1)$-restricted MIP verifier V_{inner} (for some function r); 3 encoding functions E_1, E_2 and E_3; and two (probabilistic) decoding functions D_1 and D_2. An inner-verifier system is* good, *if for every $\epsilon > 0$ there exists a $\gamma > 0$ such for every pair of positive integers m, n, the following hold:*

Completeness *If $a \in [n]$, $b \in [m]$ and $\pi : [m] \to [n]$ satisfy $\pi(b) = a$ then V_{inner}, on input (m, n, π, ϵ) accepts the provers $P_1 = E_1(a)$, $P_2 = E_2(b)$, and $P_3 = E_3(b)$ with probability at least $1 - \epsilon$.*

Soundness *For each P_1, P_2, P_3, $D_2(P_2, P_3) \in [m]$ and $D_1(P_1) \in [n]$. If V_{inner} on input (m, n, π, ϵ) accepts provers P_1, P_2, P_3 with probability $\frac{1}{2} + \epsilon$, then $\pi(D_2(P_2, P_3)) = D_1(P_1)$ with probability at least γ (over the coin tosses of the decoding procedures D_1 and D_2).*

To get the intuition of this definition, one should think of a and b as long answers given by provers in a two-prover protocol. The purpose of the inner verifier is to transform the reading of all of a and b to a much more efficient procedure by interacting with the three provers. The encoding function gives the procedure how to transform answers by provers in the two-prover protocol to provers in this new protocol and the decoding functions do the translation in the other direction. The function π captures the acceptance condition in the two-prover protocol.

For the readers more familiar with [9] we point out that n codes all assignments on the set U, m the assignments on W satisfying the chosen clauses and each of the encoding functions E_i is the long code of [5]. For readers not familiar with either [14] or [9] these notions are defined in the proof of Lemma 17 below.

The following lemma is a standard application of the paradigm of recursive proof composition [2], applied to the state-of-the-art constructions of 2-prover proof systems [11] together with the formalism of our inner verifier. It is the same construction that is used in [9] but since the formalism used here is different we also sketch the proof.

Lemma 17 *If there exists a $(O(\log n), 3, 1)$-good inner-verifier system then, for every $\epsilon > 0$, $NP = MIP_{1-\epsilon, \frac{1}{2}+\epsilon}[O(\log n), 3, 1]$.*

Proof [Sketch]: We first use the result of [1] to observe that it suffices to obtain a 3-prover 1-bit proof system verifying satisfiability of a 3-CNF formula φ, under the promise that either φ is satisfiable, or that no assignment satisfies more than a c-fraction of the clauses of φ, for some

$c < 1$. We first create a $V_{2\text{ip}}$ for a 2-prover constant-bit verifier $V_{2\text{ip}}$ for this (promise) problem as follows: For a constant u to be chosen shortly, $V_{2\text{ip}}$ picks a set of u random clauses of φ and let W be the set of variables appearing in these clauses. The verifier then picks a set U of u variables by picking one variable at random from each chosen clause. The set U is sent to the first prover and the set W to the second. The two provers respond with assignments of the variables in the two sets and $V_{2\text{ip}}$ accepts iff the assignments are consistent on U and the picked clauses are satisfied. We clearly have perfect completeness, i.e., if φ is satisfiable, then there exist provers that are always accepted by $V_{2\text{ip}}$. Using [11] it follows that the soundness is at most c_1^u for some $c_1 < 1$.

The inner-verifier system is designed to reduce the query complexity of the verification of $V_{2\text{ip}}$. Given a $(r, 3, 1)$-good MIP inner-verifier V_{inner}, we compose $V_{2\text{ip}}$ with V_{inner} to obtain V_{comp}. For each set U and W we have tables as follows. We let $n = 2^u$ where each element corresponds to an assignment on U and let m be the number of assignments on W that satisfies the picked clauses and we number these in some arbitrary way to get a correspondence between such assignments on W and $[m]$. The function π is defined as the natural projection of assignments.

The composed verifier V_{comp} interacts with three provers P_{I}, P_{II} and P_{III} where P_{I} is supposed to, for each U, provide an encoding of an assignment on U while P_{II} and P_{III} are supposed to provide encodings of assignments on W for every set W. Given φ, V_{comp} picks sets U and W as above and then let us consider $P_1(\cdot) = P_{\text{I}}(U, \cdot)$, $P_2(\cdot) = P_{\text{II}}(W, \cdot)$, and $P_3(\cdot) = P_{\text{III}}(W, \cdot)$ as three provers for V_{inner}. V_{comp} invokes V_{inner} on input $(n, m, \pi, \epsilon/2)$ with oracles P_1, P_2 and P_3, accepting iff V_{inner} does.

The completeness follows immediately (by completeness of $V_{2\text{ip}}$ and V_{inner}). To see the soundness, we claim that if $P_{\text{I}}, P_{\text{II}}, P_{\text{III}}$ are accepted by V_{comp} with probability $\frac{1}{2} + \epsilon$, then the pair of provers P_A, P_B given by $P_A(U) = D_1(P_1(U, \cdot))$ and $P_B(W) = D_2(P_{\text{II}}(W, \cdot), P_{\text{III}}(W, \cdot))$ are accepted by $V_{2\text{ip}}$

with probability at least $\frac{\epsilon\gamma}{2}$. The lemma follows from this claim by setting u s.t. $c_1^u < \frac{\epsilon\gamma}{2}$.

To verify the claim, we first apply Markov's inequality to observe that for at least a $\frac{\epsilon}{2}$-fraction of choices of U, W, the invocation of V_{inner} accepts P_1, P_2, P_3 with probability at least $\frac{1}{2} + \frac{\epsilon}{2}$. For all such choices consider the output of D_1 and D_2. Since these are assignments on U and W these are legitimate answers of P_A and P_B. By the definition of m, the outputs of P_B always satisfy the chosen clauses. Finally by the definition of π whenever $\pi(D_2(P_2, P_3)) = D_1(P_1)$ the answers are consistent on U and hence V_{2ip} accepts. This completes the proof. ∎

Proof (of Theorem 14): By Lemma 17 it suffices to establish the existence of a $(O(\log n), 3, 1)$-good inner-verifier system. We describe the three components of the inner-verifier system in order; and then analyze the system.

The inner verifier Given (n, m, π, ϵ), V_{inner} picks three functions $f : [n] \to \{+1, -1\}$, $g : [m] \to \{+1, -1\}$ and $\eta : [m] \to \{+1, -1\}$ such that $f(1) = g(1) = 1$ and otherwise f and g are random and unbiased while η is random with bias $1 - \epsilon$, i.e., for every input $j \in [m]$, $\eta(j)$ is 1 with probability $1 - \epsilon$ and -1 with probability ϵ, independently. Let $b = f(\pi(1))\eta(1)$ and g' be the function given by $g'(j) = bf(\pi(j))g(j)\eta(j)$. The verifier sends f to P_1, g to P_2 and g' to P_3. If the responses are $a_1, a_2, a_3 \in \{+1, -1\}$, then V_{inner} accepts if $a_1 a_2 a_3 = b$. As in [9], g' should be thought of as a perturbation (given by η) of the product of f and g. The variable b is introduced only to make sure that $g'(1) = 1$. The main difference between this verifier and that of [9] is that this verifier sends the queries g and g' to two different provers, while the verifier of [9] sent it to a (single) oracle.

For the sake of the analysis it will be cleaner to use an alternate description of the above verifier. For this description, notice first that $f : [n] \to \{+1, -1\}$ may also be viewed as a vector $f \in \{+1, -1\}^n$. Thus P_1 may be viewed as a function from $\{+1, -1\}^n$ to $\{+1, -1\}$. Actually, P_1 (resp. P_2, P_3) is never queried with any function f with $f(1) = -1$, but extending P_1 to be defined also for such f

by setting $P_1(f) \stackrel{\triangle}{=} -P_1(-f)$ whenever $f(1) = -1$ gives a more symmetric situation and makes the situation more similar to that in [9]. Thus we assume from now on that the functions P_1, P_2 and P_3 are defined for all inputs and preserve negation. We may now think of V_{inner} as if it picks f and g totally at random, η as before and lets g' be the function $g'(j) = f(\pi(j))g(j)\eta(j)$. It sends f to P_1, g to P_2 and g' to P_3 and accepts iff $P_1(f)P_2(g)P_3(g') = 1$.

It is easy to check that this yields exactly the same protocol as described above. The only reason for our slightly more complicated description is that enables us to assume that $P_i(h) = -P_i(-h)$ for any i and h.

Encoding The encoding functions are just the "long codes" (see [5, 9, 14]). I.e., $E_1(a)$ is the function P_1 that on input $f : [n] \to \{+1, -1\}$ responds with $f(a)$, while $E_2(b)$ (as also $E_3(b)$) is the function P_2 that on input $g : [m] \to \{+1, -1\}$ responds with $g(b)$. The completeness of the protocol follows immediately.

Decoding The decoding function D_1 is from [9, 14]. The decoding function is based on the Fourier coefficients of the functions P_i where we use

$$P_i(h) = \sum_{\alpha} \hat{P}_{i,\alpha} \ell_\alpha(h).$$

$D_1(P_1)$ works as follows: Pick $\alpha \subseteq [n]$ with probability $\hat{P}^2_{1,\alpha}$, and output a random element of α. Note that α is never empty, since $\hat{P}_{1,\emptyset} = 0$ for any function P_1 satisfying $P_1(f) = -P_1(-f)$.

The new element of our proof is the decoding function D_2. $D_2(P_2, P_3)$ works as follows: Pick $\beta \subseteq [m]$ with probability $|\hat{P}_{2,\beta} \cdot \hat{P}_{3,\beta}|$ and output a random element of β. Notice that the probabilities of picking the sets β add up to at most 1. This is true since by the inequality between the geometric

and arithmetic mean

$$\sum_{\beta} |\hat{P}_{2,\beta} \cdot \hat{P}_{3,\beta}| \leq \sum_{\beta} \frac{\hat{P}_{2,\beta}^2 + \hat{P}_{3,\beta}^2}{2} \leq 1.$$

If the sum of the probabilities is less than 1 we do nothing in the remaining case.

Analysis We now relate the performance of these decoding functions with the acceptance probability of the inner verifier V_{inner}. First we express the latter quantity in terms of the Fourier expansions of the functions P_i.

The fact that V_{inner} accepts with probability $1/2 + \epsilon$ implies that

$$\begin{aligned}
2\epsilon &= \mathbf{E}_{f,g,\eta}[P_1(f)P_2(g)P_3(g')] \\
&= \mathbf{E}_{f,g,\eta}\left[\sum_{\alpha\subseteq[n]} \hat{P}_{1,\alpha}\ell_\alpha(f) \sum_{\beta\subseteq[m]} \hat{P}_{2,\beta}\ell_\beta(g) \sum_{\beta'\subseteq[m]} \hat{P}_{3,\beta'}\ell_{\beta'}(g')\right] \\
&= \sum_{\alpha,\beta,\beta'} \hat{P}_{1,\alpha}\hat{P}_{2,\beta}\hat{P}_{3,\beta'}\mathbf{E}_{f,g,\eta}\left[\ell_\alpha(f)\ell_\beta(g)\ell_{\beta'}(\eta g f(\pi))\right] \\
&= \sum_{\alpha,\beta,\beta'} \hat{P}_{1,\alpha}\hat{P}_{2,\beta}\hat{P}_{3,\beta'}\mathbf{E}_f\left[\ell_\alpha(f)\ell_{\beta'}(f(\pi))\right]\mathbf{E}_g\left[\ell_\beta(g)\ell_{\beta'}(g)\right]\mathbf{E}_\eta\left[\ell_{\beta'}(\eta)\right].
\end{aligned}$$

Clearly the second expected value is 0 unless $\beta = \beta'$ in which case it is 1. The third expected value is, by a small calculation, seen to be $(1-2\epsilon)^{|\beta'|}$. Finally the first expected value is 0 unless it is the case that each $a \in \alpha$ has an odd number of $b \in \beta'$ such that $\pi(b) = a$ while for each $a \notin \alpha$ this number is even. We denote this condition by $\pi_2(\beta') = \alpha$ since it is naturally a "mod 2" extension of π to sets. Summing up, we have

$$2\epsilon = \sum_{\beta} \hat{P}_{1,\pi_2(\beta)}\hat{P}_{2,\beta}\hat{P}_{3,\beta}(1-2\epsilon)^{|\beta|}$$

The partial sum over all β with $|\hat{P}_{1,\pi_2(\beta)}| < \epsilon$ is at most ϵ. So, we conclude that

$$\epsilon \leq \sum_{\beta \text{ s.t. } |\hat{P}_{1,\pi_2(\beta)}|>\epsilon} \hat{P}_{1,\pi_2(\beta)}\hat{P}_{2,\beta}\hat{P}_{3,\beta}(1-2\epsilon)^{|\beta|} \leq \epsilon^{-1}\sum_{\beta} \hat{P}^2_{1,\pi_2(\beta)}\hat{P}_{2,\beta}\hat{P}_{3,\beta}(1-2\epsilon)^{|\beta|}. \tag{5}$$

Let us now estimate the probability that $D_1(P_1) = \pi(D_2(P_2, P_3))$ when D_1 and D_2 are defined as above. We claim that whenever D_2 chooses β and D_1 chooses $\pi_2(\beta)$ then probability that we get $\pi(b) = a$ is at least $1/|\beta|$. This is true since for any choice of D_1 of an element $a \in \pi_2(\beta)$ there is at least one $b \in \beta$ such that $\pi(b) = a$. The probability that D_2 chooses this element is at least $1/|\beta|$. Now note that the probability that D_1 chooses $\pi_2(\beta)$ is $\hat{P}^2_{1,\pi_2(\beta)}$ and the probability that D_2 chooses β is $|\hat{P}_{2,\beta}\hat{P}_{3,\beta}|$ and thus, by the above argument, we have

$$\begin{aligned}
\Pr[D_1(P_1) = \pi(D_2(P_2, P_3))] &\geq \sum_{\beta} \hat{P}^2_{1,\pi_2(\beta)}\hat{P}_{2,\beta}\hat{P}_{3,\beta}|\beta|^{-1} \\
&\geq \epsilon \sum_{\beta} \hat{P}^2_{1,\pi_2(\beta)}\hat{P}_{2,\beta}\hat{P}_{3,\beta}(1-2\epsilon)^{|\beta|} \\
&\qquad \text{(Using } x^{-1} \geq \epsilon(1-2\epsilon)^x \text{ for any } x \geq 1,\ \epsilon \geq 0) \\
&\geq \epsilon^3 \quad \text{(Using (5))}
\end{aligned}$$

Thus setting $\gamma = \epsilon^3$ we have established the desired properties of V_{inner}.

Now we just note that Theorem 14 follows from Lemma 17 and the constructed inner verifier. ∎

Acknowledgments

We would like to thank the reviewers of RANDOM'99 as well as the referee of the current paper for numerous comments and corrections.

References

[1] S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501-555, 1998.

[2] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM,* 45(1):70-122, 1998.

[3] Y. Aumann and M. O. Rabin. Manuscript. 1999.

[4] M. Bellare, D. Coppersmith, J. Håstad, M. Kiwi and M. Sudan. Linearity testing in characteristic two. *IEEE Transactions on Information Theory,* 42(6): 1781-1795, 1996.

[5] M. Bellare, O. Goldreich and M. Sudan. Free bits, PCPs, and non-approximability – towards tight results. *SIAM Journal on Computing,* 27(3):804-915, 1998.

[6] M. Bellare, S. Goldwasser, C. Lund and A. Russell. Efficient probabilistically checkable proofs and applications to approximation. *Proceedings of the Twenty-Fifth Annual ACM Symposium on the Theory of Computing,* pages 294-304, San Diego, California, 16-18 May 1993.

[7] M. Blum and S. Kannan. Designing programs that check their work. *Journal of the ACM,* 42(1):269-291, 1995.

[8] M. Blum, M. Luby and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences,* 47(3):549-595, 1993.

[9] J. Håstad. Some optimal inapproximability results. *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing,* pages 1-10, El Paso, Texas, 4-6 May 1997. Complete version accepted for publication in *Journal of ACM.*

[10] J. Håstad and A. Wigerson. Simple analysis of graph tests. Manuscript. December 2000.

[11] R. Raz. A parallel repetition theorem. *SIAM Journal on Computing,* 27(3):763-803, 1998.

[12] A. Samorodnitsky and L. Trevisan. A PCP characterization of NP with optimal amortized query complexity. *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 181-190, Portland, Oregon, 21-23 May 2000.

[13] L. Trevisan. Positive linear programming, parallel approximation, and PCP's. *Proceedings of the 4th European Symposium on Algorithms*, pages 62–75. LNCS 1136, Springer-Verlag, 1996.

[14] L. Trevisan. Recycling queries in PCPs and in linearity tests. *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing,* pages 299-308, Dallas, Texas, 23-26 May 1998.

[15] U. Zwick. Approximating algorithms for constraint satisfaction problems involving at most three variables per constraint. *Proceedings of the ninth ACM-SIAM Symposium on Discrete Algorithms,* 1998.