Linear Consistency Testing

Yonatan Aumann*, Johan Håstad**, Michael O. Rabin***, and Madhu Sudan†

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. We give an application of this problem (and of our results): A (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 $x, y \in G$,

* 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 an NSF Career Award.

$f(x) + f(y) = f(x + y)$. 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 of [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 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 $i \in \{1, 2, 3\}$, $\Pr_{x \in \mathbb{F}_2}[f_i(x) \neq \tilde{f}_i(x)] \leq \frac{1}{2} - \frac{2\gamma}{3}$.

[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 this is equivalent we set $f_3'(z) = -f_3(-z)$.

Motivation: We believe that the linear-consistency test is a natural variant of the linearity test and will potentially find similar applications in general. In fact, 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 appear naturally, and where we do have a concrete application, is the study of "multiple-prover one-round proof systems for NP".

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 $\mathrm{NP} = \mathrm{PCP}_{1-\epsilon,\frac{1}{2}}[\log, 3]$, 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.

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 sketch the MIP construction.

2 Definitions

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

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

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

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

A triple of functions (f_1, f_2, f_3) is defined to be *linear-consistent* if there exists a homomorphism $\phi \in \textsc{Hom}_{G\to H}$ 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 $\phi \in \textsc{Hom}_{G\to H}$ 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 $\phi \in \textsc{Hom}_{G\to H}$. 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 analysis of Blum, Luby and Rubinfeld [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 $\Delta(f,g)$ 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 $i \in \{1,2,3\}$, $\epsilon_i \stackrel{\triangle}{=} \Delta(f_i, g_i) \leq \delta$. 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 $\textsc{Corr}^f(x;y)$ to be $f(x+y)-f(y)$. Define $\tilde{f}(x) = \textsc{Plurality}_{i\in\{1,2,3\},y\in G}\{\textsc{Corr}^{f_i}(x;y)\}$ (where $\textsc{Plurality}(S)$ for a multiset S is the most commonly occurring element in S, with ties being broken arbitrarily). For $i \in \{1,2,3\}$ and $x \in G$, let $\gamma_i(x) \stackrel{\triangle}{=} \Pr_{y\in G}\left[\tilde{f}(x) \neq \textsc{Corr}^{f_i}(x;y)\right]$. 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 $\gamma = \mathbf{E}_x[\gamma(x)]$.

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 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 $i \neq j \in \{1,2,3\}$,*

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

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

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

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 (by the union bound). 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: Let $p_\alpha = \Pr_{y\in G}[\textsc{Corr}^{f_i}(x;y) = \alpha]$ and $q_\alpha = \Pr_{y\in G}[\textsc{Corr}^{f_j}(x;y) = \alpha]$. We start by showing that $\max_{\alpha\in H}\{p_\alpha\}$ is very large. Observe that

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

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 $\Pr_{y_1,y_2}[\textsc{Corr}^{f_i}(x;y_1) = \textsc{Corr}^{f_j}(x;y_2)]$ by $\tilde{p}(1-\tilde{q})+(1-\tilde{q})\tilde{q}$. 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 $\textsc{Plurality}_y\{\textsc{Corr}^{f_i}(x;y)\}$ points to the same value for every $i \in \{1,2,3\}$; and this value is the value of $\tilde{f}(x)$. 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

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

and then using Claim 4 to lower bound the left-hand side by $1 - 2\delta$. Part (3) follows by some algebraic manipulation. Details omitted. ∎

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 omitted.

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 [Sketch]: The claim is proven by showing that there exist $i \in \{1,2,3\}$ and $u \in G$ such that none of the following are true: (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) $\tilde{f}(x+y) \neq f_i(x+u) - f_i(u-y)$. The existence of such i, u is shown by picking them at random and showing probability of (A) or (B) or (C) happening is bounded away from 1. It is easy to show that if none of the events occur, then $\tilde{f}(x) + \tilde{f}(y) = \tilde{f}(x+y)$. Details omitted. ∎

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 $\alpha_i = f_i(a_0) - \tilde{f}(a_0)$.

The second part is shown by assuming, for contradicion, that $\alpha_1 + \alpha_2 \neq \alpha_3$ and then showing that that a random choice of x, y leads to the event "$f_i(x) = \tilde{f}(x) + \alpha_i$ for every $i \in \{1,2,3\}$" with probability greater than δ. But when this event happens, the test rejects, and this contradicts the fact that the rejection probability is at most δ. Details omitted. ∎

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

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. $\gamma_1, \gamma_2, \gamma_3 \leq \delta$.

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 $\{+1, -1\}$. Thus 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 $\ell_\alpha(\boldsymbol{x}) = \prod_{i \in \alpha} x_i$. The set of affine functions is given by $\textsc{Aff}_n = \{\ell_\alpha | \alpha \subseteq [n]\} \cup \{-\ell_\alpha | \alpha \subseteq [n]\}$. The homomorphisms now preserve $\ell_\alpha(\boldsymbol{x})\ell_\alpha(\boldsymbol{y}) = \ell_\alpha(\boldsymbol{x} \cdot \boldsymbol{y})$, where $\boldsymbol{x} \cdot \boldsymbol{y}$ 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

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

Then $\langle \ell_\alpha, \ell_\alpha \rangle = 1$ and $\langle \ell_\alpha, \ell_\beta \rangle = 0$ if $\alpha \neq \beta$. Thus 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 $f(\boldsymbol{x}) = \sum_{\alpha \subseteq [n]} \hat{f}_\alpha \ell_\alpha(\boldsymbol{x})$, 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 $f : \{+1,-1\}^n \to \{+1,-1\}$:*

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

[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 $\frac{1}{2}$.

– $\epsilon_{\mathrm{AFF}}(f) \stackrel{\triangle}{=} \min_{g \in \mathrm{AFF}_n}\{\Delta(f,g)\} = \min_{\alpha \subseteq [n]}\{\frac{1-|\hat{f}_\alpha|}{2}\}$.

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*

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

for every $i \in \{1,2,3\}$, $\epsilon_{\mathrm{AFF}}(f_i) \leq \delta$. Furthermore, there exists a triple of linear-consistent functions g_1, g_2, g_3 such that for every $i \in \{1,2,3\}$, $\Delta(f_i, g_i) \leq \frac{1}{2} - \frac{2\gamma}{3}$, 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 [Sketch]: The proof is obtained by modifying a proof of [4]. We omit the details, but the main steps are as follows. By arithmetizing the acceptance condition of the test we show that the rejection probability equals

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

We then use the orthogonality of the Fourier basis to show that $1 - 2\delta = \sum_{\alpha \subseteq [n]} \hat{f}_{1,\alpha}\hat{f}_{2,\alpha}\hat{f}_{3,\alpha}$. Some algebraic manipulation, using in particular Parseval's identity, yields $\max_\alpha |\hat{f}_{i,\alpha}| \geq 1 - 2\delta$ and $\max_\alpha\{\min\{|\hat{f}_{1,\alpha}|, |\hat{f}_{2,\alpha}|, |\hat{f}_{3,\alpha}|\}\} \geq \frac{2}{3}(1-2\delta)$. Applying Proposition 11 to these two bounds yields the two conclusions of the theorem. ∎

5 3-prover 1-bit proof systems

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 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 $x \notin L$ V does not accept any tuple of p-provers 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, that is tight in that none the parameters c, s can be improved.

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

We only sketch the proof here. Our verifier and analysis are simple variants of the verifier and analysis of Håstad [9]. As is usual in many of the recent PCP constructions, we start with the powerful 2-prover 1-round proof system of Raz [10] for NP, and then apply the technique of recursive proof checking [2]. To apply this technique, we define an appropriate "inner verifier system". The main point of difference in our construction from the construction of [9] is in the inner verifier that we construct and in the "decoding procedure" used in the construction. The formalism for the inner verifier system is derived from that of Trevisan [11]. Theorem 14 follows from the existence of a good inner-verifier system.

Definition 15 *An 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 $\epsilon, \delta > 0$ 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 *If V_{inner} on input (m, n, π, ϵ) accepts provers P_1, P_2, P_3 with probability $\frac{1}{2} + \delta$, 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).*

The inner verifier V_{inner} is derived directly from [9]. 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 $f(1) = g(1) = \eta(1) = 1$ 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 $b = f(\pi(1))$ and g' be the function given by $g'(j) = bf(\pi(j))g(j)\eta(j)$. The verifier sends the questions 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 $a_1 a_2 a_3 = b$. (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.)

The encoding functions are just the "long codes" (see [5, 9, 11]). 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 follows immediately.

The decoding function D_1 is from [11], which is in turn based on [9]. To describe this decoding, we notice 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

$f(1) = -1$. Thus we may set $P_1(-f) = -P_1(f)$ for every f, without altering the acceptance probability of V_{inner}. We assume here onwards that P_1 (resp. P_2, P_3 are such functions.) The decoding function is then based on the Fourier coefficients of P_1. $D_1(P_1)$ works as follows: Pick $\alpha \subseteq [n]$ with probability $\hat{P}_{1,\alpha}^2$, and output a random element of α (α is never empty, since $\hat{P}_{1,\phi} = 0$ 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 $\alpha \subseteq [m]$ with probability $|\hat{P}_{2,\alpha} \cdot \hat{P}_{3,\alpha}|$ and output a random element of α. Notice that the probabilities of picking the sets α add up to at most 1. (If the sum is smaller, we do nothing in the remaining case.)

A proof similar to that of [9] with modifications (and, in particular, the use of the Cauchy-Schwartz inequality) as in the proof of Theorem 12 provide the analysis of the soundness condition, thus yielding Theorem 14.

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.
10. R. Raz. A parallel repetition theorem. *SIAM Journal on Computing,* 27(3):763-803, 1998.
11. L. Trevisan. Recycling queries in PCPs and in linearity tests. *STOC*, 1998.