Lecture Notes on Linearity (Group Homomorphism) Testing

Oded Goldreich[*]

April 5, 2016

Summary: These notes present a linearity tester that, on input a description of two groups G, H and oracle access to a function $f : G \to H$, queries the function at three points and satisfies the following conditions:

1. If f is a homomorphism from G to H then the tester accepts with probability 1.
2. If f is δ-far from the set of all homomorphisms from G to H, then the tester rejects with probability at least $\min(0.5\delta, 0.1666)$.

The three queries are $x, y, x+y$, where x and y are selected uniformly at random in G.

These notes are based on the work of Blum, Luby, and Rubinfeld [4], a work which pioneered the study of property testing.

1 Preliminaries

Let G and H be two groups. For simplicity, we denote by $+$ the group operation in each of these groups. A function $f : G \to H$ is called a (group) homomorphism if for every $x, y \in G$ it holds that $f(x+y) = f(x) + f(y)$.

One important special case of interest is when H is a finite field and G is a vector space over this field; that is, $G = H^m$ for some natural number m. In this case, a homomorphism f from G to H can be presented as $f(x_1, ..., x_m) = \sum_{i=1}^{m} c_i x_i$, where $x_1, ..., x_m, c_1, ..., c_m \in H$. In this case, f is a linear function over H^m, which explains why testing group homomorphism is often referred to as linearity testing.

Group homomorphisms are among the simplest and most basic classes of finite functions. They may indeed claim the title of the most natural algebraic functions. This chapter addresses the problem of testing whether a given function is a group homomorphism or is far from any group homomorphism.

2 The tester

The definition of being a homomorphism is presented as a conjunction of $|G|^2$ local conditions, where each local condition refers to the value of the function on three points. Interestingly, this

[*] Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel.

definition is robust in the sense that the fraction of satisfied local conditions can be related to the distance of the function from being a homomorphism. In other words, a tester for this property is obtained by checking a single local condition that is selected at random.

Algorithm 1 (testing whether f is a homomorphism): *Select uniformly $x, y \in G$, query f at the points $x, y, x+y$, and accept if and only if $f(x+y) = f(x) + f(y)$.*

It is clear that this tester accepts each homomorphism with probability 1, and that each non-homomorphism is rejected with positive probability. The non-obvious fact is that, in the latter case, the rejection probability is linearly related to the distance of the function from the class of all homomorphisms. We first prove a weaker lower bound on the rejection/detection probability.

Theorem 2 (a partial analysis of Algorithm 1): *Algorithm 1 is a* (one-sided error) *proximity oblivious tester with detection probability $3\delta - 6\delta^2$, where δ denotes the distance of the given function from being a homomorphism from G to H.*

The lower-bound $3\delta - 6\delta^2 = 3(1-2\delta)\cdot\delta$ increases with δ only when $\delta \in [0, 1/4]$. Furthermore, this lower-bound is useless when $\delta \geq 1/2$. Thus, an alternative lower-bound is needed when δ approaches 1/2 (or is larger than it). Such a bound is provided in Theorem 3; but, let us prove Theorem 2 first.

Proof: Suppose that h is a homomorphism closest to f (i.e., $\delta = \mathbf{Pr}_{x\in G}[f(x) \neq h(x)]$). We first observe that the rejection probability (i.e., $\mathbf{Pr}_{x,y\in G}[f(x) + f(y) \neq f(x+y)]$) is lower-bounded by

$$\mathbf{Pr}_{x,y\in G}[f(x) \neq h(x) \wedge f(y) = h(y) \wedge f(x+y) = h(x+y)] \tag{1}$$

$$+\mathbf{Pr}_{x,y\in G}[f(x) = h(x) \wedge f(y) \neq h(y) \wedge f(x+y) = h(x+y)] \tag{2}$$

$$+\mathbf{Pr}_{x,y\in G}[f(x) = h(x) \wedge f(y) = h(y) \wedge f(x+y) \neq h(x+y)], \tag{3}$$

because these three events are disjoint, whereas $f(x) + f(y) \neq f(x+y)$ mandates that f and h disagree on some point in $\{x, y, x+y\}$ (since $h(x) + h(y) = h(x+y)$).[1] We lower-bound Eq. (1), while noting that Eq. (2)&(3)) can be lower-bounded analogously.

$$\begin{aligned}
&\mathbf{Pr}_{x,y}[f(x) \neq h(x) \wedge f(y) = h(y) \wedge f(x+y) = h(x+y)] \\
&\quad= \mathbf{Pr}_{x,y}[f(x) \neq h(x)] - \mathbf{Pr}_{x,y}[f(x) \neq h(x) \ \wedge\ (f(y) \neq h(y) \vee f(x+y) \neq h(x+y))] \\
&\quad\geq \mathbf{Pr}_{x,y}[f(x) \neq h(x)] \\
&\qquad - (\mathbf{Pr}_{x,y}[f(x) \neq h(x) \wedge f(y) \neq h(y)] + \mathbf{Pr}_{x,y}[f(x) \neq h(x) \wedge f(x+y) \neq h(x+y)]) \\
&\quad= \delta - \delta^2 - \delta^2
\end{aligned}$$

where the last equality follows since x and y are independently and uniformly distributed in G (and ditto w.r.t x and $x+y$). ■

[1] **Advanced comment:** Indeed, this lower bound is typically not tight, since we ignored the event in which f and h disagree on more than one point, which may also lead to rejection. For example, if H is the two-element set with addition modulo 2, then disagreement on three points (i.e., $f(x) \neq h(x) \wedge f(y) \neq h(y) \wedge f(x+y) \neq h(x+y)$) also leads to rejection (since in this case $f(x) + f(y) - f(x+y) = h(x) + 1 + h(y) + 1 - (h(x+y) + 1) = 1$).

Theorem 3 (full analysis of Algorithm 1): *Algorithm 1 is a* (one-sided error) *proximity oblivious tester with detection probability* $\min(0.5\delta, 1/6)$*, where* δ *denotes the distance of the given function from being a homomorphism from* G *to* H*.*

Proof: Let ρ denote the probability that f is rejected by the test, and suppose that $\rho < 1/6$ (since otherwise we are done). We shall show that in this case f is 2ρ-close to some homomorphism (and $\rho \geq \delta/2$ follows).[2]

The intuition underlying the proof is that the hypothesis regarding f (i.e., that it is rejected with probability $\rho < 1/6$) implies that f can be modified (or "corrected") into a homomorphism by modifying f on relatively few values (i.e., on at most $2\rho|G|$ values). Specifically, the hypothesis that $\mathbf{Pr}_{x,y\in G}[f(x)=f(x+y)-f(y)] = 1-\rho > 5/6$ suggests that a "corrected" version of f that is determined according to the most frequent value of $f(x+y)-f(y)$, when considering all possible choices of $y \in G$, is a homomorphism that is relatively close to f. Suppose, for illustration, that f is obtained by selecting an arbitrary homomorphism h and corrupting it on relatively few points (say on less than one fourth of G). Then, the corrected version of f will equal h (since for every $x \in G$ it holds that $\mathbf{Pr}_{y\in G}[f(x+y)-f(y)=h(x+y)-h(y)] > 1/2$) and both claims hold (i.e., h is a homomorphism that is relatively close to f). Needless to say, we cannot start with the foregoing assumption[3], but should rather start from an arbitrary f that satisfies

$$\mathbf{Pr}_{x,y\in G}[f(x)=f(x+y)-f(y)] = 1-\rho > 5/6. \tag{4}$$

We now turn to the actual proof.

Define the vote of y regarding the value of f at x as $\phi_y(x) \stackrel{\rm def}{=} f(x+y)-f(y)$, and define $\phi(x)$ as the corresponding plurality vote (with ties broken arbitrarily); that is,

$$\phi(x) \stackrel{\rm def}{=} \operatorname{argmax}_{v\in H}\{|\{y\in G : \phi_y(x)=v\}|\}. \tag{5}$$

We shall show that ϕ is 2ρ-close to f, and that ϕ is a homomorphism.

Claim 3.1 (closeness): *The function* ϕ *is* 2ρ*-close to* f*.*

Proof: This is merely an averaging argument, which counts as bad any point x such that $f(x)$ disagrees with at least half of the votes (regarding the value of f at x), while noting that otherwise f agrees with ϕ on x. Specifically, denoting $B = \{x\in G : \mathbf{Pr}_{y\in G}[f(x) \neq \phi_y(x)] \geq 1/2\}$, we get

$$\begin{aligned}\rho &= \mathbf{Pr}_{x,y}[f(x) \neq f(x+y)-f(y)] \\ &= \mathbf{Pr}_{x,y}[f(x) \neq \phi_y(x)] \\ &\geq \mathbf{Pr}_x[x \in B] \cdot \min_{x\in B}\{\mathbf{Pr}_y[f(x) \neq \phi_y(x)]\} \\ &\geq \frac{|B|}{|G|} \cdot \frac{1}{2}\end{aligned}$$

which implies that $|B| \leq 2\rho \cdot |G|$. On the other hand, if $x \in G \setminus B$, then $f(x) = \phi(x)$ (since $\mathbf{Pr}_y[f(x) = \phi_y(x)] > 1/2$, whereas $\phi(x)$ equals the most frequent vote). ■

[2] Hence, either $\rho \geq 1/6$ or $\rho \geq \delta/2$, which implies $\rho \geq \min(0.5\delta, 1/6)$ as claimed.

[3] The gap between the foregoing illustration and the actual proof is reflected in the fact that the illustration refers to $\delta < 1/4$, whereas the actual proof uses $\rho < 1/6$.

Recall that $\phi(x)$ was defined to equal the most frequent vote (i.e., the most frequent $\phi_y(x)$ over all possible $y \in G$). Hence, $\phi(x)$ occurs with frequency at least $1/|H|$. Actually, we just saw (in the proof of Claim 3.1) that on at least $1-2\rho$ of the x's it holds that $\phi(x)$ is the majority value. We next show that $\phi(x)$ is much more frequent: it occurs in a strong majority (for all x's).

Teaching note: The rest of the analysis is easier to verify in the case of Abelian groups, since in this case one does not need to be careful about the order of summations.

Claim 3.2 (strong majority): *For every $x \in G$, it holds that $\mathbf{Pr}_y[\phi_y(x) = \phi(x)] \geq 1 - 2\rho$.*

Proof: Fixing x, we consider the random variable $Z_x = Z_x(y) \stackrel{\text{def}}{=} f(x+y) - f(y)$, while noting that $\phi(x)$ was defined as the most frequent value that this random variable assumes. We shall show that the collision probability of Z_x (i.e., $\sum_v \mathbf{Pr}[Z_x = v]^2$) is high, and it will follow that Z_x must assume a single value (indeed $\phi(x)$) with high probability.

Recalling that the collision probability of a random variable equals the probability that two independent copies of it assume the same value, we observe that the collision probability of Z_x equals

$$\mathbf{Pr}_{y_1,y_2}[Z_x(y_1) = Z_x(y_2)] = \mathbf{Pr}_{y_1,y_2}[f(x+y_1) - f(y_1) = f(x+y_2) - f(y_2)]. \qquad (6)$$

Call a pair (y_1, y_2) good if both $f(y_1) + f(-y_1+y_2) = f(y_2)$ and $f(x+y_1) + f(-y_1+y_2) = f(x+y_2)$ hold. (Note that $y_1 + (-y_1+y_2) = y_2$ and $(x+y_1) + (-y_1+y_2) = (x+y_2)$.) Now, on the one hand, a random pair is good with probability at least $1-2\rho$, since

$$\mathbf{Pr}_{y_1,y_2}[f(y_1) + f(-y_1+y_2) = f(y_1 + (-y_1+y_2))] = 1 - \rho$$

and

$$\mathbf{Pr}_{y_1,y_2}[f(x+y_1) + f(-y_1+y_2) = f((x+y_1) + (-y_1+y_2))] = 1 - \rho,$$

where the equalities rely on the fact that the pair $(y_1, -y_1+y_2)$ (resp., the pair $(x+y_1, -y_1+y_2)$) is uniformly distributed in G^2 when (y_1, y_2) is uniformly distributed in G^2. On the other hand, for a good (y_1, y_2), it holds that $Z_x(y_1) = Z_x(y_2)$, since

$$\begin{aligned} Z_x(y_2) &= f(x+y_2) - f(y_2) \\ &= (f(x+y_1) + f(-y_1+y_2)) - (f(y_1) + f(-y_1+y_2)) \\ &= f(x+y_1) - f(y_1) \ = \ Z_x(y_1). \end{aligned}$$

It follows that the collision probability of Z_x is lower-bounded by $1-2\rho$. Observing that $\sum_v \mathbf{Pr}[Z_x = v]^2 \leq \max_v\{\mathbf{Pr}[Z_x = v]\}$, it follows that $\mathbf{Pr}[Z_x = \phi(x)] \geq 1-2\rho$, since $\phi(x)$ is the most frequent value assigned to Z_x. ■

Claim 3.3 (ϕ is a homomorphism): *For every $x, y \in G$, it holds that $\phi(x) + \phi(y) = \phi(x+y)$.*

Proof: Fixing any $x, y \in G$, we prove that $\phi(x) + \phi(y) = \phi(x+y)$ holds by considering the somewhat fictitious expression $p_{x,y} \stackrel{\text{def}}{=} \mathbf{Pr}_{r \in G}[\phi(x) + \phi(y) \neq \phi(x+y)]$, and showing that $p_{x,y} < 1$ (and hence

$\phi(x)+\phi(y)\neq\phi(x+y)$ is false).[4] We prove that $p_{x,y}<1$, by showing that

$$p_{x,y} \;\leq\; \mathbf{Pr}_r\left[\begin{array}{l} \phi(x)\neq f(x+r)-f(r) \\ \vee\ \phi(y)\neq f(r)-f(-y+r) \\ \vee\ \phi(x+y)\neq f(x+r)-f(-y+r) \end{array}\right] \tag{7}$$

and observing that equality in all three cases implies that $\phi(x)+\phi(y)=(f(x+r)-f(r))+(f(r)-f(-y+r))=f(x+r)-f(-y+r)=\phi(x+y)$. Using Claim 3.2 (and some variable substitutions), we upper-bound the probability of each of the three events in Eq. (7) holds by $2\rho<1/3$. Details follow.

Recall that Claim 3.2 asserts that for every $z\in G$ it holds that $\mathbf{Pr}_s[\phi(z)=f(z+s)-f(s)]\geq 1-2\rho$. It follows that

$$\begin{array}{rcl} \mathbf{Pr}_r[\phi(x)\neq f(x+r)-f(r)] & \leq & 2\rho \\ \mathbf{Pr}_r[\phi(y)\neq f(r)-f(-y+r)] & = & \mathbf{Pr}_s[\phi(y)\neq f(y+s)-f(s)] \;\leq\; 2\rho \\ \mathbf{Pr}_r[\phi(x+y)\neq f(x+r)-f(-y+r)] & = & \mathbf{Pr}_s[\phi(x+y)\neq f(x+y+s)-f(s)] \;\leq\; 2\rho \end{array}$$

where in both equalities we use $s=-y+r$ (equiv., $r=y+s$). Hence, $p_{x,y}\leq 3\cdot 2\rho<1$, and the claim follows. ■

Combining Claim 3.1 and 3.3, the theorem follows. ■

Digest. The proof of Theorem 3, which provides an analysis of Algorithm 1, is based on the *self-correction paradigm* (cf. [4]). In general, this paradigm refers to functions f for which the value of f at any fixed point x can be reconstructed based on the values of f at few random points. We stress that each of these points is uniformly distributed in the function's domain, but they are not independent of one another. For example, in the proof of Theorem 3, we use the fact that, when f is close to a linear function f', the value of $f'(x)$ can be reconstructed from $\phi_y(x)=f(x+y)-f(y)$, where y is uniformly distributed in G. (Note that, in this case, $x+y$ is uniformly distributed in G, but $x+y$ depends on y, since x is fixed.) Specifically, if f is ϵ-close to the linear function f', then $\mathbf{Pr}_{y\in G}[f'(x)=\phi_y(x)]\geq 1-2\epsilon$ for every $x\in G$. We note that here self-correction is only used in the analysis of an algorithm (see the proof of Claim 3.2), whereas in other cases (see, e.g., testing the long-code [3]) it is used in the algorithm itself. Furthermore, self-correction is used for reducing worst-case to average-case (see, e.g., [6, Sec. 7.1.3] and [6, Sec. 7.2.1.1]), and some of these applications predate the emergence of property testing.

3 Chapter notes

Fixing groups G and H, for every $f:G\to H$, we denote by $\delta_{G,H}(f)$ the distance of f from the set of homomorophisms, and by $\rho_{G,H}(f)$ the probability that Algorithm 1 rejects f. Recall that Theorem 2 asserts that $\rho_{G,H}(f)\geq 3\delta_{G,H}(f)-6\delta_{G,H}(f)^2$, whereas Theorem 3 asserts that $\rho_{G,H}(f)\geq\min(0.5\delta_{G,H}(f),1/6)$. These are not the best bounds known. In particular, it is known

[4] Indeed, the definition of $p_{x,y}$ is fictitious, since the event $\phi(x)+\phi(y)\neq\phi(x+y)$ does not depend on r. In particular, $p_{x,y}\in\{0,1\}$. An alternative presentation starts with the event $E_{x,y,r}$ captured by Eq. (7) and deduces from the existence of $r\in G$ that satisfies $\neg E_{x,y,r}$ that $\phi(x)+\phi(y)=(f(x+r)-f(r))+(f(r)-f(-y+r))=f(x+r)-f(-y+r)=\phi(x+y)$.

that $\rho_{G,H}(f) \geq 2/9$ for every f such that $\delta_{G,H}(f) \geq 1/4$ (see [2, 4]). Hence, for every f it holds that $\rho_{G,H}(f) \geq \beta(\delta_{G,H}(f))$, where

$$\beta(x) \stackrel{\rm def}{=} \begin{cases} 3x - 6x^2 & \text{if } x \leq \tau \\ 2/9 & \text{if } x \geq \tau \end{cases} \tag{8}$$

and $\tau = 0.25 + \sqrt{33}/36 \approx 0.41$ is the positive root of $3x - 6x^2 = 2/9$ (cf. [2]). This bound is depicted in Figure 1. Surprisingly enough, for some groups G and H, the bound $\rho_{G,H}(f) \geq \beta(\delta_{G,H}(f))$ is tight in the sense that for every $v \in [0, 5/16]$ there exists f such that $\delta_{G,H}(f) \approx v$ and $\rho_{G,H}(f) = \beta(v) = 3\delta_{G,H}(f) - 6\delta_{G,H}(f)^2$ (cf. [2]). Hence, in these groups, the decrease of β in $[1/4, 5/16]$ represent the actual behavior of the tester: The detection probability of Algorithm 1 does not necessarily increase with the distance of the function from being homomorphic.

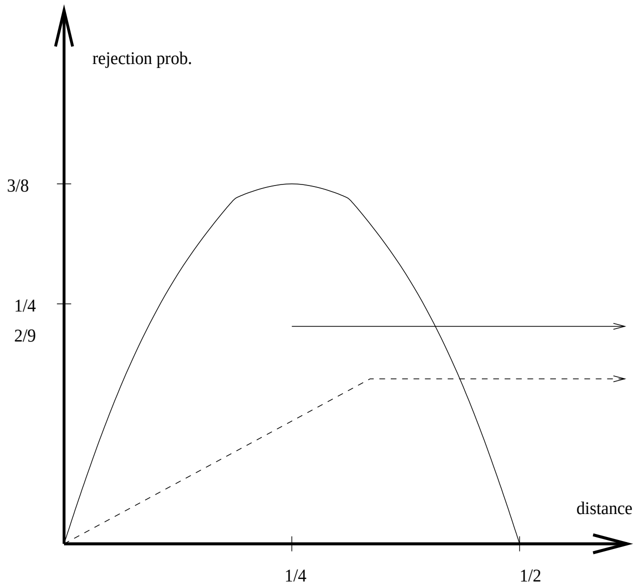

Figure 1: The lower bounds on the rejection probability of f as a function of of distance of f from a homomorphism, for general groups. The two solid lines show the bounds underlying $\beta(\cdot)$, whereas the broken dashed line shows the bound $\min(0.5x, 1/6)$.

In the special case where H is the two-element field GF(2) and $G = \mathrm{GF}(2)^m$, Bellare *et al.* [2] showed that $\rho_{G,H}(f) \geq \delta_{G,H}(f)$ and that $\rho_{G,H}(f) \geq 45/128$ for every f such that $\delta_{G,H}(f) \geq 1/4$. Thus, for every f it holds that $\rho_{G,H}(f) \geq \beta'(\delta_{G,H}(f))$, where

$$\beta'(x) \stackrel{\rm def}{=} \begin{cases} 3x - 6x^2 & \text{if } x \leq 5/16 \\ 45/128 & \text{if } x \in [5/16, 45/128] \\ x & \text{if } x \geq 45/128 \end{cases} \tag{9}$$

(This three-segment bound is depicted in Figure 2.) Furthermore, Bellare *et al.* [2] showed that the bound $\rho_{G,H}(f) \geq \beta'(\delta_{G,H}(f))$ is also tight for every value of $\rho_{G,H}(f) \in [0, 5/16]$; that is, the first segment of the bound β', which decreases in the interval $[1/4, 5/16]$ represent the actual behavior of the tester. In contrast, it is known that the bound $\rho_{G,H}(f) \geq \beta'(\delta_{G,H}(f))$ is *not* tight in the interval $(44.997/128, 0.5)$; in fact, $\rho_{G,H}(f) \geq (1 + \mathrm{poly}(1 - 2\delta_{G,H}(f)) \cdot \delta_{G,H}(f)$, where the extra term is really tiny (see [10]).[5] Still, this indicates that the known bounds used in the second and third segments of β' do not represent the actual behavior of the tester. Determining the exact behavior of $\rho_{G,H}(f)$ as a function of $\delta_{G,H}(f)$ is an open problem (even in this special case where $H = \mathrm{GF}(2)$ and $G = \mathrm{GF}(2)^m$).

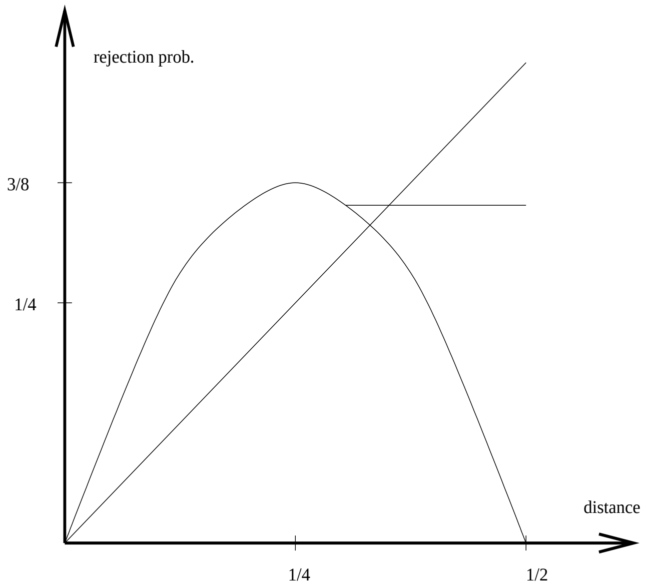

Figure 2: The lower bounds on the rejection probability of f as a function of of distance of f from a homomorphism, for $H = \mathrm{GF}(2)$ and $G = \mathrm{GF}(2)^m$.

Open Problem 4 (determining the exact behavior of Algorithm 1): *For any two groups G and H, and for every $x \in (0,1]$, what is the minimum value of $\rho_{G,H}(f)$ when taken over all $f : G \to H$ such that $\delta_{G,H}(f) = x$?*

Note that for some groups G and H, the bound $\rho_{G,H}(f) \geq \beta(\delta_{G,H}(f))$ may not be tight even for $\delta_{G,H}(f) < 5/16$.

The PCP connection. We comment that the foregoing linearity test (i.e., Algorithm 1) has played a key role in the construction of PCP systems, starting with [1]. Furthermore, a good analysis of this test was important in some of these constructions (see, e.g., [3, 8, 9][6]).

Variations. While the randomness complexity of Algorithm 1 is $2\log_2 |G|$, it is possible to show that a saving of randomness is possible (i.e., $\log_2 |G| + \log\log |H|$ bits suffice) [7].[7] A computational efficient tester was presented in [11].

A different version is considered by David *et al.* [5]. Referring to the special case where $H = \mathrm{GF}(2)$ and $G = \mathrm{GF}(2)^m$, for any $k \in [m]$, they consider functions $f : W_k \to H$, where W_k is the

[5] The additive $\mathrm{poly}(1 - 2\delta_{G,H}(f))$ term is always smaller than 0.0001.

[6] Actually, Hastad [8, 9] relies on a good analysis of the Long Code (suggested by [3]), but such an analysis would have been inconceivable without a good analysis of linearity tests (i.e., tests of the Hadamard code).

[7] This claim ignores the computational complexity of the tester. On the other hand, we note that $\log_2(|G|/q) - O(1)$ random bits are necessary for any tester that makes q queries.

set of m-dimensional Boolean vectors of weight k, and seek to test whether f agrees with a group homomorphism. That is, given oracle access to a function $f : W_k \to H$, the task is to test whether there exists a homomorphism $h : G \to H$ such that $f(x) = h(x)$ for every $x \in W_k$.

Yet another variant consists of testing affine homomorphisms (also known as translations of homomorphisms). A function $f : G \to H$ is called an affine homomorphism if there exists a group homomorphism $h : G \to H$ such that $f(x) = h(x) + f(0)$. (An equivalent definition requires that for every $x, y \in G$, it holds that $f(x + y) = f(x) - f(0) + f(y)$.)[8] Testing whether f is an affine homomorphism reduces to testing whether $h(x) \overset{\rm def}{=} f(x) - f(0)$ is a homomorphism, since if f is ϵ-far from being an affine homomorphism then h is ϵ-far from being a homomorphism.[9]

References

[1] S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy. Proof Verification and Intractability of Approximation Problems. *Journal of the ACM*, Vol. 45, pages 501–555, 1998. Preliminary version in *33rd FOCS*, 1992.

[2] M. Bellare, D. Coppersmith, J. Hastad, M.A. Kiwi, and M. Sudan. Linearity testing in characteristic two. *IEEE Transactions on Information Theory*, Vol. 42(6), pages 1781–1795, 1996.

[3] M. Bellare, O. Goldreich and M. Sudan. Free Bits, PCPs and Non-Approximability – Towards Tight Results. *SIAM Journal on Computing*, Vol. 27, No. 3, pages 804–915, 1998. Extended abstract in *36th FOCS*, 1995.

[4] M. Blum, M. Luby and R. Rubinfeld. Self-Testing/Correcting with Applications to Numerical Problems. *Journal of Computer and System Science*, Vol. 47, No. 3, pages 549–595, 1993.

[5] R. David, I. Dinur, E. Goldenberg, G. Kindler, and I. Shinkar. Direct Sum Testing. In the proceedings of the *6th ITCS*, pages 327–336, 2015.

[6] O. Goldreich. *Computational Complexity: A Conceptual Perspective.* Cambridge University Press, 2008.

[7] O. Goldreich and O. Sheffet. On The Randomness Complexity of Property Testing. *Computational Complexity*, Vol. 19 (1), pages 99–133, 2010. In *7th RANDOM*, pages 341–353, 2003.

[8] J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, Vol. 182, pages 105–142, 1999. Preliminary versions in *28th STOC* (1996) and *37th FOCS* (1996).

[9] J. Hastad. Getting optimal in-approximability results. *Journal of the ACM*, Vol. 48, pages 798–859, 2001. Extended abstract in *29th STOC*, 1997.

[8] Note that satisfying the first condition (i.e., $f(x) = h(x)+f(0)$ for sume homomorphism h) implies that $f(x+y) = h(x + y) + f(0) = h(x) + h(y) + f(0) = f(x) - f(0) + f(y)$ for all $x, y \in G$. On the other hand, if $f(x + y) = f(x) - f(0) + f(y)$ holds for all $x, y \in G$, then defining $h(x) \overset{\rm def}{=} f(x) - f(0)$ we get $h(x + y) = f(x + y) - f(0) = f(x) - f(0) + f(y) - f(0) = h(x) + h(y)$ for all $x, y \in G$.

[9] Suppose that h is ϵ-close to a homomorphism h'. Then, f is ϵ-close to f' such that $f'(x) = h'(x) + f(0)$, which means that f' is an affine homomorphism (since $f(0) = h'(0) + f(0) = f'(0)$).

[10] T. Kaufman, S. Litsyn, and N. Xie. Breaking the Epsilon-Soundness Bound of the Linearity Test over GF(2). *SIAM Journal on Computing*, Vol. 39 (5), pages 1988–2003, 2010.

[11] A. Shpilka and A. Wigderson. Derandomizing Homomorphism Testing in General Groups. *SIAM Journal on Computing*, Vol. 36 (4), pages 1215–1230, 2006.