Linearity Testing in Characteristic Two*

M. Bellare † D. Coppersmith ‡ J. Håstad § M. Kiwi ¶ M. Sudan ||

Abstract

Let $\mathsf{Dist}(f,g) = \Pr_u[\, f(u) \neq g(u)\,]$ denote the relative distance between functions f, g mapping from a group G to a group H, and let $\mathsf{Dist}(f)$ denote the minimum, over all linear functions (homomorphisms) g, of $\mathsf{Dist}(f,g)$. Given a function $f\colon G \to H$ we let $\mathsf{Err}(f) = \Pr_{u,v}[\, f(u)+f(v) \neq f(u+v)\,]$ denote the rejection probability of the BLR (Blum-Luby-Rubinfeld) linearity test. *Linearity testing* is the study of the relationship between $\mathsf{Err}(f)$ and $\mathsf{Dist}(f)$, and in particular the study of lower bounds on $\mathsf{Err}(f)$ in terms of $\mathsf{Dist}(f)$.

The case we are interested in is when the underlying groups are $G = \mathrm{GF}(2)^n$ and $H = \mathrm{GF}(2)$. In this case the collection of linear functions describe a Hadamard code of block length 2^n and for an arbitrary function f mapping $\mathrm{GF}(2)^n$ to $\mathrm{GF}(2)$ the distance $\mathsf{Dist}(f)$ measures its distance to a Hadamard code (normalized so as to be a real number between 0 and 1). The quantity $\mathsf{Err}(f)$ is a parameter that is "easy to measure" and linearity testing studies the relationship of this parameter to the distance of f.

The code and corresponding test are used in the construction of efficient probabilistically checkable proofs and thence in the derivation of hardness of approximation results. In this context, improved analyses translate into better non-approximability results. However, while several analyses of the relation of $\mathsf{Err}(f)$ to $\mathsf{Dist}(f)$ are known, none is tight.

We present a description of the relationship between $\mathsf{Err}(f)$ and $\mathsf{Dist}(f)$ which is nearly complete in all its aspects, and entirely complete (i.e. tight) in some. In particular we present functions $L, U\colon [0,1] \to [0,1]$ such that for all $x \in [0,1]$ we have $L(x) \leq \mathsf{Err}(f) \leq U(x)$ whenever $\mathsf{Dist}(f) = x$, with the upper bound being tight on the whole range, and the lower bound tight on a large part of the range and close on the rest.

Part of our strengthening is obtained by showing a new connection between the linearity testing problem and Fourier analysis, a connection which may be of independent interest. Our results are used by Bellare, Goldreich and Sudan to present the best known hardness results for Max-3SAT and other MaxSNP problems [7].

Index Terms — Probabilistically checkable proofs, approximation, program testing, Hadamard codes, error detection, linearity testing, MaxSNP.

---

* A preliminary version of this paper appeared in Proceedings *of the* 36th *Symposium on Foundations of Computer Science*, IEEE, 1995.

† Department of Computer Science & Engineering, Mail Code 0114, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093. `mihir@cs.ucsd.edu`.

‡ Research Division, IBM T.J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598, USA. `copper@watson.ibm.com`.

§ Department of Computer Science, Royal Institute of Technology, 10044 Stockholm, Sweden. `johanh@nada.kth.se`. Part of this work was done while the author was visiting MIT.

¶ Depto. de Ingeniería Matemática, Fac. de Cs. Físicas y Matemáticas, U. de Chile. `mkiwi@dim.uchile.cl`. This work was done while the author was at the Dept. of Applied Mathematics, Massachusetts Institute of Technology. Supported by AT&T Bell Laboratories PhD Scholarship, NSF Grant CCR–9503322, FONDECYT grant No. 1960849, and Fundación Andes.

|| Research Division, IBM T.J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598, USA. `madhu@watson.ibm.com`.

1 Introduction

One of the contributions of computational complexity theory has been to re-examine the classical notion of what constitutes a proof of a mathematical statement. The complexity class NP introduced the notion of an *efficiently verifiable* proof. It asks that the proof, which is a sequence of written symbols, not only be verifiable, but be verifiable quickly, namely in polynomial time. Over the last decade or so, researchers have furthered this avenue in many ways. One of the many notions that has been developed, and explored, allows the verifier (of the claimed proof) to be probabilistic in its actions. The new verifier is also allowed to err in its judgment, as long as it doesn't do so too often— proofs of false statements can be accepted with small probability. (This probability is measured over coin tosses made by the verifier and not over any distribution over theorems/proofs.) As a tradeoff, the notion restricts the access of the verifier into the proof, allowing a verifier to only *query* or probe the proof in a small number of bits and studies the behavior of the number of bits that are needed to be probed in any proof system as a function of the error probability. Such a proof system, i.e., the verifier and its associated format for valid proofs, is referred to as a probabilistically checkable proof system — PCP, for short. Along with the development of this notion, the research has also yielded a series of technical developments, which have constructed PCP verifiers which examine only a constant number of bits, C, of a purported proof and reject proofs of incorrect statements with probability $\frac{1}{2}$. This constant is a universal constant, and independent of the length of the theorem or the proof. The new proof systems, do require valid proofs to be longer than traditional (deterministic) proof systems would allow for. However the size of the new proofs are only polynomially larger than the size of the traditional proofs.

Apart from the inherent interest in the construction and performance of PCP systems, a major motivating factor for the study of PCP systems is their use in the derivation of non-approximability results for combinatorial optimization problems. The theory of NP-completeness has been employed as an important tool in the analysis of the complexity of finding *optimal* solutions to discrete (or combinatorial) optimization problems. For many optimization problems, the NP-complete or NP-hard ones, this theory can be used to show that no polynomial time solution solves this problem *exactly*, unless NP = P. However the possibility that solutions to these problems which approximate the optimum to within a relative error of ϵ for every $\epsilon > 0$ may be found in polynomial time, remained open. A new connection uses the PCP constructions mentioned above to show that for many interesting problems, even such approximate solutions can not be found in polynomial time unless NP = P. This connection further serves to motivate the study of PCP systems and in particular, their efficiency (for instance, the parameter C above) since improved efficiency translates into stronger non-approximability results.

The prime motivation for the problem to be studied in this paper is these PCP constructions and the ensuing hardness of approximation results. However, a full explanation of the details of these results is beyond the scope of this paper — in fact, we will not even attempt to formalize the definitions above. The interested reader is directed towards any of a number of surveys which have appeared on this topic.[1] Fortunately, the problem to be studied in this paper can be formulated cleanly without reference to the above mentioned results and furthermore have an interesting implication in a coding theoretic setting. We first describe this setting and then procedd to formally define the problem of interest.

Central to many of the construction of efficient PCPs has been the construction and analysis

[1] It isn't possible to provide an exhaustive list of the dozen or so surveys available but if you are on the web check out `http://www-cse.ucsd.edu/users/mihir/pcp.html`.

of error-correcting codes and probabilistic "error-detection" algorithms for these error-correcting codes. These algorithms function as follows: Given a word w which is supposed to be a codeword of some error-correcting code, the algorithm probabilistically chooses a small (sometimes constant) number of bits of the word w to examine, computes a (simple) boolean function of these bits and outputs a verdict ACCEPT/REJECT. The guarantee obtainable from such algorithms is weaker than the guarantee expected ¿from classical error-detection algorithms. In particular, the guarantees behave as follows: Given a valid codeword, the algorithm must output ACCEPT with probability 1. On the other hand if the input is far from any valid codeword (i.e., the distance is more than some specified constant fraction of the minimum distance of the code), then the algorithm must output REJECT with some positive probability, bounded away from 0. Most of the codes used in these constructions are well-known ones, with Hadamard Codes and variants of Reed-Solomon being the most commonly used ones. Much of the technical development in this area is directed towards the construction and analysis of the probabilistic error-correcting algorithms. This area of study, collectively referred to as *testing* in the PCP literature is the origin of the problem considered in this paper.

It is a feature of the area that while tests are easy to specify, they are notoriously hard to analyze, especially to analyze well. Yet, good analyses are, for several reasons, worth striving for. There is, first, the inherent mathematical interest of getting the best possible analysis and understanding of a well-defined combinatorial problem. But, there is a more pragmatic reason: better analyses typically translate into improved (increased) factors shown non-approximable in hardness of approximation results.

The specific problem considered here is called the *linearity testing* problem. We wish to look at a particular test, called the BLR test, with was the first ever proposed. Our focus is the case of most importance in applications, when the underlying function maps between groups of characteristic two. Several analyses have appeared, yet none is tight. Each improved analysis implies improved factors shown non-approximable in hardness of approximation results.

Let us begin by describing the linearity testing problem and past work more precisely.

1.1 The Problem

The linearity testing problem is a problem related to homomorphisms between groups. Let G, H be finite groups. A function $g\colon G \to H$ is said to be *linear* if $g(u) + g(v) = g(u+v)$ for all $u, v \in G$. (That is, g is a group homomorphism.) We will use the notation $u \stackrel{R}{\leftarrow} G$ to represent a random variable u chosen uniformly at random from the (finite) group G. Here are some basic definitions:

- $\textsc{Lin}(G,H)$ — Set of all linear functions of G to H
- $\mathsf{Dist}(f,g) \stackrel{\text{def}}{=} \Pr_{u \stackrel{R}{\leftarrow} G}[\, f(u) \neq g(u)\,]$ — (relative) distance between $f, g\colon G \to H$
- $\mathsf{Dist}(f) \stackrel{\text{def}}{=} \min\{\, \mathsf{Dist}(f,g) : g \in \textsc{Lin}(G,H)\,\}$ — Distance of f to its closest linear function.

The BLR Test. Blum, Luby and Rubinfeld [9] suggest a probabilistic method to "test" if a function f is really a *linear* function. This test, henceforth referred to as the BLR test, is the following [9]— Given a function $f\colon G \to H$, pick $u, v \in G$ at random and reject if $f(u) + f(v) \neq f(u+v)$. Let

$$\mathsf{Err}(f) \stackrel{\text{def}}{=} \Pr_{u,v \stackrel{R}{\leftarrow} G}[\, f(u) + f(v) \neq f(u+v)\,]$$

denote the probability that the BLR test rejects f. The issue in linearity testing is to study how $\mathsf{Err}(f)$ behaves as a function of $x = \mathsf{Dist}(f)$. In particular, one would like to derive good lower bounds on $\mathsf{Err}(f)$ as a function of x.

$\textsc{Rej}(\cdot)$. A convenient way to capture the above issues is via the *rejection probability function* $\textsc{Rej}_{G,H}\colon [0,1] \to [0,1]$ of the test. It associates to any number x the minimum value of $\mathsf{Err}(f)$, taken over all functions f of distance x from the space of linear functions. Thus,

$$\textsc{Rej}_{G,H}(x) \stackrel{\text{def}}{=} \min\{\, \mathsf{Err}(f) : f\colon G \to H \text{ s.t. } \mathsf{Dist}(f) = x \,\} .$$

The graph of $\textsc{Rej}_{G,H}$ —namely $\textsc{Rej}_{G,H}(x)$ plotted as a function of x— is called the *linearity testing curve*.[2] This curve depends only on the groups G, H.

By definition it follows that $\textsc{Rej}_{G,H}(x) > 0$ if $x > 0$. However it is not easy to see if any other quantitative statements can be made about $\textsc{Rej}_{G,H}(x) > 0$ for larger values of x. The most general problem in linearity testing is to determine the function $\textsc{Rej}_{G,H}(\cdot)$ for given G, H. Much of the work that has been done provides information about various aspects of this function.

THE KNEE OF THE CURVE. At first glance, it may be tempting to believe that $\textsc{Rej}_{G,H}(\cdot)$ will be a monotone non-decreasing function. One of the most surprising features of $\textsc{Rej}_{G,H}$ is that this is not necessarily true. It turns out (and we will see such an example presently) that there exist groups G, H such that $\textsc{Rej}_{G,H}(\frac{1}{4}) \geq \frac{3}{8}$, but $\textsc{Rej}_{G,H}(\frac{2}{3}) = \frac{2}{9}$. The threshold of $x = \frac{1}{4}$ turns out to be significant in this example and an important parameter that emerges in the study of linearity testing is how low $\textsc{Rej}_{G,H}(x)$ can be for $x \geq \frac{1}{4}$. In this paper we call this parameter, identified in [2, 6, 7, 8], the *knee* of the curve. Formally:

$$\textsc{Knee}_{G,H} \stackrel{\text{def}}{=} \min\{\, \textsc{Rej}(x) : x \geq \tfrac{1}{4} \,\} .$$

1.2 Error detection in Hadamard codes

In this paper we look at the performance of the BLR test when the underlying groups are $G = \mathrm{GF}(2)^n$ and $H = \mathrm{GF}(2)$ for some positive integer n. For notational simplicity we now drop the groups G, H from the subscripts, writing $\textsc{Rej}(x)$ and $\textsc{Knee}$— it is to be understood that we mean $G = \mathrm{GF}(2)^n$ and $H = \mathrm{GF}(2)$.

This special case is of interest because of the following reason: In this case the family of functions $\textsc{Lin}(\mathrm{GF}(2)^n, \mathrm{GF}(2))$ actually defines a Hadamard code of block length 2^n. Notice that every linear function l is specified by a vector α from $\mathrm{GF}(2)^n$ such that $l(x) = \langle \alpha, x \rangle$ (where $\langle \alpha, x \rangle = \sum_{i=1}^n \alpha_i x_i$ denotes the inner product of vectors α, x). Thus we can associate with each of the 2^n linear functions l, a codeword which is the 2^n bit sequence $(l(x) : x \in \mathrm{GF}(2)^n)$. Any two distinct codewords differ in exactly 2^{n-1} positions, making this a $(2^n, 2^n, 2^{n-1})$-code. For further details see MacWilliams and Sloane [18, pages 48–49].

For an arbitrary function f, the parameter $\mathsf{Dist}(f)$ simply measures its distance to the above mentioned Hadamard code, normalized by 2^n. Estimating $\mathsf{Dist}(f)$ is thus related to the classical task of error-detection. The parameter $\mathsf{Err}(f)$ on the other hand simply defines a quantity that can be estimated to fairly good accuracy by a probabilistic algorithm, which probes f in a few places (or reads a few bits of the purported codeword). The algorithm repeats the following step several times: It picks random $x, y \in \mathrm{GF}(2)^n$ and tests to see if $f(x) + f(y) = f(x+y)$. At the end it reports the average number of times this test fails. It can be verified easily that this provides an

[2] Actually the function $\textsc{Rej}_{G,H}(x)$ is only defined for finitely many values, namely the integral multiples of $\frac{1}{|G|}$, and undefined for infinitely many values. Thus the linearity testing curve is not really a curve in the real plane, but simply describes a function of finitely many points.

estimate on $\mathsf{Err}(f)$, and the accuracy of this estimate improves with the number of iterations. The advantage of this algorithm is that it probes f in very few places in order to compute its output (in particular the number of probes can be independent of n). The aim of Linearity Testing is to turn this estimate on $\mathsf{Err}(f)$ into an estimate on $\mathsf{Dist}(f)$. This would thus yield an algorithm which probes f in few places and yet yields some reasonable estimates on $\mathsf{Dist}(f)$, and in particular solves the earlier mentioned probabilistic error-detection task. This is the ingredient which makes this test useful in the applications to PCPs and motivates our study.

1.3 Previous work

The first investigation of the shape of the linearity testing curve, by Blum, Luby and Rubinfeld [9], was in the general context where G, H are arbitrary finite groups. Their analysis showed that $\textsc{Rej}_{G,H}(x) \geq \frac{2}{9}x$ [9]. (They indicate that this is an improvement of their original analysis obtained jointly with Coppersmith.) Interest in the tightness of the analysis begins with Bellare, Goldwasser, Lund and Russell [6] in the context of improving the performance of PCP systems. They showed that $\textsc{Rej}_{G,H}(x) \geq 3x - 6x^2$. It turns out that, with very little effort, the result of [9] can be used to show that $\textsc{Rej}_{G,H}(x) \geq \frac{2}{9}$ for $x \geq \frac{1}{3}$. This claim appears in Bellare and Sudan [8], without proof. A proof is included in the appendix of this paper, for the sake of completeness. Of the three bounds above, the last two bounds supercede the first, so that the following theorem captures the state of knowledge.

Theorem 1.1 [6, 9, 10] Let G, H be arbitrary finite groups. Then:

(1) $\textsc{Rej}_{G,H}(x) \geq 3x - 6x^2$.

(2) $\textsc{Knee}_{G,H} \geq \frac{2}{9}$.

As indicated above, an improved lower bound for the knee would lead to better PCP systems. But in this general setting, we can do no better. The following example of Coppersmith [10] shows that the above value is in fact tight in the case of general groups. Let m be divisible by three. Let f be a function from $\mathcal{Z}_m^n$ to $\mathcal{Z}_m$ such that $f(u) = 3k$, if $u_1 \in \{3k-1, 3k, 3k+1\}$. Then, $\mathsf{Dist}(f) = \frac{2}{3}$. Furthermore, $f(u) + f(v) \neq f(u+v)$ only if $u_1 = v_1 = 1 \pmod 3$, or $u_1 = v_1 = -1 \pmod 3$, i.e. $\mathsf{Err}(f) = \frac{2}{9}$.

This leads into our research. We note that the problem to which linearity testing is applied in the proof system constructions of [2, 6, 7, 8] is that of testing Hadamard codes (in the first three works) and the long code (in the last work). But this corresponds to the above problem in the special case where $G = \mathrm{GF}(2)^n$ and $H = \mathrm{GF}(2)$. (G is regarded as an additive group in the obvious way. Namely, the elements are viewed as n-bit strings or vectors over GF(2), and operations are component-wise over GF(2).) For this case, the example of Coppersmith does *not* apply, and we can hope for better results.

1.4 New results and techniques

As pointed out earlier we focus on the case where the domain and range are of characteristic two and in particular $G = \mathrm{GF}(2)^n$ and $H = \mathrm{GF}(2)$. We provide two new analyses of $\textsc{Rej}(x)$ in this case.

Fourier analysis. We establish a new connection between linearity testing and Fourier analysis. We provide an interpretation of $\mathsf{Dist}(f)$ and $\mathsf{Err}(f)$ in terms of the Fourier coefficients of an appropriate transformation of f. We use this to cast the linearity testing problem in the language of Fourier series. This enables us to use Fourier analysis to study the BLR test. The outcome is the following:

Theorem 1.2 For every real number $x \leq \frac{1}{2}$, $\textsc{Rej}(x) \geq x$.

Apart from lending a new perspective to the linearity testing problem, the result exhibits a feature which distinguishes it from all previous results. Namely, it shows that $\textsc{Rej}(x) \to \frac{1}{2}$ as $x \to \frac{1}{2}$.[3] (According to the previous analysis, namely Theorem 1.1, $\textsc{Rej}(x)$ may have been bounded above by $\frac{2}{9}$ for all $x \geq \alpha$, where α is the larger root of the equation $3z - 6z^2 = \frac{2}{9}$.) Furthermore we can show that the analysis is tight (to within $o(1)$ factors) at $x = \frac{1}{2} - o(1)$.

This result can also be combined with Part (1) of Theorem 1.1 to show that $\textsc{Knee} \geq \frac{1}{3}$. However this is not tight. So we focus next on finding the right value of the knee.

Combinatorial analysis. The analysis to find the knee is based on combinatorial techniques. It leads us to an isoperimetric problem about a 3-regular hypergraph on the vertices of the n-dimensional hypercube. We state and prove a Summation Lemma which provides a tight isoperimetric inequality for this problem. We then use it to provide the following exact value of the knee of $\textsc{Rej}(x)$.

Theorem 1.3 $\textsc{Knee} = \frac{45}{128}$.

Tightness of the analysis. We provide examples to indicate that, besides the knee value, the lower bounds on $\textsc{Rej}(x)$ as indicated by our and previous results are tight for a number of points. In particular, the curve is tight for $x \leq \frac{5}{16}$, and the bound at $x = \frac{1}{2} - o(1)$ is matched up to within $o(1)$ factors (i.e., there exist functions $f_n : \mathrm{GF}(2)^n \to \mathrm{GF}(2)$ such that as n goes to ∞, $\mathsf{Err}(f_n)$ and $\mathsf{Dist}(f_n)$ go to $\frac{1}{2}$).

Other results. The isoperimetric inequality underlying Theorem 1.3 turns out to reveal other facts about $\textsc{Rej}(x)$ as well. In particular it helps establish a tight *upper bound* on $\mathsf{Err}(f)$ as a function of $\mathsf{Dist}(f)$. This result is presented in Section 3.

Also, while the main focus of this paper has been the BLR test, we also present in Section 5 a more general result about testing for total degree one in characteristic two. The purpose is to further illustrate the strength and elegance of the Fourier analysis technique, as well as its more general applicability to the problem of analyzing program testers.

Graph. Figure 1 summarizes the results of this work. The points $\{ (\mathsf{Dist}(f), \mathsf{Err}(f)) : f \}$ lie in the white region of the first graph. The dark shaded region represents the forbidden area before our work, and the light shaded region represents what we add to the forbidden area. Note we both extend the lower bound and provide upper bounds. The dots are actual computer constructed examples; they indicate that perhaps the lower bound may be improved, but not by much.[4] In particular, the knee value is tight. Furthermore the upper bound is tight.

The second graph indicates lower bounds on $\textsc{Rej}(x)$. The line $\frac{2}{9}x$ represents the result of [9]. The parabola is the curve $3x - 6x^2$ representing the result of [6]. The curve $\frac{2}{3}x$ when $x \leq \frac{1}{3}$ and $\frac{2}{9}$ when $x > \frac{1}{3}$ represents the result of [8]. Our additions are the 45 degree line of x and the horizontal line at $\frac{45}{128}$ for the new knee value.

[3] Note that $\mathsf{Dist}(f) \leq \frac{1}{2}$ for all $f\colon G \to H$ because we are working over $\mathrm{GF}(2)$, so only the portion $x \in [0, \frac{1}{2}]$ of the curve is interesting.

[4] More precisely, we have a randomized procedure that with high probability can construct, for each plotted point, a function f such that $(\mathsf{Dist}(f), \mathsf{Err}(f))$ is arbitrarily close to the point in question.