Bounds on 2-Query Locally Testable Codes with Affine Tests[†]

Gillat Kol   Ran Raz

Abstract

We study *Locally Testable Codes* (*LTCs*) that can be tested by making two queries to the tested word using an affine test. That is, we consider LTCs over a finite field $\mathbb{F}$, with codeword testers that only use tests of the form $av_i + bv_j = c$, where v is the tested word and $a, b, c \in \mathbb{F}$.

We show that such LTCs, with high minimal distance, must be of constant size. Specifically, we show that every 2-query LTC with affine tests over $\mathbb{F}$, that has minimal distance at least $\frac{9}{10}$, completeness at least $1 - \epsilon$, and soundness at most $1 - 3\epsilon$, is of size at most $|\mathbb{F}|$.

Our main motivation in studying LTCs with affine tests is the *Unique Games Conjecture* (*UGC*), and the close connection between LTCs and PCPs. We mention that all known PCP constructions use LTCs with corresponding properties as building blocks, and that many of the LTCs used in PCP constructions are affine. Furthermore, the UGC was shown to be equivalent to the UGC with affine tests [13], thus the UGC implies the existence of a low-error 2-query PCP with affine tests.

1 Introduction

1.1 Locally Testable Codes

Locally Testable Codes (*LTCs*) are error correcting codes for which the proximity of a given word to a codeword can be probabilistically tested with good confidence. The

---

test should be done by an efficient procedure (a "*tester*") that only reads a *constant* number of locations in the given word. We say that a tester has *completeness* $1 - \epsilon$, if it accepts every codeword with probability at least $1 - \epsilon$. The tester has *soundness s*, if it accepts any word that is "far" from every codeword with probability at most s.

We mention that one of the main motivations in studying LTCs is their central role in PCP constructions.

1.1.1 Unique Locally Testable Codes

In [16] we study the Unique Games Conjecture by examining LTCs with testers that only read two locations in the tested word, and only use *unique tests*. That is, given the value read from the first queried location, there is a unique value for the second location that makes the tester accept, and vice versa.

In other words, if T is a tester for an LTC $C \subseteq \Sigma^n$, we require that for every pair of coordinates $(i, j) \in [n]^2$ that may be queried by T there exists a permutation π_{ij} over Σ for which the following holds. If $v \in \Sigma^n$ is the tested word, then after querying v_i and v_j, T accepts if and only if $v_j = \pi_{ij}(v_i)$.

1.1.2 Affine Locally Testable Codes

In this work we consider the special case of LTCs with unique tests, where the tests are also *affine*. That is, we assume that the alphabet set is a finite field $\Sigma = \mathbb{F}$, and require every permutation π_{ij} to be of the form $\pi_{ij}(x) = ax + b$, where $a, b \in \mathbb{F}$. In other words, the tests carried out by the testers are of the form $av_i + bv_j = c$, where $v \in \mathbb{F}^n$ is the tested word, $i, j \in [n]$, and $a, b, c \in \mathbb{F}$.

We next give the formal definition of an affine LTC. We use the following notations. For a natural number $t \in \mathbb{N}$, denote $[t] = \{1, ..., t\}$. Let $\mathbb{F}$ be a finite field, and let $n \in \mathbb{N}$ be a natural number. The distance between two words $u, w \in \mathbb{F}^n$ is defined as $\Delta(u, w) = \frac{1}{n}|\{i \in [n] \mid u_i \neq w_i\}|$. A subset $C \subseteq \mathbb{F}^n$ is called a code. The relative distance of the code C is $\min_{u \neq w \in C} \{\Delta(u, w)\}$, and the distance of a word $v \in \mathbb{F}^n$ from the code C is $\Delta(v, C) = \min_{u \in C} \{\Delta(v, u)\}$. Let $\mathcal{A}_{\mathbb{F}}$ be the set of all affine functions over $\mathbb{F}$, $\mathcal{A}_{\mathbb{F}} = \{f : \mathbb{F} \to \mathbb{F} \mid f(x) = ax + b, \ a, b \in \mathbb{F}, \ a \neq 0\}$.

Definition 1 (Affine Local Tester, Affine LTC). Let $\mathbb{F}$ be a finite field, $n \in \mathbb{N}$ be a natural number, and $C \subseteq \mathbb{F}^n$ be a code. Let $\epsilon, s \in [0, 1)$ be real numbers. Assume that T is a probabilistic, non-adaptive, oracle machine with access to a string $v \in \mathbb{F}^n$. In

addition, assume that T makes at most two queries to v, and outputs either `accept` or `reject`. Then, T is an (ϵ, s)-affine local tester for C if it satisfies the following conditions:

- **Affineness:** For every pair of coordinates $(i, j) \in [n]^2$ that may be queried by T in a single execution, there exists an affine function $f_{ij} \in \mathcal{A}_{\mathbb{F}}$ such that the following holds. After querying v_i and v_j, T outputs `accept` if and only if $v_j = f_{ij}(v_i)$.

  In addition, if T makes a single query to coordinate $i \in [n]$, then there exits $a \in \mathbb{F}$ such that the following holds. After querying v_i, T outputs `accept` if and only if $v_i = a$.

- **Completeness:** If $v \in C$, then $\Pr[T^v = \texttt{accept}] \geq 1 - \epsilon$.

- **Soundness:** If $\Delta(v, C) \geq \frac{1}{5}$, then $\Pr[T^v = \texttt{accept}] < s$. In other words, if $\Pr[T^v = \texttt{accept}] \geq s$, then there exists a codeword u such that $\Delta(v, u) < \frac{1}{5}$.

A code C is an (α, ϵ, s)-affine LTC if it has relative distance at least $1 - \alpha$, and has an (ϵ, s)-affine local tester.

1.2 Motivation

1.2.1 The PCP Theorem and the Unique Games Conjecture

The celebrated PCP Theorem, discovered in 1992 [3, 9, 2, 1], states that any NP membership statement (e.g., $\varphi \in SAT$) has a proof that can be probabilistically checked with good confidence. The check can be performed by an efficient procedure (a "*verifier*") that only reads a *constant* number of locations in the proof.

The PCP Theorem was a major turning point in the research of hardness of approximation. However, for some fundamental problems, optimal inapproximability results are still not known. To cope with such problems, a strengthening of the PCP Theorem, called the *Unique Games Conjecture* (*UGC*), was introduced by Subhash Khot in 2002 [12]. The conjecture, or variants of it, was shown to imply many improved inapproximability results [12, 15, 17, 7, 8, 13, 14, 18].

The UGC considers a special type of PCP verifiers: verifiers that read at most *two* locations in the proof, and only make *unique tests*. The conjecture predicts the existence of such a verifier, that errs (rejects a correct proof or accepts a false statement) with arbitrarily small probability. We mention that the UGC with affine tests is equivalent to the UGC [13].

1.2.2 The Relations Between PCPs and LTCs

As mentioned above, LTCs and PCPs are closely related. All known PCP construc-
tions use LTCs as building blocks. For example, variants of Reed-Muller, Hadamard,
and the Long Code are codes that have been extensively used in PCP constructions.
Furthermore, to obtain PCPs with certain properties, one usually uses LTCs with cor-
responding properties.

In the opposite direction, some PCP constructions were shown to imply LTCs [10].
Moreover, the existence of a special type of PCP, called *PCP of Proximity (PCPP)*,
is known to imply LTCs [4]. For further discussion of the relations between PCPs and
LTCs, see [10].

In light of the strong connection between PCPs and LTCs, in [16] we study LTCs
analogues to the UGC. Specifically, we consider LTCs with testers that only use *unique
tests*. Roughly speaking, we show that such LTCs with low error are of constant size.

In this work we consider the special case of LTCs with unique tests, where the tests
are also *affine*. Roughly speaking, we show that such LTCs are of constant size, even if
the error is a large constant. As mentioned above, the special case of UGC with affine
tests was shown to be equivalent to the UGC [13]. In addition, many of the LTCs used
in PCP constructions are affine, e.g., variants of Reed-Muller, Hadamard, and the Long
Code.

1.3 Our Result

The main result of this paper is the following Theorem 2. The theorem states that
every affine LTC with minimal distance at least $\frac{9}{10}$, completeness at least $1 - \epsilon$, and
soundness at most $1 - 3\epsilon$, is of size at most $|\mathbb{F}|$.

Theorem 2 (Main). *Let $\epsilon \in \left[0, \frac{1}{3}\right)$ be a real number, $\mathbb{F}$ be a finite field, and $n \in \mathbb{N}$ be a
natural number. Then, every $\left(\frac{1}{10}, \epsilon, 1 - 3\epsilon\right)$-affine locally testable code $C \subseteq \mathbb{F}^n$ satisfies
$|C| \leq |\mathbb{F}|$.*

Remark. The upper bound on the size of the code suggested by the theorem is tight:
Let $\mathbb{F}$ be a finite field, and $n \in \mathbb{N}$ be a natural number. Consider the code $C = \{a^n\}_{a \in \mathbb{F}} \subseteq \mathbb{F}^n$, and the following affine local tester for C: Randomly select $(i, j) \in [n]^2$,
and test $v_i = v_j$. The code has minimal distance 1, completeness 1, and constant
soundness $s < 1$, and $|C| = |F|$.

4

1.4 Previous Works

Unique Locally Testable Codes. In [16] we consider LTCs with properties similar to the ones required from a PCP by the UGC. Specifically, we study LTCs with *arbitrarily small constant error* (soundness close to 0), that only make unique tests. We show that such LTCs must be of constant size where the constant depends on the alphabet size, the completeness and the soundness.

The current work shows that in the affine case, even LTCs with *large constant error* (soundness close to 1), are of constant size. Furthermore, it gives a better upper bound on the size of such LTCs. The bound only depends on the alphabet size. We mention that the current work uses a different (and simpler) technique.

Linear 2-Query Locally Testable Codes. In [5], *linear* 2-query LTCs with constant minimal distance $1 - \alpha$ are studied, where linear means that the code is a linear subspace of the vector space $\mathbb{F}^n$. It is shown that such codes must be of size at most $|\mathbb{F}|^{\frac{3}{1-\alpha}}$.

We mention that, w.l.o.g., every tester of a linear LTC is affine [6]. However, there are non-linear LTCs that have affine testers, e.g., the Long Code. Thus, our notion of an affine tester is more general. Furthermore, we achieve a better upper bound on the size of the code (for a certain range of the parameters).

General 2-Query Locally Testable Codes. The same bound of $|\mathbb{F}|^{\frac{3}{1-\alpha}}$ is obtained in [5] also for general (non-linear) 2-query LTCs, when *perfect completeness* and *binary* alphabet are assumed. In [11], *binary* 2-query LTCs with almost-perfect completeness (say $1-\epsilon$, for ϵ approaching 0) are investigated. It is shown that such codes must either have a constant size, or have soundness $s = 1 - O(\epsilon)$.

These results are incomparable to ours. On the one hand, we consider the more general case of LTCs over *arbitrarily large alphabet sets*, and allow *imperfect completeness*. On the other hand, we only consider LTCs with affine tests.

2 Definitions

An affine constraint graph (ACG) is a quadruple $G = (H = (V, E), \mathbb{F}, A, S)$, where H is an undirected multigraph (we allow parallel edges), $\mathbb{F}$ is a finite field, $A : E \to \mathcal{A}_{\mathbb{F}}$ is a function mapping edges to affine functions, and $S \subseteq [n] \times \mathbb{F}$ is a multiset (we allow

duplicate values) satisfying $|\{a \in \mathbb{F} \mid (i, a) \in S\}| \leq 1$ for every $i \in [n]$. We will refer to S as the unitary constraint multiset. We denote $V(G) = V$, $E(G) = E$, $A(G) = A$, and $S(G) = S$. The constraints set of G is $cons(G) = E(G) \cup S(G)$.

For a finite field $\mathbb{F}$, and a finite set V, we denote by $\mathbb{F}^V$ the set of all $|V|$-element vectors over $\mathbb{F}$, indexed by the elements of V. Let $G = (H = (V, E), \mathbb{F}, A, S)$ be an ACG, and let $v \in \mathbb{F}^V$ be a word. Let $e = (i, j) \in E$ be an edge of G, and let f_{ij} be the affine function that e is mapped to by A. We say that v satisfies e if $v_j = f_{ij}(v_i)$. Let $s = (i, a) \in S$. We say that v satisfies s if $v_i = a$. We say that v satisfies G if v satisfies every constraint in $cons(G)$.

Let T be an affine local tester for a code $C \subseteq \mathbb{F}^n$, and let $v \in \mathbb{F}^n$ be a given word. When testing v, the tester T queries at most two coordinates of v and runs an affine test on the answers. Assume that each test is carried out with an equal probability. Then, T gives rise to the following natural ACG, called G_T. The vertex set of G_T is $[n]$, and its field is $\mathbb{F}$. The pair $(i, j) \in [n]^2$ is in $E(G_T)$ if T may run an affine test on v_i and v_j. In this case, the function $A(G_T)$ maps the edge (i, j) to the affine function f_{ij} promised by the affineness property of T. The pair $(i, a) \in [n] \times \mathbb{F}$ is in $S(G_T)$ if T may run the test $v_i = a$.

If T performs different tests with different probabilities, then the number of edges between $(i, j) \in [n]^2$ in the multiset $E(G_T)$, and the number of occurrences of the pair $(i, a) \in [n] \times \mathbb{F}$ in the multiset $S(G_T)$, are proportional to the probability of the corresponding tests being carried out.

Note that G_T fully characterizes the behavior of T, as T can be thought of as operating as follows: Randomly select a constraint from $cons(G_T)$. Query the necessary coordinates of v, and accept if v satisfies the chosen constraint.

3 Proof of Main Theorem

In this section we prove our main result, Theorem 2. We assume to be given a $\left(\frac{1}{10}, \epsilon, 1 - 3\epsilon\right)$-affine LTC $C \subseteq \mathbb{F}^n$, where $\epsilon \in \left[0, \frac{1}{3}\right)$ is a real number, $\mathbb{F}$ is a finite field, and $n \in \mathbb{N}$ is a natural number. We denote by T an $(\epsilon, 1 - 3\epsilon)$-affine local tester T for C. For simplicity of notations, we denote $p = |\mathbb{F}|$. Our goal is to show that $|C| \leq p$.

In order to prove Theorem 2, we show the following two lemmas. The lemmas use following definitions. Let $G = (H = (V, E), \mathbb{F}, A, S)$ and $G' = (H' = (V', E'), \mathbb{F}, A', S')$ be a pair of ACGs. We say that G' is a subgraph of G if $V' \subseteq V$, $E' \subseteq E \cap (V')^2$, A'

is the restriction of A to E', and $S' \subseteq S \cap (V' \times \mathbb{F})$. If, in addition, $V = V'$ then G' is called an edges subgraph of G.

Let $G = (H = (V, E), \mathbb{F}, A, S)$ be an ACG, and let $U \subseteq V$ be a connected component of G. We denote by G_U the subgraph of G induced by U. That is, the vertex set of G_U is U, the edge set is $E \cap U^2$, the mapping of edges to affine functions is the restriction of A to $E \cap U^2$, and the unitary constraint set is $S \cap (U \times \mathbb{F})$. We call U unfixed if there are at least two different words in $\mathbb{F}^U$ that satisfy G_U. Otherwise, U is called fixed.

The following lemma shows that any unfixed connected component of a subgraph of G_T is satisfied by p different words (rather than only two).

Lemma 3. *Let G be a subgraph of G_T, and let U be an unfixed connected component of G. Then, there are exactly p words in $\mathbb{F}^U$ satisfying G_U. Furthermore, any two such words differ on every coordinate.*

Proof. The graph G_U induces a system of linear equations over $\mathbb{F}$: Every pair $(i, a) \in S(G_U)$ induces the equation $v_i = a$. Every edge $(i, j) \in E(G_U)$ induces the equation $v_j = av_i + b$, where $f_{ij}(x) = ax + b$ is the affine function that (i, j) is mapped to by $A(G_U)$.

Since U is unfixed, the system is guaranteed to have at least two solutions. Therefore, the system has at least one free variable. This variable can be set to any of the p field elements, and still allow a solution.

Since U is a connected component of G, an assignment to one of the system's variables, determines the assignment to the rest of the variables. Thus, there are exactly p solutions to the system, and any two such solutions assign a different value to each of the variables. $\qquad \square$

A connected component of an ACG G is called large if it contains more than half of G's vertices. The following lemma shows that any edges subgraph of G_T, that contains a sufficiently large fraction of G_T's constraints and is satisfied by at least two codewords, has a large connected component.

Lemma 4. *Let G be an edges subgraph of G_T satisfying $|cons(G)| \geq (1 - 3\epsilon)|cons(G_T)|$. Assume that G is satisfied by at least two different codewords. Then, G has exactly one large unfixed connected component.*

Proof. We denote by $fix(G)$ the set of all vertices that are contained in fixed connected components of G. We first claim that $|fix(G)| < \frac{4}{5}n$. Assume for contradiction that

$|fix(G)| \geq \frac{4}{5}n$. Let u and u' be two different codewords satisfying G. Since u and u' must agree on $fix(G)$, it holds that $\Delta(u, u') \leq \frac{1}{5}$. But this is impossible, as the minimal distance of the code is $\frac{9}{10}$, and u and u' are different codewords.

Let $U_1, ..., U_k$ be all the unfixed connected components of G. Observe that $\left| \bigcup_{t \in [k]} U_t \right| = |[n] \setminus fix(G)| > \frac{1}{5}n$. Assume for contradiction that none of the components $U_1, ..., U_k$ is large, i.e., $|U_t| \leq \frac{1}{2}n$ for every $t \in [k]$. In this case, there exists a subset $K \subseteq [k]$ such that $W = \bigcup_{t \in K} U_t$ satisfies $\frac{1}{5}n < |W| \leq \left(\frac{1}{5} + \frac{1}{2}\right) n = \frac{7}{10}n$.

For every $t \in K$, the connected component U_t is unfixed. Therefore, using Lemma 3, there exists a word $x^t \in \mathbb{F}^{U_t}$ that satisfies G_{U_t} and differs from u on every coordinate of U_t. We construct a new word $v \in \mathbb{F}^n$ using u as follows: For every $i \in [n] \setminus W$, let $v_i = u_i$. For every $t \in K$ and $i \in U_t$, let $v_i = x_i^t$. Note that $\Delta(u, v) = \frac{|W|}{n}$.

Since v satisfies G, it is accepted by T with probability at least $1 - 3\epsilon$. Due to T's soundness, there exists a codewords w such that $\Delta(v, w) < \frac{1}{5}$. Therefore,

$$\Delta(u, w) \geq \Delta(u, v) - \Delta(v, w) > \frac{1}{5} - \frac{1}{5} = 0,$$

implying $u \neq w$. Since the relative distance of the code is at least $\frac{9}{10}$, and u and w are different codewords, it must holds that $\Delta(u, w) \geq \frac{9}{10}$. But,

$$\Delta(u, w) \leq \Delta(u, v) + \Delta(v, w) < \frac{7}{10} + \frac{1}{5} = \frac{9}{10},$$

a contradiction. We conclude that one of the unfixed connected components $U_1, ..., U_k$ is large. $\qquad\square$

Let G be an ACG, and let v be a word. We denote by $G(v)$ the edges subgraph of G containing all edges and unitary constraints of G satisfied by v. Let G_1 and G_2 be a pair of subgraphs of G. We denote by $G_1 \cap G_2$ the subgraph of G whose vertex set is $V(G_1) \cap V(G_2)$, edge set is $E(G_1) \cap E(G_2)$, mapping of edges to affine functions is the restriction of $A(G)$ to $E(G_1) \cap E(G_2)$, and unitary constraint set is $S(G_1) \cap S(G_2)$.

We are now ready to prove our main theorem.

Proof of Theorem 2. Let u and u' be a pair of different codewords, and let $G = G(u) \cap G(u')$. Due to T's completeness, it holds that $|cons(G)| \geq (1 - 2\epsilon)|cons(G_T)|$. Let U be the large unfixed connected component of G promised by Lemma 4. Using Lemma 3, there exist p words $x^1, ..., x^p \in \mathbb{F}^U$ that differ on every coordinate and satisfy G_U. We construct p different words $v^1, ..., v^p \in \mathbb{F}^n$. For $t \in [p]$, the word v^t is

8

constructed as follows: For every $i \in [n] \setminus U$, let $v_i^t = u_i$. For every $i \in U$, let $v_i^t = x_i^t$.

For every $t \in [p]$, the word v^t is accepted by T with probability at least $1 - 2\epsilon$, as v^t satisfies G. Let u^t be the codeword satisfying $\Delta(v^t, u^t) < \frac{1}{5}$ promised by T's soundness. Let w be a codeword. We next show that there exits $t \in [p]$ such that $w = u^t$ and conclude $|C| \le p$.

Consider $G' = G \cap G(w)$. Due to T's completeness, it holds that $|cons(G')| \ge (1 - 3\epsilon)|cons(G_T)|$. Let U' be the large unfixed connected component of G' promised by Lemma 4. Note that $U \cap U' \neq \phi$, as both $U, U' \subseteq [n]$ and each contains more than $\frac{1}{2}n$ vertices. Let $i \in U \cap U'$. For $t \neq r \in [p]$, it holds that $v_i^t \neq v_i^r$, as $x_i^t \neq x_i^r$. Therefore, the set $\{v_i^1, ..., v_i^p\}$ contains p different values. Specifically, there exists $t \in [p]$ such that $w_i = v_i^t$. Recall that U' is a connected component of G', and that w and v^t both satisfy G'. Therefore, w and v^t must agree on every coordinate of U', implying $\Delta(w, v^t) < \frac{1}{2}$. It holds that

$$\Delta(w, u^t) \le \Delta(w, v^t) + \Delta(v^t, u^t) < \frac{1}{2} + \frac{1}{5} < \frac{9}{10}.$$

Since the relative distance of the code is at least $\frac{9}{10}$, and both w and u^t are codewords, it must be the case that $w = u^t$. □

Remark. We note that the affineness property of T was only used in the proof of Lemma 3. Therefore, a more general version of Theorem 2, requiring the tester to satisfy the claim of Lemma 3 instead of the affineness property, holds.

References

[1] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[2] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[3] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.

[4] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.

[5] Eli Ben-Sasson, Oded Goldreich, and Madhu Sudan. Bounds on 2-query codeword testing. *RANDOM-APPROX*, pages 216–227, 2003.

[6] Eli Ben-Sasson, Prahladh Harsha, and Sofya Raskhodnikova. Some 3CNF properties are hard to test. *SIAM Journal on Computing*, 35(1):1–21, 2005.

[7] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Computational Complexity*, 15(2):94–114, 2006.

[8] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *STOC*, pages 344–353, 2006.

[9] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268–292, 1996.

[10] Oded Goldreich and Madhu Sudan. Locally testable codes and PCPs of almost-linear length. *Journal of the ACM*, 53(4):558–655, 2006.

[11] Venkatesan Guruswami. On 2-query codeword testing with near-perfect completeness. *ISAAC*, pages 267–276, 2006.

[12] Subhash Khot. On the power of unique 2-prover 1-round games. *STOC*, pages 767–775, 2002.

[13] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *Journal of the ACM*, 37(1):319–357, 2007.

[14] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-ϵ. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.

[15] Subhash Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. *FOCS*, pages 53–62, 2005.

[16] Gillat Kol and Ran Raz. Locally testable codes analogues to the unique games conjecture do not exist. *ECCC Report TR09-128*, 2009.

[17] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *FOCS*, pages 21–30, 2005.

[18] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? *STOC*, pages 245–254, 2008.