

Lower Bounds for Constant Query Affine-Invariant LCCs and LTCs

ARNAB BHATTACHARYYA, Indian Institute of Science
SIVAKANTH GOPI, Princeton University

Affine-invariant codes are codes whose coordinates form a vector space over a finite field and which are invariant under affine transformations of the coordinate space. They form a natural, well-studied class of codes; they include popular codes such as Reed-Muller and Reed-Solomon. A particularly appealing feature of affine-invariant codes is that they seem well suited to admit local correctors and testers. In this work, we give lower bounds on the length of locally correctable and locally testable affine-invariant codes with constant query complexity. We show that if a code $C \subset \Sigma^{\mathbb{K}^n}$ is an r -query affine invariant locally correctable code (LCC), where \mathbb{K} is a finite field and Σ is a finite alphabet, then the number of codewords in C is at most $\exp(O_{\mathbb{K},r,|\Sigma|}(nr^{-1}))$. Also, we show that if $C \subset \Sigma^{\mathbb{K}^n}$ is an r -query affine invariant locally testable code (LTC), then the number of codewords in C is at most $\exp(O_{\mathbb{K},r,|\Sigma|}(nr^{-2}))$. The dependence on n in these bounds is tight for constant-query LCCs/LTCs, since Guo, Kopparty, and Sudan (ITCS'13) constructed affine-invariant codes via lifting that have the same asymptotic tradeoffs. Note that our result holds for non-linear codes, whereas previously, Ben-Sasson and Sudan (RANDOM'11) assumed linearity to derive similar results. Our analysis uses higher-order Fourier analysis. In particular, we show that the codewords corresponding to an affine-invariant LCC/LTC must be far from each other with respect to Gowers norm of an appropriate order. This then allows us to bound the number of codewords, using known decomposition theorems, which approximate any bounded function in terms of a finite number of low-degree non-classical polynomials, up to a small error in the Gowers norm.

CCS Concepts: • **Theory of computation** → **Error-correcting codes**; • **Mathematics of computing** → *Coding theory*;

Additional Key Words and Phrases: Locally correctable code, locally testable code, affine invariance, Gowers uniformity norm

ACM Reference Format:

Arnab Bhattacharyya and Sivakanth Gopi. 2017. Lower bounds for constant query affine-invariant LCCs and LTCs. *ACM Trans. Comput. Theory* 9, 2, Article 7 (April 2017), 17 pages.
DOI: <http://dx.doi.org/10.1145/3016802>

1. INTRODUCTION

Error-correcting codes which admit local algorithms are of significant interest in theoretical computer science. A code is called a locally correctable code (LCC) if there is a randomized algorithm that, given an index i and a received word w close to a codeword c in Hamming distance, outputs c_i by querying only a few positions of w . A code is called a locally testable code (LTC) if there is a randomized algorithm that, given a received word w , determines whether w is in the code or whether w is far in Hamming distance

The first author was supported in part by DSTO1358 Ramanujan Fellowship. The second author was partially supported by NSF grants CCF-1523816 and CCF-1217416.

Authors' addresses: A. Bhattacharyya, Department of Computer Science and Automation, Indian Institute of Science, Bangalore-560012, India; S. Gopi, Computer Science Department, Princeton University, 35 Olden St, Princeton, NJ-08540, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2017 ACM 1942-3454/2017/04-ART7 \$15.00

DOI: <http://dx.doi.org/10.1145/3016802>

from every codeword, based on queries to a small number of locations of w . The number of positions of the received word queried is called the query complexity of the LCC or LTC.

The notions of local correctability and local testability have a long history in computer science by now. Also called “self-correction,” the idea of local correction originated in works by Lipton [1990] and by Blum and Kannan [1995] on program checkers. LCCs are closely related to locally decodable codes (LDCs), where the goal is to recover a symbol of the underlying message when given a corrupted codeword, using a small number of queries [Katz and Trevisan 2000]. LDCs and LCCs have found applications in private information retrieval schemes [Chor et al. 1998; Barkol et al. 2007] and derandomization [Sudan et al. 1999]. See Yekhanin [2011] for a detailed survey on LDCs and LCCs. Research on LTCs implicitly started with Blum, Luby, and Rubinfeld’s seminal discovery [Blum et al. 1993] that the Hadamard code is an LTC with query complexity 3; they were first formally defined by Goldreich and Sudan in Goldreich and Sudan [2006]. LTCs have been used (implicitly and explicitly) in many contexts, most notably in the construction of probabilistically checkable proofs (PCPs) [Arora and Safra 1998; Arora et al. 1998; Dinur 2007].

In spite of the wide interest in them, some basic questions about LCCs and LTCs remain unanswered. We restrict ourselves throughout to the setting where the query complexity is a constant (independent of the length of the code) and consider the tradeoff between query complexity and code length. The current best constant-query LCCs have exponential length, while the current best constant-query LTCs have near-linear length but they are quite complicated [Ben-Sasson and Sudan 2008; Dinur 2007; Meir 2009; Viderman 2015]. Getting sub-exponential length LCCs or linear length LTCs with constant query complexity are major open problems in the area.

Intuitively, for LCCs and LTCs with constant query complexity, there must be a lot of redundancy in the code, since every symbol of the codeword must satisfy local constraints with most other symbols in the codeword. A systematic way to generate redundancy is to make sure that the code has a large group of *invariances*.¹ Formally, given a code $\mathcal{C} \subset \Sigma^N$ of length N over alphabet Σ , a codeword $c \in \mathcal{C}$ can be naturally viewed as a function $c : [N] \rightarrow \Sigma$. Then, we say that \mathcal{C} is invariant under a set² $G \subset \{[N] \rightarrow [N]\}$ if for every $\pi \in G$ and codeword $c \in \mathcal{C}$, $c \circ \pi$ also describes a codeword $c' \in \mathcal{C}$. Now, the key observation is that if for every codeword $c \in \mathcal{C}$, if there is a constraint among $c(i_1), \dots, c(i_k)$ for some $i_1, \dots, i_k \in [N]$, then for every $c \in \mathcal{C}$, there must also be a constraint among $c(\pi(i_1)), \dots, c(\pi(i_k))$ for any π in the invariance set G , since $c \circ \pi$ is itself another codeword. Hence, if G is large, the presence of one local constraint immediately implies presence of many and suggests the possibility of local algorithms for the code. This connection between invariance and correctability/testability was first explicitly examined by Kaufman and Sudan [2008]. One is then motivated to understand more clearly the possibilities and limitations of local correctors/testers for codes possessing natural symmetries.

We focus on affine-invariant codes, for which the domain $[N]$ is an n -dimensional vector space \mathbb{K}^n over a finite field \mathbb{K} and the code $\mathcal{C} \subset \{\mathbb{K}^n \rightarrow \Sigma\}$ is invariant under affine transformations $A : \mathbb{K}^n \rightarrow \mathbb{K}^n$. Affine invariance is a very natural symmetry for “algebraic codes” and has long been studied in coding theory [Kasami et al. 1967]. The study of affine-invariant LCCs and LTCs was initiated in Kaufman and Sudan [2008] and has been investigated in several follow-up works [Ben-Sasson and Sudan 2011; Guo

¹A quite different way to generate redundancy is through *tensoring*; see Ben-Sasson and Sudan [2004]. Invariances and tensoring are essentially the only two “generic” reasons known to cause local correctability/testability.

² $\{A \rightarrow B\}$ and B^A denote the set of all functions from A to B .

2013; Ben-Sasson et al. 2012; Guruswami et al. 2015]. The hope is that because affine-invariant codes have a large group of invariance and, at the same time, are conducive to non-trivial algebraic constructions, they may contain a code that improves current constructions of LCCs or LTCs.

The current best parameters for constant-query affine-invariant LCCs and LTCs are achieved by the lifted codes of Guo, Kopparty, and Sudan [Guo et al. 2013]. They construct an affine-invariant code $\mathcal{F} \subset \{\mathbb{F}_2^n \rightarrow \mathbb{F}_2\}$ with $\exp(\Theta(n^{r-2}))$ codewords that is an $(r-1)$ -query LCC and an r -query LTC, where $r = 2^\ell$. The $\Theta(\cdot)$ notation hides factors that depend on r but not n . For LCCs, the same asymptotic tradeoff between query complexity and code length is achieved by the Reed-Muller code. For every $r \geq 2$, the Reed-Muller code of order $r-1$ (i.e., polynomials over \mathbb{F}_q on n variables of total degree $\leq r-1$ with $q > r$) is an affine-invariant r -query LCC with $\exp(\Theta(n^{r-1}))$ codewords. In fact, even if we drop the affine-invariance requirement, Reed-Muller codes and the construction of Guo et al. [2013] achieve the best known codeword length for constant query LCCs.³

In this work, we show that the parameters for the lifted codes of Guo et al. [2013] are, in fact, tight for affine-invariant LCCs/LTCs in $\{\mathbb{K}^n \rightarrow \Sigma\}$ for any fixed finite field \mathbb{K} and any fixed finite alphabet Σ .

THEOREM 1.1 (MAIN RESULT, INFORMAL).

- (i) Let $\mathcal{C} \subset \{\mathbb{K}^n \rightarrow \Sigma\}$ be an r -query affine-invariant LCC. Then $|\mathcal{C}| \leq \exp(O_{\mathbb{K},r,|\Sigma|}(n^{r-1}))$.
- (ii) Let $\mathcal{C} \subset \{\mathbb{K}^n \rightarrow \Sigma\}$ be an r -query affine-invariant LTC. Then $|\mathcal{C}| \leq \exp(O_{\mathbb{K},r,|\Sigma|}(n^{r-2}))$.

Note that a local constraint among t coordinates can be used to correct one of the coordinates using $t-1$ queries by a local corrector whereas a local tester needs to make t queries to check the constraint. This explains the difference in the dependence of r in the bounds for LCCs and LTCs.

1.1. Related Work

In the work of Ben-Sasson and Sudan [2011], a similar result as Theorem 1.1 was obtained when the code is assumed to be linear, i.e., when the codewords form a vector space. They showed that if $\mathcal{C} \subset \{\mathbb{K}^n \rightarrow \mathbb{F}\}$ is an $(r-1)$ -query locally correctable or r -query locally testable *linear*, affine-invariant code, where \mathbb{K} and \mathbb{F} are finite fields of characteristic $p > 0$ with \mathbb{K} as an extension of \mathbb{F} , then the dimension of \mathcal{C} as a vector space over \mathbb{F} is at most $(n \log_p |\mathbb{K}|)^{r-2}$. When \mathbb{K} is fixed (as in Guo et al.'s [2013] construction of constant query LCCs/LTCs), the result of Ben-Sasson and Sudan [2011] is a very special case of our Theorem 1.1. On the other hand, the result of the work done by Ben-Sasson and Sudan [2011] also applies when the size of \mathbb{K} is growing (as long as \mathbb{K} extends \mathbb{F}), whereas ours does not.

There are several works which study lower bounds for constant query LCCs [Katz and Trevisan 2000; Goldreich et al. 2012; Dvir and Shpilka 2007; Kerenidis and de Wolf 2003; Barak et al. 2011; Bhattacharyya et al. 2011; Woodruff 2012; Dvir et al. 2014]. For general (non-affine-invariant) LCCs, tight lower bounds are known only for 2-query LCCs. Kerenidis and deWolf [2003] prove that if $\mathcal{C} \subset \{\{0, 1\}^n \rightarrow \Sigma\}$ is a 2-query LCC,⁴ then $|\mathcal{C}| \leq \exp(O(n|\Sigma|^5))$. This is tight for constant Σ and achieved by the Hadamard code. For r -query LCCs where $r > 2$, the lower bounds known are much weaker. The

³In contrast, there exist non-affine-invariant LTCs of constant query complexity and inverse polylogarithmic rate. This corresponds to a constant query LTC $\mathcal{C} \subset \{\{0, 1\}^n \rightarrow \{0, 1\}\}$ with $\exp(2^n/\text{poly}(n))$ codewords, while the affine-invariant LTC of Guo et al. [2013] and Reed-Muller codes have $\exp(\text{poly}(n))$ codewords for constant query complexity.

⁴Their lower bound also holds for the weaker notion of LDC.

best known bounds, due to Kerenidis and de Wolf [2003] and Woodruff [2007], show that if $\mathcal{C} \subset \{\{0, 1\}^n \rightarrow \{0, 1\}\}$ is an r -query LCC, then

$$|\mathcal{C}| \leq \exp(2^{n/(1+1/(\lceil r/2 \rceil + 1)) + o(n)}).$$

Higher-order Fourier analysis was applied to other problems in coding theory in Bhowmick and Lovett [2015b] and Tulsiani and Wolf [2014].

1.2. Proof Overview

Our arguments are based on standard techniques from higher-order Fourier analysis [Tao 2012], but they are new in this context. We show that if an affine-invariant code is an r -query LCC, then its codewords are far from each other in the U^r -norm, the *Gowers norm of order r* . Similarly, we show that the codewords of an affine-invariant r -query LTC are far from each other in the U^{r-1} -norm. Therefore, we can upper bound the number of LCC/LTC codewords in terms of the size of a net that is fine enough with respect to the Gowers norm of an appropriate order. We bound the size of such a net by explicitly constructing one using a standard decomposition theorem (analogous to Szemerédi's regularity lemma): any bounded function $f : \mathbb{K}^n \rightarrow \mathbb{C}$ can be approximated, up to a small error in the Gowers norm, by a composition of a bounded number of low-degree non-classical polynomials [Tao and Ziegler 2012].

The way we argue that two codewords f and g of an r -query LCC are far in the Gowers norm is that if $\|f - g\|_{U^r} < \epsilon$, then for small enough ϵ (with respect to r , $|\Sigma|$, and correctness probability), the local corrector when applied to f can act as if it is applied to g . The argument is, briefly, as follows. On the one hand, the codewords f and g must be far in Hamming distance, because the definition of LCC implies that there is a unique codeword close to any string. So, with constant probability over choice of $y \in \mathbb{K}^n$, the local corrector's guess for $f(y)$ must differ from $g(y)$. On the other hand, we can lower bound, by a constant, the probability of the event that the corrector outputs $g(y)$ when it queries coordinates of f , because f and g are close in the $\|\cdot\|_{U^r}$ norm. This last calculation uses the affine invariance of the code and the *generalized von Neumann inequality*, which bounds by $\|f_0\|_{U^k}$ the expectation over $z_1, \dots, z_m \in \mathbb{K}^n$ of the product $\prod_{i=0}^k f_i(\mathcal{L}_i(z_1, \dots, z_m))$, where the \mathcal{L}_i 's are arbitrary linear forms so that no two are linearly dependent and $f_i : \mathbb{K}^n \rightarrow \mathbb{C}$ are arbitrary functions with $|f_i| \leq 1$.

The argument for r -query LTCs is similar. Suppose f and g are close in the $\|\cdot\|_{U^{r-1}}$ norm. Consider the random function H such that for every x independently, $H(x)$ equals $f(x)$ with probability $1/2$ and $g(x)$ with probability $1/2$. H itself is far from a codeword with high probability. But we show that since the local tester accepts f , it will also accept $H \circ \ell$ for a random invertible affine map $\ell : \mathbb{K}^n \rightarrow \mathbb{K}^n$ with good probability. This implies that with good probability, $H \circ \ell$ is close to a codeword and by affine-invariance, H itself is close to a codeword which gives a contradiction. To draw this conclusion, we again use the generalized von Neumann inequality as well as a hybrid argument.

Organization. Section 2 contains preliminaries that lay the foundations of our analysis. Section 3 proves the first part of our main result about LCCs, while Section 4 proves the second part about LTCs.

2. PRELIMINARIES

2.1. Error-correcting Codes

Let \mathcal{X} be a finite set called the set of coordinates and Σ be another finite set called the alphabet. Let $\Sigma^{\mathcal{X}}$ denote the set of all functions from $\mathcal{X} \rightarrow \Sigma$. A subset $\mathcal{C} \subset \Sigma^{\mathcal{X}}$ is called a code and its elements are called *codewords*.

Definition 2.1 (Hamming Distance). Given $f, g \in \Sigma^\mathcal{X}$, we define the (normalized) Hamming distance between f and g as $\Delta(f, g) := \Pr_{x \in \mathcal{X}}[f(x) \neq g(x)]$ where x is uniformly chosen from \mathcal{X} . For a code $\mathcal{C} \subset \Sigma^\mathcal{X}$, we define the minimum distance of \mathcal{C} as $\min_{f, g \in \mathcal{C}, f \neq g} \Delta(f, g)$.

Let $\blacktriangle_\Sigma = \{q : \Sigma \rightarrow \mathbb{R}_{\geq 0} : \sum_{i \in \Sigma} q(i) = 1\}$ denote the probability simplex on Σ . We embed Σ into \blacktriangle_Σ by sending $i \in \Sigma$ to e_i , which is the i^{th} coordinate vector in \mathbb{R}^Σ . This also lets us extend functions $f : \mathcal{X} \rightarrow \Sigma$ to $\hat{f} : \mathcal{X} \rightarrow \blacktriangle_\Sigma$ using the embedding. We call \hat{f} the simplex extension of f . Now given $f, g \in \Sigma^\mathcal{X}$, we can write the Hamming distance between them as

$$\Delta(f, g) = 1 - \Pr_{x \in \mathcal{X}}[f(x) = g(x)] = 1 - \mathbb{E}_{x \in \mathcal{X}} \langle \hat{f}, \hat{g} \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product in \mathbb{R}^Σ .

Definition 2.2 (Affine Invariance). Let \mathcal{X} be a finite dimensional vector space over some finite field \mathbb{K} , then $\mathcal{C} \subset \Sigma^\mathcal{X}$ is called affine invariant if for every $f \in \mathcal{C}$ and every invertible affine map $\ell : \mathcal{X} \rightarrow \mathcal{X}$, $f \circ \ell \in \mathcal{C}$.

Locally correctable and testable codes are defined formally in Sections 3 and 4, respectively.

2.2. Higher Order Fourier Analysis

Fix a finite field \mathbb{F}_p of prime order p , and let $\mathbb{K} = \mathbb{F}_q$ where $q = p^t$ for a positive integer t . \mathbb{K} is then a vector space of dimension t over \mathbb{F}_p . We denote by $\text{Tr} : \mathbb{K} \rightarrow \mathbb{F}_p$ the *trace function*:

$$\text{Tr}(x) = x + x^p + x^{p^2} + \dots + x^{p^{t-1}}.$$

Also, we use $|\cdot|$ to denote the obvious map from \mathbb{F}_p to $\{0, 1, \dots, p-1\}$.

Given functions $f, g : \mathbb{K}^n \rightarrow \mathbb{C}$, we define their inner product as $\langle f, g \rangle = \mathbb{E}_x[\overline{f(x)}g(x)]$ where x is chosen uniformly from \mathbb{K}^n . We define $\|\cdot\|_p$ -norm on such functions as $\|f\|_p = \mathbb{E}_x[|f(x)|^p]^{1/p}$. We say a function $f : \mathbb{K}^n \rightarrow \mathbb{C}$ is *bounded* if $|f| \leq 1$. Let \mathbb{T} denote the circle group \mathbb{R}/\mathbb{Z} and $e : \mathbb{T} \rightarrow \mathbb{C}$ be the map given by $e(x) = \exp(2\pi ix)$.

Definition 2.3 (Non-Classical Polynomials). A non-classical polynomial of degree $< d$ is a function $f : \mathbb{K}^n \rightarrow \mathbb{T}$ if

$$\forall h_1, h_2, \dots, h_d \in \mathbb{K}^n \quad D_{h_1} D_{h_2} \dots D_{h_d} f = 0,$$

where D_h is the difference operator defined as $D_h f(x) = f(x+h) - f(x)$. For such an f , the function $e(f)$ is called a *non-classical phase polynomial of degree $< d$* .

Note that the derivative operator is linear, and so the multiplicative structure of the field \mathbb{K} is ignored here. Because, as an additive group, \mathbb{K}^n is isomorphic to \mathbb{F}^{tn} , a non-classical polynomial $P : \mathbb{K}^n \rightarrow \mathbb{T}$ over \mathbb{K} can also be identified as a non-classical polynomial $P : \mathbb{F}^{tn} \rightarrow \mathbb{T}$ over \mathbb{F} .

Let $\alpha_1, \dots, \alpha_t \in \mathbb{K}$ be a basis for \mathbb{K} when viewed as a vector space over \mathbb{F}_p . It is known [Tao and Ziegler 2012; Bhattacharyya and Bhowmick 2015] that non-classical polynomials of degree $\leq d$ are exactly those functions $P : \mathbb{K}^n \rightarrow \mathbb{T}$, which have the

following form:

$$P(x_1, \dots, x_n) = \theta + \sum_{k \geq 0} \sum_{\substack{0 \leq d_{i,j} < p \forall i \in [n], j \in [t]; \\ 0 < \sum_{i=1}^n \sum_{j=1}^t d_{i,j} \leq d - k(p-1)}} \frac{c_{d_{1,1}, \dots, d_{n,t}, k} \prod_{i=1}^n \prod_{j=1}^t |\text{Tr}(\alpha_j x_i)|^{d_{i,j}}}{p^{k+1}} \pmod{1} \quad (1)$$

for some $c_{d_{1,1}, \dots, d_{n,t}, k} \in \{0, 1, \dots, p-1\}$ and $\theta \in \mathbb{T}$. Next, we define the Gowers norm for arbitrary functions $f : \mathbb{K}^n \rightarrow \mathbb{C}$.

Definition 2.4 (Gowers uniformity norm [Gowers 2001]). For a function $f : \mathbb{K}^n \rightarrow \mathbb{C}$, the Gowers norm of order r , denoted by $\| \cdot \|_{U^r}$, is defined as

$$\|f\|_{U^r} = \left(\mathbb{E}_{x, h_1, \dots, h_r \in \mathbb{K}^n} [\Delta_{h_1} \Delta_{h_2} \dots \Delta_{h_r} f(x)] \right)^{1/2^r},$$

where Δ_h is the multiplicative difference operator defined as $\Delta_h f(x) = f(x+h) \overline{f(x)}$.

The Gowers norm is an actual norm when $r \geq 2$. It also satisfies a useful monotonicity property: for any function $f : \mathbb{K}^n \rightarrow \mathbb{C}$,

$$|\mathbb{E}[f(x)]| = \|f\|_{U^1} \leq \|f\|_{U^2} \leq \dots \leq \|f\|_{U^r} \leq \dots \leq \|f\|_{U^\infty}.$$

See Tao [2012] for more on Gowers norm. Observe that if $f : \mathbb{K}^n \rightarrow \mathbb{C}$ is a non-classical phase polynomial of degree $< r$, then $\|f\|_{U^r} = 1$. The inverse Gowers theorem is a partial converse to this. It shows that the Gowers norm of order r of a function is in direct correspondence with its correlation with non-classical phase polynomials of degree $< r$. In particular:

LEMMA 2.5 (INVERSE GOWERS THEOREM [TAO AND ZIEGLER 2012]). *For any bounded⁵ $f : \mathbb{K}^n \rightarrow \mathbb{C}$, if $\|f\|_{U^r} > \delta$, then there exists a non-classical polynomial P of degree $< r$ such that*

$$|\langle f, e(P) \rangle| \geq c(\delta, \mathbb{K}, r),$$

where $c(\delta, \mathbb{K}, r)$ is a constant depending only on δ, \mathbb{K}, r .

A linear form on m variables is a vector $\mathcal{L} = (w_1, \dots, w_m) \in \mathbb{K}^m$ that is interpreted as a function $\mathcal{L} : (\mathbb{K}^n)^m \rightarrow \mathbb{K}^n$ via the map $(x_1, \dots, x_m) \mapsto \sum_{i=1}^m w_i x_i$. A key reason that the Gowers norm is useful in applications is that if a function has small Gowers norm of the appropriate order, then it behaves pseudorandomly in a certain way with respect to linear forms.

LEMMA 2.6 (GENERALIZED VON NEUMANN INEQUALITY (EXERCISE 1.3.23 IN TAO [2012])). *Let $f_0, f_1, f_2, \dots, f_k : \mathbb{K}^n \rightarrow \mathbb{C}$ be bounded functions and let $\mathcal{L} = \{\mathcal{L}_0, \mathcal{L}_1, \dots, \mathcal{L}_k\}$ be a system of $k+1$ linear forms in m variables such that no form is a multiple of another. Then*

$$\left| \mathbb{E}_{z_1, \dots, z_m \in \mathbb{K}^n} \left[\prod_{i=0}^k f_i(\mathcal{L}_i(z_1, \dots, z_m)) \right] \right| \leq \min_{0 \leq i \leq k} \|f_i\|_{U^k}.$$

See Appendix A for proof.

2.3. A Net for Gowers Norm

The goal of this section is to establish the following claim.

⁵Note that bounded means $|f| \leq 1$.

THEOREM 2.7 (ϵ -NET FOR U^r NORM). *The metric induced by the $\|\cdot\|_{U^r}$ norm on the space of all bounded functions $\{f : \mathbb{K}^n \rightarrow \mathbb{C}\}$ has an ϵ -net of size $\exp(O_{\epsilon, \mathbb{K}, r}(n^{r-1}))$.*

For the proof, we need the following definitions.

Definition 2.8 (*Polynomial Factors*). A polynomial factor \mathcal{B} is a sequence of non-classical polynomials $P_1, \dots, P_k : \mathbb{K}^n \rightarrow \mathbb{T}$. We also identify it with the function $\mathcal{B} : \mathbb{K}^n \rightarrow \mathbb{T}^k$ mapping $x \mapsto (P_1(x), \dots, P_k(x))$. The partition induced by \mathcal{B} is the partition of \mathbb{K}^n given by $\{\mathcal{B}^{-1}(y) : y \in \mathbb{T}^k\}$. The complexity of \mathcal{B} is the number of defining polynomials, $|\mathcal{B}| = k$. The degree of \mathcal{B} is the maximum degree among its defining polynomials P_1, \dots, P_k . A function $f : \mathbb{K}^n \rightarrow \mathbb{C}$ is called \mathcal{B} -measurable if it is constant in each cell of the partition induced by \mathcal{B} or equivalently f can be written as a $\tau(P_1, \dots, P_k)$ for some function $\tau : \mathbb{T}^k \rightarrow \mathbb{C}$.

Definition 2.9 (*Conditional Expectations*). Given a polynomial factor \mathcal{B} , the conditional expectation of $f : \mathbb{K}^n \rightarrow \mathbb{C}$ over \mathcal{B} , denoted by $\mathbb{E}[f|\mathcal{B}]$, is the \mathcal{B} -measurable function defined by

$$\mathbb{E}[f|\mathcal{B}](x) = \mathbb{E}_{y \in \mathcal{B}^{-1}(\mathcal{B}(x))}[f(y)].$$

Definition 2.10 (*Factor Refinement*). Given two polynomial factors $\mathcal{B}, \mathcal{B}'$, we say \mathcal{B}' is a refinement of \mathcal{B} , denoted by $\mathcal{B}' \preceq \mathcal{B}$, if every cell in the partition induced by \mathcal{B}' is contained in some cell in the partition induced by \mathcal{B} .

The definition of refinement immediately implies:

LEMMA 2.11 (**PYTHAGORAS THEOREM**). *Let $\mathcal{B}, \mathcal{B}'$ be polynomial factors such that $\mathcal{B}' \preceq \mathcal{B}$, then for any function $f : \mathbb{K}^n \rightarrow \mathbb{C}$,*

$$\|\mathbb{E}[f|\mathcal{B}']\|_2^2 = \|\mathbb{E}[f|\mathcal{B}]\|_2^2 + \|\mathbb{E}[f|\mathcal{B}'] - \mathbb{E}[f|\mathcal{B}]\|_2^2.$$

The next claim shows that any bounded function is “close” to being measurable by a polynomial factor of bounded complexity. Precisely:

LEMMA 2.12 (**DECOMPOSITION THEOREM**). *Any bounded $f : \mathbb{K}^n \rightarrow \mathbb{C}$ can be approximated in $\|\cdot\|_{U^r}$ by a function of a small number of degree $< r$ non-classical polynomials, i.e., for any $\epsilon > 0$, there exists non-classical polynomials P_1, P_2, \dots, P_k of degree $< r$ with $P_i(\vec{0}) = 0 \forall i$ and a bounded function $\tau : \mathbb{T}^k \rightarrow \mathbb{C}$ such that*

$$\|f - \tau(P_1, P_2, \dots, P_k)\|_{U^r} \leq \epsilon,$$

where $k = k(\epsilon, \mathbb{K}, r)$ is a constant depending only on ϵ, \mathbb{K}, r .

PROOF. The proof is similar to the proof of the Quadratic Koopman-von Neumann decomposition which is Prop 3.7 in Green [2006] but using the full Inverse Gowers Theorem (Lemma 2.5) and similar claims are implicit elsewhere, but for completeness, we give the proof.

The main idea is to approximate the function f using its conditional expectation over a suitable polynomial factor \mathcal{B} of degree $< r$. We will start with the trivial factor $\mathcal{B}_0 = (1)$ and iteratively construct more refined partitions $\mathcal{B}_i \preceq \mathcal{B}_{i-1}$ until we find a factor \mathcal{B}_k which satisfies $\|f - \mathbb{E}[f|\mathcal{B}_k]\|_{U^r} \leq \epsilon$. To bound the number of iterations needed to achieve this, we will show that the energy $\|\mathbb{E}[f|\mathcal{B}_i]\|_2^2$, which is bounded above by 1, increases by a fixed constant in every step.

Suppose that after step $i - 1$, we still have $\|f - \mathbb{E}[f|\mathcal{B}_{i-1}]\|_{U^r} > \epsilon$. Let $g = f - \mathbb{E}[f|\mathcal{B}_{i-1}]$, then by the inverse Gowers theorem (Lemma 2.5), we have some non-classical polynomial P_i of degree $< r$ such that $|\langle g, e(P_i) \rangle| \geq \kappa = c(\epsilon, p, r)$. We can assume that $P_i(\vec{0}) = 0$. Refine the factor \mathcal{B}_{i-1} by adding the polynomial P_i to obtain $\mathcal{B}_i \preceq \mathcal{B}_{i-1}$. Now

consider the energy increment,

$$\|\mathbb{E}[f|\mathcal{B}_i]\|_2^2 - \|\mathbb{E}[f|\mathcal{B}_{i-1}]\|_2^2 = \|\mathbb{E}[f|\mathcal{B}_i] - \mathbb{E}[f|\mathcal{B}_{i-1}]\|_2^2 = \|\mathbb{E}[g|\mathcal{B}_i]\|_2^2,$$

where we used the Pythagoras theorem (Lemma 2.11) and the fact that $\mathbb{E}[\mathbb{E}[f|\mathcal{B}_{i-1}]|\mathcal{B}_i] = \mathbb{E}[f|\mathcal{B}_{i-1}]$ since $\mathcal{B}_i \leq \mathcal{B}_{i-1}$. So,

$$\begin{aligned} \kappa^2 &\leq |\mathbb{E}[g \cdot e(P_i)]|^2 = |\mathbb{E}[\mathbb{E}[g \cdot e(P_i)|\mathcal{B}_i]]|^2 = |\mathbb{E}[e(P_i)\mathbb{E}[g|\mathcal{B}_i]]|^2 \\ &\leq \|\mathbb{E}[g|\mathcal{B}_i]\|_1^2 \leq \|\mathbb{E}[g|\mathcal{B}_i]\|_2^2 = \|\mathbb{E}[f|\mathcal{B}_i]\|_2^2 - \|\mathbb{E}[f|\mathcal{B}_{i-1}]\|_2^2. \end{aligned}$$

Thus, the energy increases by κ^2 every step. But since the energy is bounded above by 1, the process should end in a finite number of steps $k \leq \frac{1}{\kappa^2}$. So, $\|f - \mathbb{E}[f|\mathcal{B}_k]\|_{U^r} \leq \epsilon$, but since $\mathbb{E}[f|\mathcal{B}_k]$ is \mathcal{B}_k -measurable, we can write $\mathbb{E}[f|\mathcal{B}_k] = \tau(P_1, \dots, P_k)$ for some function τ with $|\tau| = |\mathbb{E}[f|\mathcal{B}_k]| \leq |f| \leq 1$. \square

We are now ready to prove Theorem 2.7.

PROOF OF THEOREM 2.7. Recall that \mathbb{K} is an extension field of dimension t over a prime field \mathbb{F}_p . The ϵ -net will be the set \mathcal{N} of all functions of the form $\tau(P_1, \dots, P_k)$, where P_1, \dots, P_k are degree $< r$ non-classical polynomials with zero constant terms, $\tau : \mathbb{T}^k \rightarrow \mathbb{C}$ is a bounded function and $k = k(\epsilon, p, r)$ is the constant given by Lemma 2.12. But we will not include all possible bounded $\tau : \mathbb{T}^k \rightarrow \mathbb{C}$. Firstly, by Equation (1), P_1, \dots, P_k take values only in $\frac{1}{p^r}\mathbb{Z}/\mathbb{Z}$. Next, we will discretize the set $\{z \in \mathbb{C} : |z| \leq 1\}$ into the ϵ -lattice, i.e., we will only consider maps $\tau : (\frac{1}{p^r}\mathbb{Z}/\mathbb{Z})^k \rightarrow \{z \in \mathbb{C} : |z| \leq 1\} \cap \epsilon(\mathbb{Z} + i\mathbb{Z})$. The number of such maps is bounded by $(4/\epsilon^2)^{p^k}$.

By Equation (1), a non-classical polynomial of degree $< r$ in n variables with zero constant term can be represented by $\leq \binom{nt+r-1}{r-1}r$ coefficients in $\{0, 1, \dots, p-1\}$. So, the number of such non-classical polynomials is bounded by $\exp(O_{r,\mathbb{K}}(n^{r-1}))$. Combining both the bounds,

$$|\mathcal{N}| \leq \exp(O_{r,\mathbb{K}}(n^{r-1}))^k \cdot (4/\epsilon^2)^{p^k} = \exp(O_{\epsilon,\mathbb{K},r}(n^{r-1})).$$

We will now prove that \mathcal{N} is a 3ϵ -net. Given any $f : \mathbb{K}^n \rightarrow [-1, 1]$, using Lemma 2.12, there is a function $\tau(P_1, \dots, P_k)$ such that

$$\|f - \tau(P_1, P_2, \dots, P_k)\|_{U^r} \leq \epsilon.$$

If we consider the $\tilde{\tau} \in \mathcal{N}$ by rounding values real and imaginary parts of τ to the nearest multiple of ϵ , we get

$$\begin{aligned} \|f - \tilde{\tau}(P_1, P_2, \dots, P_k)\|_{U^r} &\leq \|f - \tau(P_1, P_2, \dots, P_k)\|_{U^r} + \|\tau(P_1, P_2, \dots, P_k) - \tilde{\tau}(P_1, P_2, \dots, P_k)\|_{U^r} \\ &\leq \epsilon + \|\tau(P_1, P_2, \dots, P_k) - \tilde{\tau}(P_1, P_2, \dots, P_k)\|_\infty \leq 3\epsilon. \quad \square \end{aligned}$$

3. LOCALLY CORRECTABLE CODES

We begin by defining locally correctable codes formally. Note that the definition below differs from the conventional one in terms of a local correction algorithm and adversarial errors (see, for instance, Yekhanin [2011]); however, our definition is certainly weaker. Therefore, this makes our lower bounds stronger.

Definition 3.1 (LCC). An (r, δ, τ) LCC is a code $\mathcal{C} \subset \Sigma^{\mathcal{X}}$ with the following property: For each $x \in \mathcal{X}$, there is a distribution \mathcal{M}_x over r -tuples of distinct⁶ coordinates such that whenever $\tilde{f} \in \Sigma^{\mathcal{X}}$ is δ -close to some codeword $f \in \mathcal{C}$ in Hamming distance,

$$\Pr_{(y_1, \dots, y_r) \sim \mathcal{M}_x} [\mathcal{D}_{x, y_1, \dots, y_r}(\tilde{f}(y_1), \tilde{f}(y_2), \dots, \tilde{f}(y_r)) = f(x)] \geq 1 - \tau,$$

where $\mathcal{D}_{x, y_1, \dots, y_r} : \Sigma^r \rightarrow \Sigma$, called the decoding operator, depends only on x, y_1, \dots, y_r .

If, furthermore, \mathcal{X} is a vector space and \mathcal{C} is affine invariant, then we call it an affine invariant LCC.

Remark 3.2. Let $|\Sigma| = m$. Without loss of generality, we can assume that $\Sigma = \{1, 2, \dots, m\}$. Then, we can extend functions $f : \mathcal{X} \rightarrow \Sigma$ to $\hat{f} : \mathcal{X} \rightarrow \blacktriangle_m$. The decoding operators $\mathcal{D} : \Sigma^r \rightarrow \Sigma$ can also be extended to $\hat{\mathcal{D}} : \blacktriangle_m^r \rightarrow \blacktriangle_m$ as follows: For $z_1, \dots, z_r \in \blacktriangle_m$, define

$$\hat{\mathcal{D}}(z_1, \dots, z_r) = \sum_{1 \leq i_1, \dots, i_r \leq m} e_{\mathcal{D}(i_1, \dots, i_r)}(z_1)_{i_1} \cdots (z_r)_{i_r},$$

where e_j stands for the j^{th} coordinate vector in \mathbb{R}^m and $(z_j)_i$ is the i^{th} coordinate of the vector z_j . Now, we can rewrite the decoding condition as:

$$\mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}_x} [\langle \hat{f}(x), \hat{\mathcal{D}}_{x, y_1, \dots, y_r}(\hat{f}(y_1), \hat{f}(y_2), \dots, \hat{f}(y_r)) \rangle] \geq 1 - \tau.$$

First, we make the observation that any LCC must have good minimum distance.

LEMMA 3.3. *Let $\mathcal{C} \subset \Sigma^{\mathcal{X}}$ be an (r, δ, τ) LCC with $\tau < 1/2$, then the minimum distance of \mathcal{C} is at least 2δ .*

PROOF. Let $f, g \in \mathcal{C}$ be two distinct codewords such that $\Delta(f, g) < 2\delta$. Let h be the midpoint of f and g , i.e., h is δ -close to both f and g . Let $x \in \mathcal{X}$ be such that $f(x) \neq g(x)$. By the LCC property,

$$\Pr_{(y_1, \dots, y_r) \sim \mathcal{M}_x} [f(x) = \mathcal{D}_{x, y_1, \dots, y_r}(h(y_1), \dots, h(y_r))] \geq 1 - \tau$$

$$\Pr_{(y_1, \dots, y_r) \sim \mathcal{M}_x} [g(x) = \mathcal{D}_{x, y_1, \dots, y_r}(h(y_1), \dots, h(y_r))] \geq 1 - \tau.$$

This is a contradiction when $\tau < \frac{1}{2}$. Therefore, every two codewords must be at least 2δ apart. \square

Now, we are ready to prove our main result of this section.

THEOREM 3.4 (LOWER BOUND FOR LCCs). *Let $\mathcal{C} \subset \Sigma^{\mathbb{K}^n}$ be an (r, δ, τ) affine-invariant LCC where $\tau < \frac{2\delta}{3}$. Then, $|\mathcal{C}| \leq \exp(O_{\delta, \mathbb{K}, r, |\Sigma|}(n^{r-1}))$.*

PROOF. Let $|\Sigma| = m$. Let \mathcal{N} be an $\epsilon/2$ -net for the space of all bounded functions $\{h : \mathbb{K}^n \rightarrow \mathbb{C}\}$ with the metric induced by $\|\cdot\|_{U^r}$ -norm where $\epsilon = \frac{2\delta}{3m^r}$. Given a bounded $h : \mathbb{K}^n \rightarrow \mathbb{C}$, define

$$\phi(h) := \operatorname{argmin}_{h' \in \mathcal{N}} \|h - h'\|_{U^r}$$

(break ties arbitrarily). Since \mathcal{N} is an $\epsilon/2$ net, we have $\|h - \phi(h)\|_{U^r} \leq \epsilon/2$. Define $\Psi : \mathcal{C} \rightarrow \mathcal{N}^m$ as

$$\Psi(f) := (\phi(\hat{f}_1), \dots, \phi(\hat{f}_m)),$$

⁶Without loss of generality, we can assume the tuples have distinct coordinates by adding dummy coordinates and modifying the decoding functions $\mathcal{D}_{x, y_1, \dots, y_r}$.

where $\hat{f}_i : \mathbb{K}^n \rightarrow \mathbb{R}_{\geq 0}$ is the i^{th} coordinate function of the simplex extension $\hat{f} : \mathbb{K}^n \rightarrow \blacktriangle_m$ of $f : \mathbb{K}^n \rightarrow \Sigma$. We claim that Ψ is an injection, which implies that $|\mathcal{C}| \leq |\mathcal{N}|^m$. Now using Theorem 2.7, the required bound follows. Suppose that Ψ is not an injection. Let $f, g \in \mathcal{C}$ be two distinct codewords such that $\Psi(f) = \Psi(g)$. This implies that

$$\forall i \in [m] \quad \|\hat{f}_i - \hat{g}_i\|_{U^r} \leq \|\hat{f}_i - \phi(\hat{f}_i)\|_{U^r} + \|\hat{g}_i - \phi(\hat{g}_i)\|_{U^r} \leq \epsilon.$$

By affine invariance of \mathcal{C} , $f \circ \ell \in \mathcal{C}$ for all invertible affine maps $\ell : \mathbb{K}^n \rightarrow \mathbb{K}^n$. So, by the local correction property,

$$\Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [f \circ \ell(y_0) = \mathcal{D}_{y_0, y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \geq 1 - \tau,$$

where ℓ ranges uniformly over all invertible affine maps from $\mathbb{K}^n \rightarrow \mathbb{K}^n$ and y_0 ranges uniformly over \mathbb{K}^n . Now consider the following difference:

$$\begin{aligned} & \Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [f \circ \ell(y_0) = \mathcal{D}_{y_0, y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \\ & \quad - \Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [g \circ \ell(y_0) = \mathcal{D}_{y_0, y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \\ & = \mathbb{E}_\ell \mathbb{E}_{y_0} \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} \left[\left(\hat{f} \circ \ell(y_0), \widehat{\mathcal{D}}_{y_0, y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \right) \right. \\ & \quad \left. - \left(\hat{g} \circ \ell(y_0), \widehat{\mathcal{D}}_{y_0, y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \right) \right] \\ & = \mathbb{E}_{y_0} \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} \left[\mathbb{E}_\ell \left[\left(\hat{f} \circ \ell(y_0) - \hat{g} \circ \ell(y_0), \widehat{\mathcal{D}}_{y_0, y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \right) \right] \right] \end{aligned}$$

Now, we fix y_0, y_1, \dots, y_r and show that inner expectation is small for each tuple (y_0, y_1, \dots, y_r) . Let us denote $\mathcal{D} = \mathcal{D}_{y_0, y_1, \dots, y_r}$ for brevity. Let $t = \text{rank}(y_0, y_1, \dots, y_r)$,⁷ thus there exist independent vectors $v_1, \dots, v_t \in \mathbb{K}^n$ such that for every $0 \leq i \leq r$, $y_i = \sum_{j=1}^t \lambda_{ij} v_j$ for some fixed $\lambda_{ij} \in \mathbb{K}$. The action of a random invertible affine map ℓ can be approximated by sampling $z_0, z_1, \dots, z_t \in \mathbb{K}^n$ uniformly and mapping $y_i \mapsto z_0 + \sum_{j=1}^t \lambda_{ij} z_j$ since, with probability $1 - o_n(1)$, z_1, \dots, z_t will be independent. Therefore,

$$\begin{aligned} & \mathbb{E}_\ell \left[\left(\hat{f} \circ \ell(y_0) - \hat{g} \circ \ell(y_0), \widehat{\mathcal{D}}_{y_0, y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \right) \right] \\ & = o_n(1) + \mathbb{E}_{z_0, z_1, \dots, z_t \in \mathbb{K}^n} \left[\left\langle \left(\hat{f} - \hat{g} \right) \left(z_0 + \sum_{j=1}^t \lambda_{0j} z_j \right), \right. \right. \\ & \quad \left. \left. \widehat{\mathcal{D}} \left(\hat{f} \left(z_0 + \sum_{j=1}^t \lambda_{1j} z_j \right), \dots, \hat{f} \left(z_0 + \sum_{j=1}^t \lambda_{rj} z_j \right) \right) \right\rangle \right] \\ & \hspace{15em} \text{(we can ignore the } o_n(1) \text{ term)} \end{aligned}$$

⁷ $\text{rank}(y_0, y_1, \dots, y_r)$ is the dimension of the subspace spanned by the vectors y_0, y_1, \dots, y_r .

$$\begin{aligned}
&= \mathbb{E}_{z_0, z_1, \dots, z_t \in \mathbb{K}^n} \left[\left\langle \left(\hat{f} - \hat{g} \right) \left(z_0 + \sum_{j=1}^t \lambda_{0j} z_j \right), \right. \right. \\
&\quad \left. \left. \left(\sum_{1 \leq i_1, \dots, i_r \leq m} e_{\mathcal{D}(i_1, \dots, i_r)} \prod_{k=1}^r \hat{f}_{i_k} \left(z_0 + \sum_{j=1}^t \lambda_{kj} z_j \right) \right) \right\rangle \right] \\
&= \mathbb{E}_{z_0, z_1, \dots, z_t \in \mathbb{K}^n} \left[\left(\sum_{1 \leq i_1, \dots, i_r \leq m} \left(\hat{f} - \hat{g} \right)_{\mathcal{D}(i_1, \dots, i_r)} \left(z_0 + \sum_{j=1}^t \lambda_{0j} z_j \right) \cdot \prod_{k=1}^r \hat{f}_{i_k} \left(z_0 + \sum_{j=1}^t \lambda_{kj} z_j \right) \right) \right] \\
&\leq \left(\sum_{0 \leq i_1, \dots, i_r \leq m-1} \|(\hat{f} - \hat{g})_{\mathcal{D}(i_1, \dots, i_r)}\|_{U^r} \right) \leq m^r \epsilon,
\end{aligned}$$

where the first inequality is obtained by applying generalized von Neumann inequality (Lemma 2.6) to each term. Therefore,

$$\begin{aligned}
&\Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [g \circ \ell(y_0) = \mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \\
&\geq \Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [f \circ \ell(y_0) = \mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] - m^r \epsilon \\
&\geq 1 - \tau - 2\delta/3.
\end{aligned}$$

On the other hand,

$$\begin{aligned}
&\Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [g \circ \ell(y_0) = \mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \\
&\leq \Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [g \circ \ell(y_0) = f \circ \ell(y_0)] \\
&\quad + \Pr_{\ell, y_0, (y_1, \dots, y_r) \sim \mathcal{M}_{y_0}} [f \circ \ell(y_0) \neq \mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] \\
&\leq \Pr_x [f(x) = g(x)] + \tau \leq 1 - 2\delta + \tau \quad (\text{By Lemma 3.3})
\end{aligned}$$

This is a contradiction when $\tau < \frac{2\delta}{3}$. \square

4. LOCALLY TESTABLE CODES

We start by defining locally testable codes in a formulation convenient for our use.

Definition 4.1 (LTC). An (r, δ, τ) LTC is a code $\mathcal{C} \subset \Sigma^{\mathcal{X}}$ with minimum distance at least δ and the following property:

There is a distribution \mathcal{M} over r -tuples of distinct⁸ coordinates such that for each codeword $f \in \mathcal{C}$,

$$\Pr_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(f(y_1), f(y_2), \dots, f(y_r)) = 1] \geq 3/4,$$

⁸Again, without loss of generality, we can assume the tuples have distinct coordinates by adding dummy coordinates and modifying the decoding functions $\mathcal{D}_{y_1, \dots, y_r}$.

and for every $g \in \Sigma^{\mathcal{X}}$, which is τ -far away from every codeword,

$$\Pr_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(g(y_1), g(y_2), \dots, g(y_r)) = 1] \leq 1/4,$$

where $\mathcal{D}_{y_1, \dots, y_r} : \Sigma^r \rightarrow \{0, 1\}$, called the testing operator, depends only on y_1, \dots, y_r .

If, furthermore, \mathcal{X} is a vector space and \mathcal{C} is affine-invariant, then we call it an affine invariant LTC.

Remark 4.2. Let $|\Sigma| = m$. Without loss of generality, we can assume that $\Sigma = \{1, 2, \dots, m\}$. We can extend $f : \mathcal{X} \rightarrow \Sigma$ to $\hat{f} : \mathcal{X} \rightarrow \blacktriangle_m$. The testing operator $\mathcal{D} : \Sigma^r \rightarrow \{0, 1\}$ can also be extended to $\widehat{\mathcal{D}} : \blacktriangle_m^r \rightarrow [0, 1]$ as follows: For $z_1, \dots, z_r \in \blacktriangle_m$, define

$$\widehat{\mathcal{D}}(z_1, \dots, z_r) = \sum_{1 \leq i_1, \dots, i_r \leq m} \mathcal{D}(i_1, \dots, i_r)(z_1)_{i_1} \cdots (z_r)_{i_r}. \quad (2)$$

Now we can rewrite the probability in terms of expectation as:

$$\begin{aligned} & \Pr_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(f(y_1), \dots, f(y_r)) = 1] \\ &= \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\widehat{\mathcal{D}}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))]. \end{aligned}$$

We are now ready to prove the main result of this section.

THEOREM 4.3 (LOWER BOUND FOR LTC'S). *Let $\mathcal{C} \subset \Sigma^{\mathbb{K}^n}$ be an $(r, \delta, \delta/3)$ affine invariant LTC, then $|\mathcal{C}| \leq \exp(O_{\delta, \mathbb{K}, r, |\Sigma|}(n^{r-2}))$.*

PROOF. Let $|\Sigma| = m$. The proof is very similar to that of Theorem 3.4. Let \mathcal{N} be an $\epsilon/2$ -net for the space of all bounded functions $\{f : \mathbb{K}^n \rightarrow \mathbb{C}\}$ with the metric induced by $\|\cdot\|_{U^{r-1}}$ -norm where $\epsilon = 1/2rm$. Define $\Psi : \mathcal{C} \rightarrow \mathcal{N}^m$ as in the proof of Theorem 3.4; it is enough to show that Ψ is an injection. Suppose that Ψ is not an injection, then there exists $f, g \in \mathcal{C}$, which are distinct such that $\Psi(f) = \Psi(g)$. This implies that

$$\forall i \in [m] \quad \|\hat{f}_i - \hat{g}_i\|_{U^{r-1}} \leq \epsilon.$$

By affine invariance of \mathcal{C} , $f \circ \ell \in \mathcal{C}$ for all invertible affine maps $\ell : \mathbb{K}^n \rightarrow \mathbb{K}^n$. So

$$\mathbb{E}_{\ell} \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), f \circ \ell(y_2), \dots, f \circ \ell(y_r))] \geq 3/4,$$

where ℓ ranges over all invertible affine maps from $\mathbb{K}^n \rightarrow \mathbb{K}^n$. Let $H \in \Sigma^{\mathcal{X}}$ be a random word where for each coordinate $x \in \mathcal{X}$ independently,

$$H(x) = \begin{cases} f(x) & \text{with probability } 1/2 \\ g(x) & \text{with probability } 1/2 \end{cases}.$$

Define $\hat{h} : \mathcal{X} \rightarrow \blacktriangle_m$ as $\hat{h}(x) = \mathbb{E}_H[\widehat{H}(x)] = \frac{\hat{f}(x) + \hat{g}(x)}{2}$ where \hat{f}, \hat{g} are the simplex extensions of the original f, g . So $\forall i \in [m] \quad \|\hat{f}_i - \hat{h}_i\|_{U^{r-1}} = \|\hat{f}_i - \hat{g}_i\|_{U^{r-1}}/2 \leq \epsilon/2$. We will now show that the test accepts $H \circ \ell$ with good probability when ℓ is a random invertible

affine map from $\mathbb{K}^n \rightarrow \mathbb{K}^n$.

$$\begin{aligned}
& \mathbb{E}_H \mathbb{E}_\ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r)) \\
& \quad - \mathcal{D}_{y_1, \dots, y_r}(H \circ \ell(y_1), \dots, H \circ \ell(y_r))] \\
& = \mathbb{E}_H \mathbb{E}_\ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \\
& \quad - \widehat{\mathcal{D}}_{y_1, \dots, y_r}(\widehat{H} \circ \ell(y_1), \dots, \widehat{H} \circ \ell(y_r))] \\
& = \mathbb{E}_\ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \\
& \quad - \widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{h} \circ \ell(y_1), \dots, \hat{h} \circ \ell(y_r))] \\
& \text{(using multilinear expansion of } \widehat{\mathcal{D}}_{y_1, \dots, y_r} \text{ (Equation (2)) and taking expectation over } H) \\
& = \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathbb{E}_\ell [\widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) \\
& \quad - \widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{h} \circ \ell(y_1), \dots, \hat{h} \circ \ell(y_r))]
\end{aligned}$$

Now we fix y_1, \dots, y_r and show that inner expectation is small for each tuple (y_1, \dots, y_r) . Let us denote $\mathcal{D} = \mathcal{D}_{y_1, \dots, y_r}$ for brevity. Let $t = \text{rank}(y_1, \dots, y_r)$; thus, there exist independent vectors $v_1, \dots, v_t \in \mathbb{K}^n$ such that for every $1 \leq i \leq r$, $y_i = \sum_{j=1}^t \lambda_{ij} v_j$ for some fixed $\lambda_{ij} \in \mathbb{K}$. The action of a random invertible affine map ℓ can be approximated by sampling $z_0, z_1, \dots, z_t \in \mathbb{K}^n$ uniformly and mapping $y_i \mapsto z_0 + \sum_{j=1}^t \lambda_{ij} z_j$ since with probability $1 - o_n(1)$, z_1, \dots, z_t will be independent. Therefore,

$$\begin{aligned}
& \mathbb{E}_\ell [\widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{f} \circ \ell(y_1), \dots, \hat{f} \circ \ell(y_r)) - \widehat{\mathcal{D}}_{y_1, \dots, y_r}(\hat{h} \circ \ell(y_1), \dots, \hat{h} \circ \ell(y_r))] \\
& = o_n(1) + \mathbb{E}_{z_0, \dots, z_t \in \mathbb{K}^n} \left[\widehat{\mathcal{D}} \left(\hat{f} \left(z_0 + \sum_{j=1}^t \lambda_{1j} z_j \right), \dots, \hat{f} \left(z_0 + \sum_{j=1}^t \lambda_{rj} z_j \right) \right) \right. \\
& \quad \left. - \mathcal{D} \left(\hat{h} \left(z_0 + \sum_{j=1}^t \lambda_{1j} z_j \right), \dots, \hat{h} \left(z_0 + \sum_{j=1}^t \lambda_{rj} z_j \right) \right) \right] \\
& = \mathbb{E}_{z_0, z_1, \dots, z_t \in \mathbb{K}^n} \left[\sum_{1 \leq i_1, \dots, i_r \leq m} \mathcal{D}(i_1, \dots, i_r) \left(\prod_{k=1}^r \hat{f}_{i_k} \left(z_0 + \sum_{j=1}^t \lambda_{kj} z_j \right) - \prod_{k=1}^r \hat{h}_{i_k} \left(z_0 + \sum_{j=1}^t \lambda_{kj} z_j \right) \right) \right] \\
& \leq r \cdot m^r \cdot \frac{\epsilon}{2} = \frac{1}{4},
\end{aligned}$$

where the last line is obtained by forming hybrids, i.e., writing

$$\begin{aligned}
& \hat{f}_{i_1} \cdot \hat{f}_{i_2} \cdots \hat{f}_{i_r} - \hat{h}_{i_1} \cdot \hat{h}_{i_2} \cdots \hat{h}_{i_r} \\
& = (\hat{f}_{i_1} - \hat{h}_{i_1}) \cdot \hat{f}_{i_2} \cdots \hat{f}_{i_r} + \hat{h}_{i_1} \cdot (\hat{f}_{i_2} - \hat{h}_{i_2}) \cdots \hat{f}_{i_r} + \cdots + \hat{h}_{i_1} \cdot \hat{h}_{i_2} \cdots (\hat{f}_{i_r} - \hat{h}_{i_r})
\end{aligned}$$

and using Lemma 2.6 for each term. Therefore,

$$\begin{aligned}
& \mathbb{E}_H \mathbb{E}_\ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(H \circ \ell(y_1), \dots, H \circ \ell(y_r))] \\
& \geq \mathbb{E}_\ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(f \circ \ell(y_1), \dots, f \circ \ell(y_r))] - \frac{1}{4} \geq \frac{3}{4} - \frac{1}{4} = \frac{1}{2}.
\end{aligned}$$

By Markov inequality,

$$\begin{aligned}
\frac{1}{4} &\leq \Pr_H \left[\mathbb{E}_{\ell} \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(H \circ \ell(y_1), \dots, H \circ \ell(y_r))] \geq \frac{1}{3} \right] \\
&\leq \Pr_H \left[\exists \ell \mathbb{E}_{(y_1, \dots, y_r) \sim \mathcal{M}} [\mathcal{D}_{y_1, \dots, y_r}(H \circ \ell(y_1), \dots, H \circ \ell(y_r))] \geq \frac{1}{3} \right] \\
&\leq \Pr_H \left[\exists \ell \Delta(H \circ \ell, \mathcal{C}) \leq \frac{\delta}{3} \right] && \text{(by the soundness of the tester)} \\
&= \Pr_H \left[\Delta(H, \mathcal{C}) \leq \frac{\delta}{3} \right] && \text{(since } \ell \text{ is invertible and } \mathcal{C} \text{ is affine invariant)}
\end{aligned}$$

Let $\mathcal{H} = \text{Supp}(H)$ be the set of words between f and g , i.e., the set of all words $e \in \Sigma^{\mathbb{K}^n}$ such that $e(x) = f(x)$ or $e(x) = g(x)$ for all $x \in \mathbb{K}^n$. We have $|\mathcal{H}| = 2^{\Delta(f,g)n}$. Since the distribution of H is uniform in \mathcal{H} , and we proved that at least $\frac{1}{4}$ fraction of words in \mathcal{H} contain a codeword in their $\delta/3$ neighborhood, let $\mathcal{H}' \subset \mathcal{H}$ denote this subset. Therefore, the $\delta/6$ neighborhoods around the points in \mathcal{H}' must be disjoint or else two distinct codewords will be $< \delta$ close to each other. The number of words in \mathcal{H} , which lie in a Hamming ball of radius $\delta/6$ around a point of \mathcal{H}' , is

$$\sum_{i=0}^{\delta n/6} \binom{\Delta(f,g)n}{i} \geq 2^{H(\delta/6)\Delta(f,g)\Delta(f,g)n - o(n)} \geq 2^{H(\delta/6)\Delta(f,g)n - o(n)},$$

where $H(\cdot)$ is the binary entropy function. By a packing argument, we can upper bound the size of \mathcal{H}' as

$$|\mathcal{H}'| \leq \frac{2^{\Delta(f,g)n}}{2^{H(\delta/6)\Delta(f,g)n - o(n)}} = o(|\mathcal{H}|).$$

This contradicts the fact that $|\mathcal{H}'| \geq |\mathcal{H}|/4$. \square

5. CONCLUSIONS

In this work, we proved tight lower bounds for constant query affine-invariant LCCs and LTCs when the number of queries r , underlying field \mathbb{K} and the alphabet Σ are constant. However, the constants in the bounds we obtain are of Ackermann-type in r , $|\mathbb{K}|$, $|\Sigma|$ because of the use of higher-order Fourier analysis. Improving the dependence on these parameters is an open problem that might require new ideas. In a recent work, Bhowmick and Lovett [2015a] obtain a “bias implies low rank” theorem for polynomials over growing fields. This might be a first step toward proving a variant of the inverse Gowers theorem (Lemma 2.5) for growing field size, which could then be used to make our lower bounds extend to the case of growing field size.

We also remark that our lower bounds work for any LCC or LTC where the queries are obtained as fixed linear combinations of uniformly chosen points from \mathbb{K}^n . Affine-invariant codes are a natural class of local codes where this is true. Relaxing these conditions to get lower bounds for a more general class of LCCs or LTCs is an open problem.

APPENDIX

A. PROOF OF GENERALIZED VON NEUMANN INEQUALITY (LEMMA 2.6)

Since the lemma is not stated in the form we want in Tao [2012], we will include a proof here for completeness. To prove Lemma 2.6, we need the following lemma first.

LEMMA A.1 (EXERCISE 1.3.22 IN TAO [2012]). *Let $f : \mathbb{K}^n \rightarrow \mathbb{C}$ be a function, and for each $1 \leq i \leq k$, let $g_i : (\mathbb{K}^n)^k \rightarrow \mathbb{C}$ be a bounded function which is independent of the i^{th} coordinate of $(\mathbb{K}^n)^k$. Then,*

$$|\mathbb{E}_{x_1, \dots, x_k \in \mathbb{K}^n} [f(x_1 + x_2 + \dots + x_k) \prod_{i=1}^k g_i(x_1, \dots, x_k)]| \leq \|f\|_{U^k}$$

PROOF. The proof is by induction on k and using Cauchy-Schwarz inequality repeatedly. The case $k = 1$ is true by definition of $\|\cdot\|_{U^1}$.

$$\begin{aligned} & \left| \mathbb{E}_{x_1, \dots, x_k \in \mathbb{K}^n} \left[f(x_1 + x_2 + \dots + x_k) \prod_{i=1}^k g_i(x_1, \dots, x_k) \right] \right| \\ &= \left| \mathbb{E}_{x_2, \dots, x_k} \left[g_1(x_1, \dots, x_k) \mathbb{E}_{x_1} \left[f(x_1 + x_2 + \dots + x_k) \prod_{i=2}^k g_i(x_1, \dots, x_k) \right] \right] \right| \\ & \hspace{15em} \text{(Since } g_1 \text{ doesn't depend on } x_1) \\ &\leq \left| \mathbb{E}_{x_2, \dots, x_k} \left[\mathbb{E}_{x'_1} \left[f(x'_1 + x_2 + \dots + x_k) \prod_{i=2}^k g_i(x'_1, x_2, \dots, x_k) \right] \right. \right. \\ & \hspace{10em} \left. \left. \cdot \mathbb{E}_{x_1} \left[\bar{f}(x_1 + x_2 + \dots + x_k) \prod_{i=2}^k \bar{g}_i(x_1, x_2, \dots, x_k) \right] \right] \right|^{1/2} \\ & \hspace{15em} \text{(By Cauchy-Schwarz inequality and the fact that } |g_1| \leq 1) \\ &= \left| \mathbb{E}_{x_1, h_1} \left[\mathbb{E}_{x_2, \dots, x_k} \left[\Delta_{h_1} f(x_1 + x_2 + \dots + x_k) \right. \right. \right. \\ & \hspace{10em} \left. \left. \cdot \prod_{i=2}^k g_i(x_1 + h_1, x_2, \dots, x_k) \bar{g}_i(x_1, x_2, \dots, x_k) \right] \right] \right|^{1/2} \\ & \hspace{15em} \text{(By substituting } x'_1 = x_1 + h_1) \\ &\leq \left| \mathbb{E}_{x_1, h_1} \left[\mathbb{E}_{h_2, \dots, h_k, z} \left[\Delta_{h_k} \dots \Delta_{h_1} f(x_1 + z) \right]^{1/2^{k-1}} \right] \right|^{1/2} \\ & \hspace{15em} \text{(By induction hypothesis and the definition of Gowers norm)} \\ &\leq \left| \mathbb{E}_{x_1, h_1, h_2, \dots, h_k, z} \left[\Delta_{h_k} \dots \Delta_{h_1} f(x_1 + z) \right] \right|^{1/2^k} \hspace{5em} \text{(By Jensen's inequality)} \\ &= \left| \mathbb{E}_{h_1, h_2, \dots, h_k, z} \left[\Delta_{h_k} \dots \Delta_{h_1} f(z) \right] \right|^{1/2^k} = \|f\|_{U^k} \quad \square \end{aligned}$$

PROOF OF LEMMA 2.6. By symmetry, it is enough to show that

$$|\mathbb{E}_{z_1, \dots, z_m \in \mathbb{K}^n} [f_0(\mathcal{L}_0(z_1, \dots, z_m)) \prod_{i=1}^k f_i(\mathcal{L}_i(z_1, \dots, z_m))]| \leq \|f_0\|_{U^k}.$$

We will make a linear change of variables so that we can use Lemma A.1 to get the required bound. For each $1 \leq i \leq k$, since \mathcal{L}_0 is not a multiple of \mathcal{L}_i , there exists a vector $v_i \in \mathbb{K}^m$ such that $\mathcal{L}_0(v_i) = 1$ and $\mathcal{L}_i(v_i) = 0$. Now, we make the following change of variables: $(z_1, \dots, z_m) \rightarrow (x_1, \dots, x_m) + \sum_{i=1}^k y_i v_i^T$ where x_1, \dots, x_m and y_1, \dots, y_k are the

new variables which range over \mathbb{K}^n .

$$\begin{aligned}
& \mathbb{E}_{z_1, \dots, z_m \in \mathbb{K}^n} [f_0(\mathcal{L}_0(z_1, \dots, z_m)) \prod_{i=1}^k f_i(\mathcal{L}_i(z_1, \dots, z_m))] \\
&= \left| \mathbb{E}_{x_1, \dots, x_m, y_1, \dots, y_k \in \mathbb{K}^n} \left[f_0 \left(\mathcal{L}_0(x_1, \dots, x_m) + \sum_{j \in [k]} y_j \right) \right. \right. \\
&\quad \left. \left. \prod_{i \in [k]} f_i \left(\mathcal{L}_i(x_1, \dots, x_m) + \sum_{j \in [k] \setminus \{i\}} y_j \mathcal{L}_i(v_j) \right) \right] \right| \\
&\quad \text{(By change of variables and linearity of } \mathcal{L}_i) \\
&\leq \mathbb{E}_{x_1, \dots, x_m \in \mathbb{K}^n} \left[\left| \mathbb{E}_{y_1, \dots, y_k \in \mathbb{K}^n} \left[f_0 \left(\mathcal{L}_0(x_1, \dots, x_m) + \sum_{j \in [k]} y_j \right) \right. \right. \right. \\
&\quad \left. \left. \prod_{i \in [k]} f_i \left(\mathcal{L}_i(x_1, \dots, x_m) + \sum_{j \in [k] \setminus \{i\}} y_j \mathcal{L}_i(v_j) \right) \right] \right| \right] \\
&\leq \|f_0\|_{U^k} \quad \square \quad \text{(By Lemma A.1)}
\end{aligned}$$

ACKNOWLEDGMENTS

We thank Madhu Sudan for helpful pointers to previous work. The second author would like to thank his advisor, Zeev Dvir, for his guidance and encouragement.

REFERENCES

- Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. 1998. Proof verification and the hardness of approximation problems. *J. ACM* 45, 3 (1998), 501–555.
- Sanjeev Arora and Shmuel Safra. 1998. Probabilistic checking of proofs: A new characterization of NP. *J. ACM* 45, 1 (1998), 70–122.
- Boaz Barak, Zeev Dvir, Amir Yehudayoff, and Avi Wigderson. 2011. Rank bounds for design matrices with applications to combinatorial geometry and locally correctable codes. *ACM*, 519–528.
- Omer Barkol, Yuval Ishai, and Enav Weinreb. 2007. On locally decodable codes, self-correctable codes, and t-private PIR. Springer, 311–325.
- Eli Ben-Sasson, Noga Ron-Zewi, and Madhu Sudan. 2012. Sparse affine-invariant linear codes are locally testable. *IEEE*, 561–570.
- Eli Ben-Sasson and Madhu Sudan. 2004. Robust locally testable codes and products of codes. 286–297.
- Eli Ben-Sasson and Madhu Sudan. 2008. Short PCPs with polylog query complexity. 38, 2 (2008), 551–607.
- Eli Ben-Sasson and Madhu Sudan. 2011. Limits on the rate of locally testable affine-invariant codes. Springer, 412–423.
- Arnab Bhattacharyya and Abhishek Bhowmick. 2015. Using higher-order Fourier analysis over general fields. *Preprint arXiv:1505.00619* (2015).
- Arnab Bhattacharyya, Zeev Dvir, Amir Shpilka, and Shubhangi Saraf. 2011. Tight lower bounds for 2-query LCCs over finite fields. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 638–647.
- Abhishek Bhowmick and Shachar Lovett. 2015a. Bias vs structure of polynomials in large fields, and applications in effective algebraic geometry and coding theory. *Preprint arXiv:1506.02047* (2015).
- Abhishek Bhowmick and Shachar Lovett. 2015b. The list decoding radius of Reed-Muller codes over small fields. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM.
- Manuel Blum and Sampath Kannan. 1995. Designing programs that check their work. *J. ACM* 42, 1 (1995), 269–291.

- Manuel Blum, Michael Luby, and Ronitt Rubinfeld. 1993. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences* 47, 3 (1993), 549–595.
- Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. 1998. Private information retrieval. *J. ACM* 45, 6 (1998), 965–981.
- Irit Dinur. 2007. The PCP theorem by gap amplification. *J. ACM* 54, 3 (2007), 12.
- Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. 2014. Breaking the quadratic barrier for 3-LCC's over the reals. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, 784–793.
- Zeev Dvir and Amir Shpilka. 2007. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. 36, 5 (2007), 1404–1434.
- Oded Goldreich, Howard Karloff, Leonard J. Schulman, and Luca Trevisan. 2012. Lower bounds for linear locally decodable codes and private information retrieval. In *Proceedings of the 17th IEEE Annual Conference on Computational Complexity*. IEEE, 175–183.
- Oded Goldreich and Madhu Sudan. 2006. Locally testable codes and PCPs of almost-linear length. *J. ACM* 53, 4 (July 2006), 558–655.
- William T. Gowers. 2001. A new proof of Szemerédi's theorem. 11, 3 (2001), 465–588.
- Ben Green. 2006. Montreal lecture notes on quadratic Fourier analysis. *Preprint arXiv:math/0604089* (2006).
- Alan Guo, Swastik Kopparty, and Madhu Sudan. 2013. New affine-invariant codes from lifting. ACM, 529–540.
- Alan Xinyu Guo. 2013. *Some Closure Features of Locally Testable Affine-invariant Properties*. Master's thesis. Massachusetts Institute of Technology.
- Venkatesan Guruswami, Madhu Sudan, Ameya Velingker, and Carol Wang. 2015. Limitations on testable affine-invariant codes in the high-rate regime. SIAM, 1312–1325.
- T. Kasami, S. Lin, and W. W. Peterson. 1967. Some results on cyclic codes which are invariant under the affine group and their applications. *Inform. and Comput.* 11, 5–6 (1967), 475–496.
- Jonathan Katz and Luca Trevisan. 2000. On the efficiency of local decoding procedures for error-correcting codes. ACM, 80–86.
- Tali Kaufman and Madhu Sudan. 2008. Algebraic property testing: The role of invariance. ACM, 403–412.
- Iordanis Kerenidis and Ronald de Wolf. 2003. Exponential lower bound for 2-query locally decodable codes via a quantum argument. ACM, 106–115.
- Richard J. Lipton. 1990. Efficient checking of computations. Springer, 207–215.
- Or Meir. 2009. Combinatorial construction of locally testable codes. 39, 2 (2009), 491–544.
- Madhu Sudan, Luca Trevisan, and Salil Vadhan. 1999. Pseudorandom generators without the XOR lemma. ACM, 537–546.
- Terence Tao. 2012. *Higher Order Fourier Analysis*. Vol. 142. American Mathematical Soc.
- Terence Tao and Tamar Ziegler. 2012. The inverse conjecture for the Gowers norm over finite fields in low characteristic. 16, 1 (2012), 121–188.
- Madhur Tulsiani and Julia Wolf. 2014. Quadratic Goldreich-Levin theorems. *SIAM J. Comput.* 43, 2 (2014), 730–766.
- Michael Viderman. 2015. Explicit strong LTCs with inverse poly-log rate and constant soundness. *Electronic Colloquium on Computational Complexity (ECCC)* 22 (2015), 20.
- David Woodruff. 2007. New lower bounds for general locally decodable codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 14.
- David P. Woodruff. 2012. A quadratic lower bound for three-query linear locally decodable codes over any field. *J. Comput. Sci. Technol.* 27, 4 (2012), 678–686.
- Sergey Yekhanin. 2011. Locally decodable codes. In *Computer Science—Theory and Applications*. Springer, 289–290.

Received February 2007; revised March 2009; accepted June 2009