List-Decodable Zero-Rate Codes

Noga Alon, Boris Bukh, and Yury Polyanskiy, *Senior Member, IEEE*

***Abstract*— We consider list decoding in the zero-rate regime for two cases—the binary alphabet and the spherical codes in Euclidean space. Specifically, we study the maximal $\tau \in [0, 1]$ for which there exists an arrangement of M balls of relative Hamming radius τ in the binary hypercube (of arbitrary dimension) with the property that no point of the latter is covered by L or more of them. As $M \to \infty$ the maximal τ decreases to a well-known critical value τ_L. In this paper, we prove several results on the rate of this convergence. For the binary case, we show that the rate is $\Theta(M^{-1})$ when L is even, thus extending the classical results of Plotkin and Levenshtein for $L = 2$. For $L = 3$, the rate is shown to be $\Theta(M^{-(2/3)})$. For the similar question about spherical codes, we prove the rate is $\Omega(M^{-1})$ and $O(M^{-(2L/L^2-L+2)})$.**

***Index Terms*— List decoding, error correction codes, Hamming space, Euclidean space.**

I. Introduction

THIS work concerns list-decoding under *worst-case* errors in the zero-rate regime. We consider the case of the binary alphabet in Sections I-VIII and the case of the unit sphere in Hilbert space in Section IX.

To motivate our study, we note that the maximum possible size of a code with given parameters is the most fundamental combinatorial problem in coding theory. Although the positive rate region is more interesting for applications, the range just above the threshold between positive rate and zero rate is intriguing and leads to challenging combinatorial problems, exhibiting, as we show in the paper, some interesting behavior.

Specifically, suppose we transmit a sequence of n symbols from $\{0, 1\}$ over a channel that can adversarially change less than a fraction τ of the symbols. The locations of corrupted symbols are unknown to the receiver. The goal is to find a code $C \subset \{0, 1\}^n$ such that the receiver can always produce a list of fewer than L messages containing the transmitted message. In other words, we seek a code C such that for every $w \in \{0, 1\}^n$ there are fewer than L codewords within Hamming distance τn from w. We call such code $<L$-*list-decodable* with radius τ. The largest such τ is denoted by $\tilde{\rho}_L(C)$ and is called the *(normalized) L-radius* of the code.

Let

$$\tau_L \stackrel{\text{def}}{=} \frac{1}{2} - \frac{\binom{2k}{k}}{2^{2k+1}} \quad \text{if } L = 2k \text{ or } L = 2k + 1. \tag{1}$$

It is known [3] that for radius $\tau < \tau_L$ the largest $<L$-list-decodable code is exponentially large in n, and for radius $\tau > \tau_L$ the largest $<L$-list-decodable code is of bounded size (independent of n). The aim of this paper is to understand how this constant varies as τ approaches τ_L from above. We define

$$\text{maxcode}_L(\varepsilon) \stackrel{\text{def}}{=} \max\big\{|C| : C \subset \{0, 1\}^n \text{ is } <L\text{-list-decodable of radius } \tau_L + \varepsilon\big\}.$$

Note that in this definition we do not restrict the block length n. The maximum is over all choices of $n \in \mathbb{N}$.

We are aware of three results on $\text{maxcode}_L(\varepsilon)$. First, a construction due to Levenshtein [9] shows that the so-called Plotkin bound is sharp in the unique decoding case, namely

$$\text{maxcode}_2(\varepsilon) = \frac{1}{4\varepsilon} + O(1).$$

Levenshtein's construction uses Hadamard matrices, and so requires ε to be of a special form. As a part of Theorem 1 below we present a construction without a condition on ε.

Second, Blinovsky [3] proved that $\text{maxcode}_L(\varepsilon)$ is finite for every L and every $\varepsilon > 0$. His proof iterates Ramsey's theorem, and gives a very large bound on $\text{maxcode}_L(\varepsilon)$ (which is not made explicit in the paper). Finally, in [5, Th. 1] Blinovsky claims an upper bound on $\text{maxcode}_L(\varepsilon)$ of the form $\text{maxcode}_L(\varepsilon) = O(1/\varepsilon)$. Below we construct a counterexample to this claim for $L = 3$.[1]

We next overview our results for the binary alphabet. The results for spherical codes are in Section IX.

A. Our Results (Binary Alphabet)

Our first result is a version of Levenshtein's construction for any fixed L. In comparison to Levenshtein's result, we have no restriction on ε, but our codes are longer (the value of n is larger).

Manuscript received November 4, 2017; revised May 16, 2018 and August 18, 2018; accepted August 23, 2018. Date of publication September 6, 2018; date of current version February 14, 2019. This work was supported in part by ISF under Grant 281/17, in part by GIF under Grant G-1347-304.6/2016, and in part by the Simons Foundation. N. Alon was supported in part by a BSF Grant, in part by an ISF Grant and a in part by GIF Grant. B. Bukh was supported in part by the Sloan Research Fellowship and in part by the U.S. taxpayers through NSF CAREER under Grant DMS-1555149, in part by NSF under Grant DMS-1301548, and in part by LabEx Bézout under Grant ANR-10-LABX-58. Y. Polyanskiy was supported in part by the NSF under Grant CCF-13-18620 and Grant CCF-17-17842 and in part by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under Grant CCF-09-39370. Part of the work was done when B. Bukh was on visit to the Université Paris-Est Marne-la-Vallée.

N. Alon is with the Department of Mathematics, Princeton University, Princeton, NJ 08544 USA, and also with the Schools of Mathematics and Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (e-mail: nogaa@tau.ac.il).

B. Bukh is with the Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: bbukh@math.cmu.edu).

Y. Polyanskiy is with the Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA 02139 USA (e-mail: yp@mit.edu).

Communicated by A. Rudra, Associate Editor for Complexity.

Digital Object Identifier 10.1109/TIT.2018.2868957

[1] The mistake appears to stem from the second paragraph of the proof of [5, Th. 1], which proposes a certain extension procedure for codes and claims that it does not decrease L-radius. A simple counter-example to the claim is a code $C = \{0, 1\}^2$ with 4-radius equal to 1, but its extension results in reduction of the 4-radius to $\frac{1}{2}$.

Theorem 1: Let $L \geq 2$, suppose m is a positive integer, and let $M = \binom{2m}{m}$. Let $c_L = 2^{-L}\lfloor L/2 \rfloor \binom{L-1}{\lfloor L/2 \rfloor}$. Then there exists an $<L$-list-decodable code in $\{0,1\}^M$ of size $2m$ and radius $\tau = \tau_L + c_L/2m + O(m^{-2})$. In particular,

$$\text{maxcode}_L(\varepsilon) \geq c_L \varepsilon^{-1} + O(1).$$

Theorem 2: Let $L \geq 2$ be even. Then

$$\text{maxcode}_L(\varepsilon) = O(\varepsilon^{-1})$$

We believe that in fact $\text{maxcode}_L(\varepsilon) = c_L\varepsilon^{-1} + O(1)$ for even L.

The case of odd L appears to be significantly harder: the principal reason for this is Lemma 8(b) below. By a different method, however, we were able to make progress for $L = 3$:

Theorem 3: We have $\text{maxcode}_3(\varepsilon) = \Theta(\varepsilon^{-3/2})$.

II. Overview of the Proof of Theorem 2

The bulk of the paper is devoted to the proof of Theorem 2. In this section, we explain the general structure of the argument.

A code with $\tilde{\rho}_L(C) = \tau$ is, equivalently, a code in which all L-tuples of distinct codewords have large circumscribed radius, where the latter is defined as the radius of the smallest ball (in Hamming space) containing the L-tuple. First of all, we show (Prop. 4) that the center of the ball can as well be sought after in $[0,1]^n$, instead of $\{0,1\}^n$. The resulting quantity is denoted $\text{rad}(x)$, cf. (4). The resulting linear relaxation is much easier to analyze.

Second, we introduce mean radius $\text{mrad}_\omega(x)$ of an L-tuple x with respect to a measure ω on $\{1,2,\ldots,L\}$. The significance of this quantity is that there exists (Lemma 6) finitely many ω's such that

$$\text{rad}(x) = \max_{\omega \in \Omega'_L} \text{mrad}_\omega(x) \quad \forall x \in (\{0,1\}^n)^L. \tag{2}$$

Among these ω's a special one is $\mathcal{U}[L]$ – a uniform distribution on $[L]$.

The mean radii of random independent bits is denoted as $\tau_{\omega,p} = \mathbb{E}[\text{mrad}_\omega(X_1,\ldots,X_L)]$, where $X_i \overset{i.i.d.}{\sim} \text{Ber}(p)$. Properties of $\tau_{\omega,p}$ are summarized in Lemma 8, where two crucial ones are that $\tau_L = \tau_{\mathcal{U}[L],1/2}$ and that for even L (!) we have $\tau_{\omega,p} \leq \tau_L - \delta_0$ for some $\delta_0 > 0$ and all $\omega \neq \mathcal{U}[L], \omega \in \Omega'_L$, $p \in [0,1]$.

To understand why $\text{maxcode}_\epsilon(L) = \Theta(\frac{1}{\epsilon})$ is a natural guess, we recall a standard averaging argument (Corollary 10):

$$\frac{1}{\binom{M}{L}} \sum_{x \in \binom{C}{L}} \text{mrad}_\omega(x) \leq \tau_{\omega,1/2} + O(1/M). \tag{3}$$

Thus, averaged over the code the mean radii (with respect to any ω) can never exceed $\tau_L + O(1/M)$. Our proof shows that (for even L) and every sufficiently large code with $\tilde{\rho}_L(C) \geq \tau_L$ averaging mean radii (with $\omega = \mathcal{U}[L]$) is equivalent to averaging the actual radius.

With these preparations our proof proceeds as follows:

a) We start with showing that a biased code (i.e. one in which the fraction of 1's among all codewords is $\leq 1/2 - \delta_0$) with radius $\tilde{\rho}_L(C) \geq \tau_L$ can have at most finite size. This is the content of Lemma 13, which is proved by appealing to a Ramsey theorem to extract a large subcode with the property that $\text{rad}(x) = \text{mrad}_\omega(x)$ for every L-tuple x and some fixed ω. Invocation of a biased-version of (3) then shows this subcode cannot be too large.

b) Next, we show (Lemma 11) that there is a dichotomy: either a code has a large biased subcode (thus contradicting above), or all but $O(1/M)$ fraction of its L-tuples must have random-like coordinate composition (i.e. out of n coordinates about $2^{-L}n$ have zeros in all L codewords, and the same holds for each binary pattern). In particular, this implies that every such L-tuple has $\text{mrad}_\omega(x) \approx \tau_{\omega,1/2}$.

c) Finally, since every $\omega \neq \mathcal{U}[L]$ yields $\text{mrad}_\omega(x) \approx \tau_{\omega,1/2} < \tau_L$ we must have that $1 - O(1/M)$ fraction of L-tuples satisfy $\text{rad}(x) = \text{mrad}_{\mathcal{U}[L]}(x)$ and hence in averaging (3) we can replace $\text{mrad}_\omega(x)$ with $\text{rad}(x)$ and conclude that $\text{rad}(x)$ of a typical L-tuple is at most $\tau_L + O(1/M)$ as claimed.

We mention that our proof of Theorem 2 also shows why Theorem 3 is perhaps surprising: for $L = 3$ we construct a code such that averaged over the codebook $\text{mrad}_\omega(x)$ is always $\leq \tau_L + O(1/M)$ (as it should be), yet the $\text{rad}(x)$ of every L-tuple is $\geq \tau_L + O(1/M^{2/3})$. The proof of Theorem 3 relies on a special relation for radii of triangles in ℓ_1-spaces, see Prop. 7.

III. Mean Radii

a) Definitions: For $x \in \mathbb{R}^n$, let $\|x\| \overset{\text{def}}{=} (1/n)\sum|x_i|$. In particular, for $x, y \in \{0,1\}^n$ the quantity $\|x - y\|$ is the (normalized) Hamming distance between bit strings x and y.

For points $x^{(1)},\ldots,x^{(L)} \in \{0,1\}^n$ let

$$\text{rad}(x^{(1)},\ldots,x^{(L)}) = \min_{y \in [0,1]^n} \max_i \|x^{(i)} - y\|. \tag{4}$$

Note that we allow the coordinates of y to be arbitrary real numbers between 0 and 1. For example, for $C = \{000, 100\}$ we have $\text{rad}(C) = 1/6$, but for every $y \in \{0,1\}^3$ one of the points of C is at distance at least $1/3$. However, this relaxation makes only slight effect, as the next proposition shows.

Proposition 4: Let $x^{(1)},\ldots,x^{(L)} \in \{0,1\}^n$. If $\tau = \text{rad}(x^{(1)},\ldots,x^{(L)})$, then there is a point $y \in \{0,1\}^n$ such that $\|x^{(i)} - y\| \leq \tau + \frac{L}{2n}$ for all i.

Proof: For any bit $z \in \{0,1\}$ and real $w \in [0,1]$ define $\ell(z,w) = w$ if $z = 0$ and $\ell(z,w) = 1 - w$ if $z = 1$. Note that with this notation, for every i, $\|x^{(i)} - y\| = \frac{1}{n}\sum_{j=1}^n \ell(x_j^{(i)}, y_j)$ is an affine function of the variables y_j.

The assumption $\text{rad}(x^{(1)},\ldots,x^{(L)}) \leq \tau$ is equivalent to the fact that the polytope in the variables y_j defined by the inequalities $0 \leq y_j \leq 1$ for all $1 \leq j \leq n$ and $\frac{1}{n}\sum_{j=1}^n \ell(x_j^{(i)}, y_j) \leq \tau$ for all $1 \leq i \leq L$ is nonempty. Hence it contains a vertex $y' = (y'_1,\ldots,y'_n)$. In this vertex there are at most L variables y'_j which are neither 0 nor 1, and the desired result is obtained by rounding each such y'_j to the closest integer y_j and by taking $y_j = y'_j$ for all other coordinates y'_j. □

Let Ω_L be the set of all probability measures on the set $[L] \stackrel{\text{def}}{=} \{1, 2, \ldots, L\}$. Suppose $\omega \in \Omega_L$ is a probability distribution on $[L]$. Then for an L-tuple $x = (x^{(1)}, \ldots, x^{(L)}) \in (\{0,1\}^n)^L$ of codewords, we define their *mean radius* with respect to ω by

$$\operatorname{mrad}_\omega(x) \stackrel{\text{def}}{=} \min_{y \in [0,1]^n} \mathop{\mathbb{E}}_{i \sim \omega} \|x^{(i)} - y\|. \tag{5}$$

Because $\mathbb{E}_{i\sim\omega}\|x^{(i)} - y\|$ can be written as a sum over the individual coordinates, the y attaining minimum in (5) may be taken to have all of its coordinates in $\{0, 1\}$. This leads to an alternative formula for $\operatorname{mrad}_\omega(x)$:

$$\begin{aligned} \operatorname{mrad}_\omega(x) &= \mathbb{E}_{j\in[n]} \min\left(\textstyle\sum_{x_j^{(i)}=0} \omega_i, \sum_{x_j^{(i)}=1} \omega_i\right) &(6)\\ &= \tfrac{1}{2} - \tfrac{1}{2}\,\mathbb{E}_{j\in[n]} \left|\textstyle\sum_{x_j^{(i)}=0} \omega_i - \sum_{x_j^{(i)}=1} \omega_i\right|. &(7)\end{aligned}$$

b) Duality: From the comparison of (4) and (5) it is clear that $\operatorname{rad}(x) \geq \operatorname{mrad}_\omega(x)$ for any ω. The key observation is that a suitable converse holds as well.

Lemma 5: For every $x = (x^{(1)}, \ldots, x^{(L)}) \in (\{0,1\}^n)^L$ we have

$$\operatorname{rad}(x) = \max_{\omega \in \Omega_L} \operatorname{mrad}_\omega(x), \tag{8}$$

where the maximum is over all probability measures ω on $\{1, 2, \ldots, L\}$.

Proof: Notice that (4) can be rewritten as

$$\operatorname{rad}(x) = \min_{y\in[0,1]^n} \max_{\omega\in\Omega_L} \mathop{\mathbb{E}}_{i\sim\omega} \|x^{(i)} - y\|.$$

Since the function

$$(y, \omega) \mapsto \mathop{\mathbb{E}}_{i\sim\omega} \|x^{(i)} - y\|,$$

is convex in y and affine in ω, von Neumann minimax theorem [13] implies

$$\min_{y\in[0,1]^n} \max_{\omega\in\Omega_L} \mathop{\mathbb{E}}_{i\sim\omega} \|x^{(i)} - y\| = \max_{\omega\in\Omega_L} \min_{y\in[0,1]^n} \mathop{\mathbb{E}}_{i\sim\omega} \|x^{(i)} - y\|.$$

Comparing with (5) completes the proof. □

Lemma 6: For every L there exists a finite set of probability measures $\Omega'_L \subset \Omega_L$ such that

$$\operatorname{rad}(x) = \max_{\omega\in\Omega'_L} \operatorname{mrad}_\omega(x) \quad \text{for all } x \in (\{0,1\}^n)^L. \tag{9}$$

Furthermore, $|\Omega'_L| \leq 4^{L^2}$.

Proof: Let $x \in \left(\{0,1\}^n\right)^L$ be any L-tuple of words. To each coordinate $j \in [n]$ we can then associate the bit string

$$T(j) \stackrel{\text{def}}{=} \left(x_j^{(1)}, x_j^{(2)}, \ldots, x_j^{(L)}\right).$$

For a bit string $T \in \{0,1\}^L$, put $P_T = \{j \in [n] : T(j) = T\}$.

Let y be a point that achieves minimum in (4). For each T, replace coordinates of y indexed by P_T by their average. This does not change $\|x^{(i)} - y\|$ and so the obtained point also achieves minimum in (4). So, we may assume that y_j depends only on $T(j)$. Let $\alpha_T = |P_T|/n$.

To each probability measure $\omega \in \Omega_L$ we can associate a *signature*, which is a function $S_\omega\colon \{0,1\}^L \to \{1, -1\}$ defined by

$$S_\omega(T) \stackrel{\text{def}}{=} \operatorname{sign}\left(\sum_{i:T_i=0} \omega_i - \sum_{i:T_i=1} \omega_i\right) \quad \text{for } T \in \{0,1\}^L.$$

Note that $\omega \to S_\omega$ is a sign of a linear function on the Ω_L, which we identify with a simplex in $\mathbb{R}^{L-1}$. Since 2^L hyperplanes partition $\mathbb{R}^{L-1}$ into at most $\sum_{j\leq L-1} \binom{2^L}{j} \leq 2^{L^2}$ regions, the number of possible signatures is at most 2^{L^2}. For each possible signature S, let $\Omega_S \stackrel{\text{def}}{=} \{\omega \in \Omega_L : S_\omega = S\}$. Since Ω_S is an intersection of halfspaces, which, in addition to $S_\omega = S$, includes the additional inequalities $\omega_i \geq 0$ for all i and $\sum_i \omega_i = 1$, it is a convex polytope.

By the preceding lemma and (6), we want to maximize $\operatorname{mrad}_\omega(x)$ over all probability measures $\omega \in \Omega_L$. However, while ω ranges over Ω_S the values of all $S_\omega(T), T \in \{0,1\}^L$ stay constant. Thus the maximum over $\omega \in \Omega_S$ of $\operatorname{mrad}_\omega(x)$ is the maximum of the following linear function in the variables ω_i:

$$\sum_T \alpha_T f(\omega, T)$$

where

$$f(\omega, T) = \sum_{i:T_i=0} \omega_i \quad \text{if } S(T) = -1$$

and

$$f(\omega, T) = \sum_{i:T_i=1} \omega_i \quad \text{if } S(T) = +1,$$

where $S(T)$ denotes the constant signature $S_\omega(T)$. This maximum is attained at a vertex of the polytope Ω_S. Thus, we may take Ω'_L to be the union of the vertex sets of all polytopes Ω_S, for all signatures S.

Each Ω_S is defined by $m \stackrel{\text{def}}{=} L + 2^L$ inequalities, and so by McMullen's upper bound theorem has at most $\binom{m-\lceil L/2\rceil}{\lfloor L/2\rfloor} + \binom{m-\lfloor L/2\rfloor-1}{\lceil L/2\rceil-1} \leq 2^{L^2}$ vertices. Multiplying by the 2^{L^2} possible signatures, we obtain the result. □

The preceding proof gives an algorithm to compute the set Ω'_L. The results of this computation for small L can be found at http://www.borisbukh.org/code/listdecoding17.html. Interestingly, for $L \leq 4$ the result is very nice. For a set $R \subset [L]$ let $\operatorname{mrad}_R(x)$ be the $\operatorname{mrad}_\omega(x)$ for the probability measure ω that is uniform on R, i.e., $\omega_i = 1/|R|$ if $i \in R$. Then for any $x \in (\{0,1\}^n)^L$

$$\operatorname{rad}(x) = \max_{|R| \text{ is even}} \operatorname{mrad}_R(x) \quad \text{if } L \leq 4. \tag{10}$$

Our proof of Theorem 3 will use (10) with $L = 3$ and so we establish it formally (generalized to arbitrary ℓ_1-vectors).

Proposition 7: For any set of three vectors x, y, z in $\mathbb{R}^n$ with respect to the ℓ_1-norm, $\operatorname{rad}(x, y, z) = \frac{1}{2}\operatorname{diam}(x, y, z)$.

Proof: Put $x = (x_1, x_2, \ldots, x_n)$, $y = (y_1, y_2, \ldots, y_n)$, $z = (z_1, z_2, \ldots, z_n)$. Let d be the diameter of the set $\{x, y, z\}$. For each i let m_i be the median of x_i, y_i, z_i and define $m = (m_1, m_2, \ldots, m_n)$. Put, also, $a = \|x - m\|, b = \|y - m\|$,

$c = \|z - m\|$, where $\|\cdot\|$ is the ℓ_1-norm. Note that, crucially $\|x - y\| = a + b$, $\|y - z\| = b + c$, $\|x - z\| = a + c$. Thus each of these three sums is at most d. If each of the quantities a, b, c is at most $d/2$ then m is a center of an ℓ_1-ball of radius $d/2$ containing x, y, z, showing that in this case the radius is indeed $d/2$, as needed. Otherwise one of the above, say a, is larger than $d/2$. In this case define $w = (1 - \frac{d}{2a})x + \frac{d}{2a}m$. It is easy to check that $x - w = \frac{d}{2a}(x - m)$ and hence $\|x - w\| = \frac{d}{2a}a = d/2$. In addition $m - w = (1 - \frac{d}{2a})(m - x)$ and hence $\|m - w\| = a - \frac{d}{2}$.

Thus, by the triangle inequality, $\|y - w\| \le \|y - m\| + \|m - w\| = b + a - \frac{d}{2} \le \frac{d}{2}$, and similarly $\|z - w\| \le \frac{d}{2}$, completing the proof. □

IV. Averaging

Averaging arguments play a major role in this paper. We collect them in this section.

A. Mean Radii of Random Bit Strings

Averaging arguments will allow us to show that, in a large code C, mean radii of codewords from C rarely exceed the mean radius of random bit strings. Because of that, we start by computing the mean radius of random bit strings for arbitrary probability measure ω. In particular, we will see that the radius threshold τ_L defined in (1) is the average radius of a random L-tuple of bit strings.

Call a random string $w \in \{0, 1\}^n$ *p-biased* if each bit of w is 1 with probability p and 0 with probability $1 - p$, and the bits are independent of each other. For a probability measure ω on $[L]$, we let

$$\tau_{\omega,p} \stackrel{\text{def}}{=} \mathbb{E}\,\mathrm{mrad}_\omega(w^{(1)}, \ldots, w^{(L)}),$$

where $w^{(1)}, \ldots, w^{(L)} \in \{0, 1\}^n$ are independent p-biased. For brevity, write τ_ω in lieu of $\tau_{\omega,1/2}$. Note that $\tau_{\omega,p}$ is independent of n, in view of (6).

Let $\mathcal{U}[L]$ denote the uniform probability measure on $[L]$.

Lemma 8:

a) *For every probability measure ω on $[L]$ and every p we have*

$$\tau_{\omega,p} \le \tau_{\mathcal{U}[L],p}.$$

b) *If L is even and $0 < p < 1$, then the equality holds if and only if $\omega = \mathcal{U}[L]$.*

c) *$\tau_{\mathcal{U}[L],1/2} = \tau_L$, where τ_L is defined in* (1)*.*

d) *We have $\tau_{\omega,p} < \tau_L$ whenever $p \ne \frac{1}{2}$.*

Proof:

Part (a)*:* Given an ω and a vector of signs $\varepsilon \stackrel{\text{def}}{=} (\varepsilon_1, \ldots, \varepsilon_L) \in \{1, -1\}^L$, define signed sum $f_\omega(\varepsilon) \stackrel{\text{def}}{=} \sum_i \varepsilon_i \omega(i)$. By (7), maximization of $\tau_{\omega,p}$ is equivalent to minimization of

$$\mathop{\mathbb{E}}_{\varepsilon} |f_\omega(\varepsilon)| \quad \text{for } p\text{-biased } \varepsilon \in \{1, -1\}^L,$$

i.e. where $\Pr[\varepsilon_i = -1] = p$.

Let Ω_p^* be the set of all probability measures on $[L]$ that maximize $\tau_{\omega,p}$. (The maximum is achieved because $\omega \mapsto \tau_{\omega,p}$ is a continuous function on a compact set.) Let $\omega \in \Omega_p^*$ be arbitrary. Suppose ω is not uniform. Without loss generality, $\omega(L - 1) \ne \omega(L)$. Let ω' be obtained from ω by replacing the values of $L - 1$ and L by their averages. If $\varepsilon_{L-1} = \varepsilon_L$, then $f_\omega(\varepsilon) = f_{\omega'}(\varepsilon)$. Also,

$$\begin{aligned} &|f_\omega(\varepsilon_1, \ldots, \varepsilon_{L-2}, 1, -1)| + |f_\omega(\varepsilon_1, \ldots, \varepsilon_{L-2}, -1, 1)| \\ &\quad \ge |f_{\omega'}(\varepsilon_1, \ldots, \varepsilon_{L-2}, 1, -1)| + |f_{\omega'}(\varepsilon_1, \ldots, \varepsilon_{L-2}, -1, 1)|. \end{aligned} \tag{11}$$

with equality only if $|\sum_{i \le L-2} \varepsilon_i \omega(i)| \ge |\omega(L - 1) - \omega(L)|$. Indeed, denoting $a = \sum_{i \le L-2} \varepsilon_i \omega(i)$ and $b = \omega(L - 1) - \omega(L)$ the inequality (11) is just

$$|a + b| + |a - b| \ge 2|a|,$$

which is a consequence of convexity of $|\cdot|$. Furthermore, it is clear that equality holds only when $a + b$ and $a - b$ have the same sign, i.e. $|a| \ge |b|$.

Since $\omega \in \Omega_p^*$, it follows that the equality does hold in (11) for every ε, and that $\omega' \in \Omega_p^*$ as well. From the condition for equality, we deduce that for any $\omega \in \Omega_p^*$ we have

$$|\omega(i) - \omega(i')| \le \min_{\varepsilon} \Big|\sum_{j \notin \{i, i'\}} \varepsilon_j \omega(j)\Big| \quad \text{for all } i \ne i' \tag{12}$$

From continuity of $\omega \mapsto \tau_\omega$, it follows by repeated pairwise averaging, that if ω' is obtained from ω by replacing the values of ω on any subset of $[L]$ by their averages, then $\omega' \in \Omega_p^*$ as well. In particular, $\mathcal{U}[L] \in \Omega_p^*$ and so (a) holds.

Part (b)*:* Suppose that L is even and (b) does not hold. Let $\omega \in \Omega_p^*$ be non-uniform. Without loss of generality, $\omega(L - 1) \ne \omega(L)$. Let ω' be obtained from ω by replacing values on $[L - 2]$ by their averages. Since $\sum_{j \le L-2} (-1)^j \omega'(j) = 0$, the measure ω' fails (12), and so $\omega' \notin \Omega_p^*$. Thus, $\omega \notin \Omega_p^*$ and hence, $\Omega_p^* = \{\mathcal{U}[L]\}$, as claimed by (b).

Part (c)*:* Consider a random walk on $\mathbb{Z}$ starting from 0 that makes a step to the right with probability p and to the left with probability $1 - p$. Let $\Delta_{s,p}$ be the position of this walk after s steps. Relation (7) implies that

$$\tau_{\mathcal{U}[L],p} = \tfrac{1}{2} - \tfrac{1}{2L}\,\mathbb{E}|\Delta_{L,p}|. \tag{13}$$

From (13) and [14, eqs. (23) and (32)] we obtain the formula (1) for τ_L.

Part (d)*:* In view of (a) we may restrict ourselves to the case $\omega = \mathcal{U}[L]$. Because of (13) and the symmetry under $p \to (1 - p)$, it suffices to prove that

$$\Pr[|\Delta_{s,p}| \ge k] \ge \Pr[|\Delta_{s,1/2}| \ge k] \quad \forall p > \tfrac{1}{2}, \quad s \ge 2 \tag{14}$$

and that (14) is a strict inequality for some k. In fact, we will show that the inequality is strict whenever $k \ge 2$ and $s \equiv k \pmod 2$.

Since $\Delta_{s,p}$ and $\Delta_{s,1/2}$ are of the same parity as s, it suffices to consider only the case $s \equiv k \pmod 2$. If $k = 0$ or $k = 1$ and $s \equiv k \pmod 2$, then both sides of (14) are equal to 1. If $(s, k) = (2, 2)$, then the inequality (14) is strict because $p^2 + (1 - p)^2 > 1/2$ for $p > 1/2$.

The general case follows by induction on s from

$$\begin{aligned}&\Pr\big[|\Delta_{s+1,p}| \ge k\big]\\ &\quad= \tfrac{1}{2}\Pr\big[|\Delta_{s,p}| \ge k-1\big] + \tfrac{1}{2}\Pr\big[|\Delta_{s,p}| \ge k+1\big]\\ &\qquad+ \big(p-\tfrac{1}{2}\big)\big(\Pr[\Delta_{s,p} \in \{k, k-1\}]\\ &\qquad- \Pr[\Delta_{s,p} \in \{-k, -k+1\}]\big),\end{aligned}$$

which is valid for $k \ge 2$. □

B. Mean Radii in Large Codes

Here we show that the average $\mathrm{mrad}_\omega(\cdot)$ over L-tuples in a large $C \subset \{0,1\}^n$ can be only slightly larger than τ_ω. In fact, we will show a generalization of this to codes of possibly small radius.

Lemma 9: Let ω be a probability measure on $[L]$. Suppose $C \subset \{0,1\}^n$ satisfies $\mathrm{rad}(C) \le p \le \frac{1}{2}$. Then

$$\mathop{\mathbb{E}}_{w^{(1)},\ldots,w^{(L)}\in C} \mathrm{mrad}_\omega(w^{(1)},\ldots,w^{(L)}) \le \tau_{\mathcal{U}[L],p},$$

where the expectation is over uniformly and independently chosen codewords $w^{(1)},\ldots,w^{(L)}$ from C.

Proof: Let $p_j = \Pr_{w\in C}[w_j = 1]$. We have $\mathrm{mrad}_{\mathcal{U}[C]}(C) \le p$ from (8). Recall that one can always take y attaining minimum in the definition of mrad to have all its coordinates in $\{0,1\}$. So, without loss of generality (otherwise invert some coordinates), we may assume that y attaining $\mathrm{mrad}_{\mathcal{U}[C]}(C)$ in (6) is $y = 0$. Then we have $\frac{1}{n}\sum_{j\in[n]} p_j \le p$. Denote by $B(q)$ the distribution on $\{1,-1\}^L$ where each coordinate is independently 1 with probability q and -1 with probability $1-q$. Given a vector of signs $\varepsilon = (\varepsilon_1,\ldots,\varepsilon_L) \in \{1,-1\}^L$ define $f_\omega(\varepsilon) \stackrel{\mathrm{def}}{=} \sum_i \varepsilon_i \omega(i)$.

From (7) and the proof of part (a) of Lemma 8 we then have

$$\begin{aligned}1 - 2\mathop{\mathbb{E}}_{w\in C^L} \mathrm{mrad}_\omega(w) &= \mathbb{E}_{j\in[n]}\,\mathbb{E}_{\varepsilon\sim B(p_j)}\,|f_\omega(\varepsilon)|\\ &\ge \mathbb{E}_{j\in[n]}\,\mathbb{E}_{\varepsilon\sim B(p_j)}\left|f_{\mathcal{U}[L]}(\varepsilon)\right|.\end{aligned}$$

By [10, Lemma 8], the function $p \mapsto \mathbb{E}_{\varepsilon\sim B(p)}\left|f_{\mathcal{U}[L]}(\varepsilon)\right|$ is convex. Jensen's inequality and the fact that $\tau_{\mathcal{U}[L],p}$ is an increasing function of p on $[0,\frac{1}{2}]$ then complete the proof. □

Corollary 10: Let ω be a probability measure on $[L]$. Suppose $C \subset \{0,1\}^n$ is of size $|C| \ge L^2 M$ and satisfies $\mathrm{rad}(C) \le p$. Then there is an L-tuple $w \in C^L$ with distinct codewords such that $\mathrm{mrad}_\omega(w) \le \tau_{\mathcal{U}[L],p} + 1/M$.

Proof: Let $X \subset C^L$ be the set of all L-tuples with distinct codewords. The corollary follows from $\Pr[w \notin X] \le \binom{L}{2}/|C|$ and $\mathbb{E}_{w\in C}\,\mathrm{mrad}_\omega(w) \ge \Pr[w \in X]\,\mathbb{E}_{w\in X}\,\mathrm{mrad}_\omega(w)$. □

V. Abundance of Random-Like L-Tuples

Lemma 11: Let $\pi : \mathbb{R}^n \to \mathbb{R}^m$ be an orthogonal projection on a set of m coordinates. Suppose that $C \subset \{0,1\}^n$ satisfies $\mathrm{rad}(\pi(C)) \le \frac{1}{2} - \varepsilon$. Then there is a $C' \subset C$ of size $|C'| \ge |C|/2$ satisfying $\mathrm{rad}(C') \le \frac{1}{2} - \frac{m}{n}\varepsilon$.

Proof: Let π' be the projection on the remaining $n-m$ coordinates. Classify codewords $c \in C$ based on whether $\|\pi'(c)\| \le \frac{1}{2}$ or $> \frac{1}{2}$. Without loss of generality, at least half of them (call it C') satisfy $\|\pi'(c)\| \le \frac{1}{2}$. Let $y_1 \in \mathbb{R}^m$ be the center attaining $\mathrm{rad}(\pi(C))$ and define $y \in \mathbb{R}^n$ to be the solution to $\pi(y) = y_1$, $\pi'(y) = 0$. We have for any $c \in C'$

$$\begin{aligned}\|y - c\| &= \tfrac{m}{n}\|\pi(y) - \pi(c)\| + \tfrac{n-m}{n}\|\pi'(y) - \pi'(c)\|\\ &\le \tfrac{m}{n}\left(\tfrac{1}{2} - \varepsilon\right) + \tfrac{n-m}{2n} = \tfrac{1}{2} - \tfrac{m}{n}\varepsilon.\end{aligned}$$

□

For an L-tuple $x = (x^{(1)},\ldots,x^{(L)}) \in (\{0,1\}^n)^L$, we define $\mathrm{type}(x) \stackrel{\mathrm{def}}{=} \big(\mathrm{type}(x)_u\big)_{u\in\{0,1\}^L}$ to be the probability distribution on $\{0,1\}^L$ with $\mathrm{type}(x)_u \stackrel{\mathrm{def}}{=} \frac{1}{n}\#\{j : x_j^{(i)} = u_i, \forall i \in [L]\}$ (Note that $\mathrm{type}(x)_u = \alpha_T$ with $T = u$ in the notation of Lemma 6). The next result shows that the only obstruction to finding a large number of L-tuples with approximately uniform $\mathrm{type}(x)$ is the existence of a large biased subcode.

Lemma 12: Let L be fixed. For every $\varepsilon > 0$ there is a $\delta > 0$ with the following property. If s is a natural number, there exist constants $M_0 = M_0(s)$ and $c = c(s)$ such that for any code $C \subset \{0,1\}^n$ with $M \stackrel{def}{=} |C| \ge M_0$ one of the following two alternatives must hold:

a) *$\exists C' \subset C$ such that $|C'| \ge s$ and $\mathrm{rad}(C') \le \frac{1}{2} - \delta$, or*
b) *there exist at least $M^L - cM^{L-1}$ many L-tuples of distinct codewords from C such that*

$$|\,\mathrm{type}(x)_u - 2^{-L}| \le \varepsilon \qquad \forall u \in \{0,1\}^L \tag{15}$$

and, in particular

$$|\,\mathrm{mrad}_\omega(x) - \tau_{\omega,1/2}| \le 2^L\varepsilon \quad \text{for every } \omega \in \Omega_L \tag{16}$$

for each of these L-tuples x.

Consequently, if C does not satisfy a), then the number of L-tuples of distinct codewords of C violating (15) is of size at most cM^{L-1}.

Proof: Set $2\delta_0 \stackrel{\mathrm{def}}{=} (1+2^L\varepsilon)^{1/L} - 1$ and note that with this choice we have $|(\frac{1}{2} \pm \delta_0)^L - 2^{-L}| \le \varepsilon$. Set also $\mu \stackrel{\mathrm{def}}{=} (\frac{1}{2} - \delta_0)^L$ and $\delta \stackrel{\mathrm{def}}{=} \delta_0\mu$. Finally, set $c(s) \stackrel{\mathrm{def}}{=} s2^{L+3}$ and $M_0(s) \stackrel{\mathrm{def}}{=} s2^{L+3}$. Note that (15) implies (16) via (6), and so we only consider (15) below.

Let us assume that a) does not hold. Then in any C'' with $|C''| \ge 4s$ and for any orthogonal projection π_A on a subset of coordinates $A \subset [n]$ there must exist a $c \in C''$ such that

$$\|\pi_A(c)\| \in (1/2 - \delta_0, 1/2 + \delta_0), \tag{17}$$

provided that $\delta_0\frac{|A|}{n} \ge \delta$. Indeed, if all $c \in C''$ violate (17), then at least half of $c \in C$ should satisfy $\|\pi_A(c)\| \le 1/2 - \delta_0$, say. Denote this collection by C_0 and observe that $\mathrm{rad}(\pi_A(C_0)) \le 1/2 - \delta_0$ and $|C_0| \ge 2s$. By Lemma 11 there must exist $C' \subset C_0$ of size $\ge s$ such that $\mathrm{rad}(C') \le \frac{1}{2} - \delta$, contradicting assumption.

It similarly follows that for any collection of subsets $A_1,\ldots,A_r$ with $|A_j| \ge \mu n$ for all $j \in [r]$, and any C'' with $|C''| \ge 4sr$ there must exist $c \in C''$ such that (17) holds simultaneously for all $A = A_j$, $j \in [r]$. Indeed, a given $c \in C''$ can violate (17) in $2r$ ways (two ways for each A_j). By the pigeonhole principle, if all codewords in C'' fail (17), then there are $2s$ that fail in the same way, and then we proceed as in the case above.

We next show that there are more than

$$N_1 = \prod_{j=0}^{L-1}(M - j - 4s \cdot 2^j)$$

L-tuples x of distinct codewords from C that satisfy (15). Indeed, at least $M - 4s$ codewords $x^{(1)}$ have $|\|x^{(1)}\| - \frac{1}{2}| \le \delta_0$. Once one such codeword $x^{(1)}$ is selected, let $A_0 = \{j \in [n] : x_j^{(1)} = 0\}$ and $A_1 = A_0^c$. Each of these two subsets has cardinality $\ge n(\frac{1}{2} - \delta_0) \ge \mu n$. By the argument above, there are more than $M - 1 - 4s \cdot 2$ codewords $x^{(2)}$ not equal to $x^{(1)}$ such that projections of $x^{(2)}$ on A_0 and A_1 both have weights $\in [\frac{1}{2} - \delta_0, \frac{1}{2} + \delta_0]$. Selecting one such $x^{(2)}$, we define partitions $A_0 = A_{00} \cup A_{01}$ and $A_1 = A_{10} \cup A_{11}$ according to the values of coordinates of $x^{(1)}$ and $x^{(2)}$. Proceeding similarly, we construct $x^{(3)}, \dots, x^{(L)}$. The resulting L-tuple has distinct elements and satisfies

$$(\tfrac{1}{2} - \delta_0)^L \le \mathrm{type}(x)_u \le (\tfrac{1}{2} + \delta_0)^L \qquad \forall u \in \{0,1\}^L,$$

which by the choice of δ_0 implies that it satisfies (15) as well.

Note that for $M \ge \max_j k_j$ we have

$$\prod_{j=0}^{L-1}(M - k_j) = M^L \prod_{j=0}^{L-1}(1 - k_j/M) \ge M^L - M^{L-1}\sum_{j=0}^{L-1} k_j.$$

Setting $k_j = s2^{j+3} \ge j + 4s \cdot 2^j$ we obtain

$$N_1 \ge M^L - cM^{L-1},$$

provided $M \ge M_0$, completing the proof of the first part.

The final statement of the lemma follows from the fact that there are at least $M^L - cM^{L-1}$ many L-tuples satisfying (15). □

VI. Proof of Theorem 2

Let L be even, and suppose $C \subset \{0,1\}^n$ is an $<L$-list-decodable code of radius $\tau_L + \varepsilon$ (in fact, we consider a sequence of codes with $\varepsilon \to 0$, so more exactly we should say C_ε, but for readability we avoid the subscript ε everywhere below). We wish to prove that $|C| = O(\varepsilon^{-1})$. Let $\rho_L(C) = \min_{x \in C^L} \mathrm{rad}(x)$ with the minimum taken over all L-tuples x with distinct elements. Unlike $\rho_L(C)$, the L-radius of a code (denoted $\tilde{\rho}_L(C)$) is not a well-behaved quantity. Sadly, our assumptions on C do not imply that $\rho_L(C) \ge \tau_L$. For example, if $L = 2$ then the radius of $\{000, 100\}$ is $1/3 > 1/4 = \tau_2$ whereas $\mathrm{rad}(000, 100) = 1/6$. To get around this, we use the pigeonhole principle to find a subcode C' of size $|C'| \ge 2^{-8L}|C|$ consisting of codewords with the same prefix of length $8L$. Removing the common prefix yields a code of block length $n - 8L$ whose L-radius is at least

$$\frac{n}{n - 8L}(\tau_L + \varepsilon) \ge (1 + 8L/n)\tau_L + \varepsilon \ge \tau_L + \varepsilon + 2L/n.$$

By Proposition 4 we have $\mathrm{rad}(x) \ge \tau_L + \varepsilon$ for every L-tuple x of distinct codewords from this new code. With slight abuse of notation, we rename this new code C (and adjust the value of n accordingly).

Lemma 13: Let C' be any code with $\rho_L(C') \ge \tau_L$. If $\mathrm{rad}(C') \le \frac{1}{2} - \delta$ then $|C'| < s$ for some s depending on δ.

Proof: Identify L-element subsets of C' with ordered L-tuples by fixing some (arbitrary) ordering on L'. Elements of such L-tuples are distinct. For every $x \in (C')^L$, which is an L-tuple with distinct elements, there is $\omega \in \Omega'_L$ that solves (9). This gives a coloring of L-element subsets of C' into $|\Omega'_L|$ colors. From finiteness of Ω'_L and the hypergraph version of Ramsey's theorem [7, Th. 2], it follows that, if C' is large enough, then there is a monochromatic subset $C'' \subset C'$ of size exceeding $\frac{L^2}{\tau_L - \tau_{\mathcal{U}[L],p}}$, where $p \stackrel{\mathrm{def}}{=} \frac{1}{2} - \delta$.

Let $\omega \in \Omega'_l$ be the color of C'', i.e., $\mathrm{mrad}_\omega(x) \ge \tau_L$ for any L-tuple $x \in (C'')^L$ with distinct elements. Since $\mathrm{rad}(C'') \le \mathrm{rad}(C') \le p$, it follows from Corollary 10 and the bound $\tau_{\mathcal{U}[L],p} < \tau_L$ of Lemma 8 that $|C''| \le \frac{L^2}{\tau_L - \tau_{\mathcal{U}[L],p}}$. The contradiction shows that C' cannot be arbitrarily large. □

Let H be the set of all L-tuples $x \in C^L$ such that $\mathrm{mrad}_\omega(x) > \tau_L$ for some $\omega \ne \mathcal{U}[L]$.

Lemma 14: We have $|H| \le c_L|C|^{L-1}$ for some constant c_L that depends only on L.

Proof: Let $\varepsilon_0 = 2^{-L}\min\{\tau_L - \tau_{\omega,1/2} : \omega \in \Omega'_L,\ \omega \ne \mathcal{U}[L]\}$. Since Ω'_L is finite, part (b) of Lemma 8 implies that $\varepsilon_0 > 0$. So, let δ be as in Lemma 12 applied with $\varepsilon = \varepsilon_0$. Let s be the bound from Lemma 13. Note that by the choice of ε_0, the set H consists entirely of the L-tuples violating (16) and hence (15). By the choice of s, alternative (a) in Lemma 12 is impossible. Therefore, by the last statement of the latter Lemma, we have $|H| \le c(s)|C|^{L-1}$. □

Proof of Theorem 2: Call an L-tuple $x \in C^L$ *good* if all of its codewords are distinct, and $x \notin H$. For a randomly chosen L-tuple $x \in C^L$, the probability that $x^{(i)} = x^{(i')}$ for some $i \ne i'$ is less than $L^2/|C|$. By the preceding lemma, the probability that $x \in H$ is also $O(1/|C|)$. So a random x is good with probability $1 - O(1/|C|)$. Lemma 9 then implies that

$$\Pr[x \text{ is good}] \mathop{\mathbb{E}}_{\text{good } x} \mathrm{mrad}_{\mathcal{U}[L]}(x) \le \tau_{\mathcal{U}[L],1/2} = \tau_L.$$

On the other hand, for a good L-tuple we have $\mathrm{rad}(x) = \mathrm{mrad}_{\mathcal{U}[L]}(x)$ and thus the expectation is lower bounded by $\tau_L + \varepsilon$. In all, we conclude that $\frac{\varepsilon}{\tau_L + \varepsilon} = O(1/|C|)$, completing the proof. □

VII. Proof of Theorem 1

Proof of Theorem 1: Recall that $M = \binom{2m}{m}$. Consider an $2m$-by-M matrix with $\{0,1\}$ entries whose columns are all possible vectors consisting of exactly m ones. The $2m$ rows of the matrix are the codewords of a code $C \subset \{0,1\}^M$. We claim that $\mathrm{mrad}_{\mathcal{U}[L]}(x) \ge \tau_L + c_L/2m + O(m^{-2})$ for every L-tuple x of distinct codewords from C.

By symmetry, $\mathrm{mrad}_{\mathcal{U}[L]}(x)$ is independent of the actual choice of x. So, fix any x, and pick j at random from $[M]$. Let 0_j be the number of these codewords that have 0 in the j'th column. Similarly, let 1_j be the number of these codewords that have 1 in the j'th column. Let $X_j = \min(0_j, 1_j)/L$. Note that $\mathrm{mrad}_{\mathcal{U}[L]}(x) = \mathbb{E}\, X_j$ by (6).

Suppose $L = 2k+1$ is odd. Then

$$\begin{aligned}\mathbb{E}_j X_j &= \frac{1}{2k+1}\sum_{l\le k} l\cdot \Pr[\min(0_i,1_i)=l]\\ &= \binom{2m}{m}^{-1}\sum_{l\le k}\frac{2l}{2k+1}\binom{2k+1}{l}\binom{2m-2k-1}{m-l}\\ &= \binom{2m}{m}^{-1}\sum_{1\le l\le k} 2\binom{2k}{l-1}\binom{2m-2k-1}{m-l}\\ &= \sum_{1\le l\le k} 2\binom{2k}{l-1}\frac{\prod_{j=m-l+1}^{m} j\cdot\prod_{j=m-2k+l}^{m} j}{\prod_{j=2m-2k}^{2m} j}\end{aligned}$$

which, as $m\to\infty$, is

$$\begin{aligned}&= \sum_{1\le l\le k}\binom{2k}{l-1}2^{-2k}\Bigg[1+\frac{1}{2m}\binom{2k+1}{2}-\frac{1}{m}\binom{l}{2}\\ &\quad -\frac{1}{m}\binom{2k-l+1}{2}+O(m^{-2})\Bigg]\\ &= \tau_{2k+1}+2^{-2k-1}k\binom{2k}{k}/2m+O(m^{-2})\end{aligned}$$

In the last equality here we used the expression for τ_{2k+1}, the formula for the variance of the binomial random variable $B(2k,1/2)$, and the known expression for the expected distance of a balanced random walk of $2k$ steps from the origin.

Similar computations hold if $L=2k$. Denote by $\sum'$ the sum in which the last summand is halved. The expected value of X_j is

$$\begin{aligned}&\binom{2m}{m}^{-1}\frac{1}{2k}\sum_{l\le k}{}' 2l\binom{2k}{l}\binom{2m-2k}{m-l}\\ &= \binom{2m}{m}^{-1}\sum_{l\le k}{}' 2\binom{2k-1}{l-1}\binom{2m-2k}{m-l}\\ &= \sum_{l\le k}2\binom{2k-1}{l-1}\frac{\prod_{j=m-l+1}^{m} j\cdot\prod_{j=m-2k+l+1}^{m} j}{\prod_{j=2m-2k+1}^{2m} j}\\ &= \sum_{l\le k}\binom{2k-1}{l-1}2^{-2k+1}\Bigg[1+\frac{1}{2m}\binom{2k}{2}-\frac{1}{m}\binom{l}{2}\\ &\qquad -\binom{2k-l}{2}+O(m^{-2})\Bigg]\\ &= \tau_{2k}+2^{-2k}k\binom{2k-1}{k}/2m+O(m^{-2}).\end{aligned}$$

□

VIII. Proof of Theorem 3

We start with the proof of the upper bound, following the approach of Konyagin in [8]. Let C be <3-list-decodable code of vectors in $\{0,1\}^n$ of radius at most $\tau_3+\varepsilon = 1/4+\varepsilon$. By Proposition 7 this implies that among any 3 codewords in C there are two of distance at least $(1/2+2\varepsilon)n$. For each codeword $x=(x_1,x_2,\ldots,x_n)$ define a vector $v=(v_1,v_2,\ldots,v_n)$ in the Euclidean space R^n by $v_i=\frac{(-1)^{x_i}}{\sqrt{n}}$. Note that each such vector is of unit norm, and among any three vectors there are two whose inner product is at most -4ε. Let V be the set of all the vectors obtained from the words in C and put $|V|=m$. Our objective is to show that $m\le O(1/\varepsilon^{3/2})$. Let $H=(V,E)$ be the graph whose set of vertices is V in which two vertices u,v are connected iff their inner product is larger than -4ε. Fix a vertex $v\in V$ and let $W=N(v)$ be the set of all its neighbors in H. Note that the inner product between any two vertices in W is at most -4ε. Therefore, if $d=|W|$ is the degree of v in H and $\|v\|$ denotes the Euclidean 2-norm of a vector v, then

$$0\le\left\|\sum_{u\in W}u\right\|^2\le d-d(d-1)4\varepsilon \tag{18}$$

implying that $d\le\frac{1}{4\varepsilon}+1$ and also implying that

$$\begin{aligned}\left\|\sum_{u\in W}u\right\|^2 &\le d-d(d-1)4\varepsilon=\frac{1}{4\varepsilon}(4\varepsilon d)(1+4\varepsilon-4\varepsilon d)\\ &\le\frac{(1+4\varepsilon)^2}{16\varepsilon}.\end{aligned}$$

Therefore, by Cauchy–Schwarz, for every $v\in V$

$$\sum_{u\in N(v)}\langle v,u\rangle\le\left\|\sum_{u\in N(v)}u\right\|\le\frac{1+4\varepsilon}{4\sqrt{\varepsilon}}. \tag{19}$$

This gives the following (which can be slightly improved, but as this only changes the error term we prefer to present the simple version below):

$$\begin{aligned}0\le\left\|\sum_{v\in V}v\right\|^2 &= m+\sum_{v\in V}\sum_{u\in N(v)}\langle v,u\rangle+\sum_{u\ne v\in V,\ uv\notin E}\langle v,u\rangle\\ &\le m+m\frac{1+4\varepsilon}{4\sqrt{\varepsilon}}-m\left(m-\frac{1}{4\varepsilon}-2\right)4\varepsilon.\end{aligned}$$

By the last inequality

$$\left(m-\frac{1}{4\varepsilon}-2\right)4\varepsilon\le 1+\frac{1+4\varepsilon}{4\sqrt{\varepsilon}},$$

implying that

$$m\le\frac{1}{16\varepsilon^{3/2}}+O\left(\frac{1}{\varepsilon}\right). \tag{20}$$

This completes the proof of the upper bound.

We proceed with the proof of the lower bound by describing an appropriate construction. Let $G=(V,E)$ be a graph on m vertices, suppose it is a Cayley graph of an elementary abelian 2-group $\mathbb{Z}_2^r$, let A be its adjacency matrix, and let $d=\lambda_1\ge\lambda_2\ge\cdots\lambda_m=-\lambda$ be its eigenvalues, where d is the degree of regularity and $-\lambda$ is the smallest eigenvalue. Assume, further, that G is triangle-free. As G is a Cayley graph of an elementary abelian 2-group, it has an orthonormal basis of eigenvectors $v_1,v_2,\ldots,v_m$ in which each coordinate of each vector is in $\{-1/\sqrt{m},1/\sqrt{m}\}$. Define $B=(A+\lambda I)/\lambda$ where I is the m-by-m identity matrix. Then B is a positive semidefinite matrix, its diagonal is the all-1 vector, its eigenvalues are $\mu_i=(\lambda_i+\lambda)/\lambda$ and the corresponding eigenvectors are the vectors v_i. Let P be the m-by-m orthogonal matrix whose columns are the vectors v_i, and note that the first v_1 is the constant vector $1/\sqrt{m}$. Let D be the diagonal matrix whose diagonal entries are the eigenvalues μ_i and let $\sqrt{D}$ denote the

diagonal matrix whose entries are $\sqrt{\mu_i}$. Then $P^t B P = D$ and thus $B = (P\sqrt{D})(\sqrt{D}P^t)$.

The rows of the matrix $P\sqrt{D}$ are row-vectors $x_1, x_2, \ldots, x_m$ where $x_i = (x_{i1}, x_{i2}, \ldots, x_{im})$. Note that for each j, $x_{ij} \in \{-\sqrt{\mu_j/m}, \sqrt{\mu_j/m}\}$ for all i, and that x_{i1} is positive for all i. In addition $x_i x_j^t = B_{ij}$ for all i, j meaning that the ℓ_2-norm of each vector x_i is 1 and that among any three vectors x_i there is an orthogonal pair. Let y_i be the vector obtained from x_i by removing its first coordinate (the one which is $\sqrt{\mu_1/m} = \sqrt{(d+\lambda)/m\lambda}$). Then each y_i is a vector of ℓ_2-norm $1 - \mu_1/m$ and among any three of them there is a pair with inner product $-\mu_1/m$. We can normalize the vectors by dividing each entry by $\sqrt{1-\mu_1/m}$ to get m unit vectors $z_1, z_2, \ldots, z_m$, where any three of them contain a pair with inner product $-\delta$, where $\delta = \mu_1/(m-\mu_1)$. Moreover, for the vectors $z_i = (z_{ij})$, for each fixed j the absolute value of all z_{ij} is the same for all i, denote this common value by t_j. We can now use the vectors z_i to define functions mapping $[0,1]$ to $\{1,-1\}$ as follows. Split $[0,1]$ into disjoint intervals I_j of length t_j^2 and define f_i to be $\mathrm{sign}(z_{ij})$ on the interval I_j. It is clear that the ℓ_2-norm of each f_i is 1 and the inner product between f_i and f_j is exactly that between z_i and z_j. In particular, each three functions f_i contain a pair whose inner product is at most $-\delta$.

One can replace the functions by vectors of $1, -1$ with essentially the same property, using an obvious rational approximation to the lengths of the intervals.

The graph in [1] is a triangle-free Cayley graph of an elementary abelian 2-group with $d = \big(1/4 + o(1)\big)m^{2/3}$ and $\lambda = \big(9 + o(1)\big)m^{1/3}$. This gives us $\delta = (1/36 + o(1))m^{-2/3}$ and hence the number of vectors is $m = \Theta\big((1/\delta)^{3/2}\big)$. Setting $\delta = 4\varepsilon$ this gives a binary code with $m = \Theta\big((1/\varepsilon)^{3/2}\big)$ codewords of length n so that among any three codewords there are two such that the Hamming distance between them is at least $(1/2+2\varepsilon)n$. According to Proposition 7, this means that the code is $<$3-list-decodable with $\tau = \frac{1}{4} + \varepsilon$, and thus $\mathrm{maxcode}_\varepsilon(L=2) = \Omega(\varepsilon^{-\frac{3}{2}})$.

IX. Spherical Codes in the Hilbert Space

Let us now consider a similar question for the case of the real Hilbert space $\mathcal{H}$ (the space of square-summable sequences of real numbers). Similar to the binary alphabet, we may motivate the question by the desire to construct a maximal number M of unit-energy signals, such that when one of them is sent and adversarial noise of bounded energy is applied, it is still possible to reconstruct the original signal, to within a list of size $< L$. We also note that results on adversarial-noise lead to bounds for the average-noise variation, as propounded in [12, Sec. XII]. We proceed to formal definitions.

We shall employ the same notation as in the rest of the paper, but with the meaning adapted to spherical codes. For example, we denote the norm in $\mathcal{H}$ by $\|\cdot\|$. We redefine $\mathrm{rad}(x)$ similarly: for an arbitrary L-tuple $x = (x^{(1)}, \ldots, x^{(L)}) \in \mathcal{H}^L$ we define

$$\mathrm{rad}(x) = \min_{y \in \mathcal{H}} \max_j \|x^{(j)} - y\|,$$

$$\mathrm{diam}(x) = \max_{i,j} \|x^{(i)} - x^{(j)}\|.$$

Recall Jung's theorem [6, (2.6)]: For any L-tuple x we have

$$\mathrm{rad}(x) \le \sqrt{\frac{L-1}{2L}}\,\mathrm{diam}(x) \tag{21}$$

with equality if and only if $x^{(1)}, \ldots, x^{(L)}$ are the vertices of an $(L-1)$-simplex, i.e., when x consists of L vectors with pairwise distances all equal.

A *spherical code* C is a finite collection of unit-norm vectors in $\mathcal{H}$ and its L-radius $\rho_L(C)$ is the minimum value of $\mathrm{rad}(x)$ among all L-tuples x of distinct elements of C. We define

$$\mathrm{maxcode}_L(\varepsilon) \stackrel{\mathrm{def}}{=} \sup\big\{|C| : \rho_L(C) \ge \tau_L + \varepsilon\big\},$$

where in this section $\tau_L \stackrel{\mathrm{def}}{=} \sqrt{\frac{L-1}{L}}$. Our formulation corresponds to the problem of packing balls $B(x,r) \stackrel{\mathrm{def}}{=} \{y \in \mathcal{H} : \|x-y\| \le r\}$ centered on the unit sphere so that no point of $\mathcal{H}$ is covered by more than $(L-1)$ of them. Another equivalent way is to consider the problem of packing spherical caps $C(x,\alpha) = \{y : \|y\| = 1, \langle y, x\rangle \ge \cos\alpha\}$, where $\|x\| = 1$, with the requirement that no point *of the unit sphere* is covered by more than $(L-1)$ of them.

A classical result of Rankin [11] solves the case $L = 2$:

$$\mathrm{maxcode}_2(\varepsilon) = \left\lceil 1 + \frac{1}{2\sqrt{2}\varepsilon + 2\varepsilon^2} \right\rceil = \Theta\left(\frac{1}{\varepsilon}\right). \tag{22}$$

For $L > 2$, Blachman and Few [2] proved that if $\mathcal{H}$ is replaced by $\mathbb{R}^n$ then codes with $\rho_L(C) > \tau_L$ have size polynomial in n, while for $\rho_L(C) < \tau_L$ exponentially large codes exist. This was improved by Blinovsky [4], who demonstrated that codes with $\rho_L(C) > \tau_L + \varepsilon$, $\varepsilon > 0$ have a finite upper bound on their size independent of n. His proof relied on the Ramsey theorem and can be condensed as follows:

Proposition 15 [4]: For any $\varepsilon > 0$ $\mathrm{maxcode}_L(\varepsilon)$ is finite.

Proof: Consider a code C with $\rho_L(C) \ge \tau_L + \varepsilon$. Fix an integer $q \ge 1$, and break $[0,2]$ into q intervals of size $\frac{2}{q}$. Consider a code C and label each pair $(c,c') \in \binom{C}{2}$ according to the interval which contains $\|c - c'\|$. By Ramsey's theorem if C is sufficiently large then there should exist a large subcode C' whose all pairwise distances are in $[a, a+\frac{2}{q})$. From (22) we have $a \le \sqrt{2} + O(1/|C'|)$ and from (21) we get that $\rho_L(C) \le \rho_L(C') \le \tau_L + O(1/|C'|) + O(1/q)$ and hence $|C'| \le O(1/\varepsilon)$ when $q = O(1/\varepsilon)$. Consequently, C cannot be arbitrary large for a given $\varepsilon > 0$. □

Our main result on spherical codes is the following.

Theorem 16: For any $L \ge 2$ there exist constants $c_1, c_2 > 0$ such that for all $\varepsilon > 0$

$$c_1 \varepsilon^{-1} \le \mathrm{maxcode}_L(\varepsilon) \le c_2 \varepsilon^{-\frac{L^2-L+2}{2L}}. \tag{23}$$

Before proving the theorem we establish two auxiliary lemmas.

Definition: Call a collection S of unit vectors an (m,ϵ)-*system* if among any m distinct elements $x_1, \ldots, x_m \in S$ there exists a pair with $\langle x_i, x_j\rangle < -\epsilon$.

Lemma 17: For each m there exists $C_m > 0$, such that the size of any (m, ϵ)-system S is at most $C_m \epsilon^{-\frac{m}{2}}$ and

$$\| \sum_{x \in S} x \| \le C_m \epsilon^{-\frac{m-1}{2}}.$$

Proof: For $m = 2$ this follows from (22) and (18). For $m = 3$ this was shown above in (19) and (20), essentially by reducing to $m = 2$. In general, for arbitrary m we can define a graph with vertices S as in the proof of (20) and notice that the neighborhood $\mathcal{N}(v)$ is an $(m-1, \epsilon)$-system and then apply induction. □

Lemma 18: For any $L \ge 3$ there exists a non-negative function $\tau(\gamma) = \frac{2\gamma}{L^2-L-2} + O_L(\gamma^2)$, $\gamma \in [0, 1]$, with the following property. Consider any L-tuple $x = (x_1, \ldots, x_L)$ of unit-norm vectors with $\mathrm{rad}(x) \ge \tau_L$. If $\langle x_1, x_2 \rangle \ge \gamma \ge 0$ then there exist i, j such that $\langle x_i, x_j \rangle \le -\tau(\gamma)$.

Proof: Entirely like in (8) we can prove

$$\begin{aligned} \mathrm{rad}(x)^2 &= \max_{\omega \in \Omega_L} \min_{y \in \mathcal{H}} \mathop{\mathbb{E}}_{i \sim \omega} \|x_i - y\|^2 \\ &= \max_{\omega \in \Omega_L} \min_{y \in \mathcal{H}} \left(1 - 2\Big\langle \sum_i \omega_i x_i, y \Big\rangle + \|y\|^2 \right) \\ &= \max_{\omega} \left(1 - \Big\| \sum_i \omega_i x_i \Big\|^2 \right) = 1 - \min_{\omega \in \Omega_L} V(\omega), \end{aligned} \quad (24)$$

where $V(\omega) = \sum_{i,j} v_{i,j} \omega_i \omega_j$ is the quadratic form corresponding to the Gram matrix of x with $v_{i,j} = \langle x_i, x_j \rangle$.

Fix some $0 \le \tau \le \frac{1}{L-1}$ and suppose now that x is such that $\langle x_i, x_j \rangle \ge -\tau$ for all i, j. We will show that for some function $\tau(\gamma)$ if $\tau < \tau(\gamma)$ then $\mathrm{rad}(x) < \tau_L$. To that end, we introduce another quadratic form $U(\omega) = \sum_{i,j} u_{i,j} \omega_i \omega_j$ with

$$u_{i,j} = \begin{cases} 1, & i = j, \\ \gamma, & \{i, j\} = \{1, 2\}, \\ -\tau, & \text{otherwise.} \end{cases} \quad (25)$$

Note that according to assumptions $v_{i,j} \ge u_{i,j}$ and, therefore, on Ω_L we have $V(\omega) \ge U(\omega)$, and

$$\min_{\Omega_L} V(\omega) \ge \min_{\Omega_L} U(\omega).$$

We next show that U is non-negative definite for all $0 \le \tau \le \frac{1}{L-1}$ and all $-\frac{1}{L-1} \le \gamma \le 1$. From convexity of the PSD cone, it is sufficient to check the four corners. For $\tau = 0$ the statement is clear. For $\tau = \frac{1}{L-1}$ we consider the two endpoints: $\gamma = -\frac{1}{L-1}$, $\gamma = 1$. For $\gamma = -\frac{1}{L-1}$ the resulting quadratic form equals $U_1(\omega) = \sum_i \omega_i^2 - \frac{1}{L-1} \sum_{i \ne j} \omega_i \omega_j$ and corresponds to the Gram matrix of unit-norm vectors forming an $(L-1)$-simplex centered at the origin. Consequently, U_1 is positive definite. Similarly, for $\gamma = 1$, the quadratic form corresponds to Gram matrix of the following collection: take unit-norm vectors forming an $(L-1)$-simplex, delete one vector and add a copy of another. The resulting quadratic form is non-negative definite.

Since U is convex, we could evaluate $\min_\omega U(\omega)$ by arguing that optimal assignment is symmetric (has equal coordinates $3, \ldots, n$ and $1, 2$). Instead we prefer to proceed indirectly and show another useful property of radii in Hilbert space.

Since $U \succeq 0$, it is a Gram matrix of some other L-tuple x' of unit-norm vectors and we know

$$\mathrm{rad}(x') \ge \mathrm{rad}(x). \quad (26)$$

We temporarily forget about the special form of U, as defined in (25), and view it as a generic Gram matrix of *some* L-tuple x' of unit-norm vectors with the property that $|\langle x'_i, x'_j \rangle| \le \theta$ for $i \ne j$. We will prove

$$\mathrm{rad}(x')^2 = \tau_L^2 - \frac{1}{L^2} \sum_{i \ne j} \langle x'_i, x'_j \rangle + O(\theta^2), \quad (27)$$

where the $O(\cdot)$ term is uniform in x'. Note that the first two terms of the expression in (27) correspond to $\omega = \mathcal{U}[L]$ in (24). As $\theta \to 0$ the L-tuple x' becomes very close to L orthogonal vectors, and hence in (24) we expect that the optimal $\omega = \mathcal{U}[L] + O(\theta)$, cf. (28). Since we are operating near the minimum of the quadratic form, the $O(\theta)$ deviation of ω translates to $O(\theta^2)$ deviation for the value of U.

Proceeding to a formal proof of (27), first notice that if $\omega_1 = 0$ then as $\theta \to 0$ we must have $1 - U(\omega) \le \tau_{L-1} + o(1)$ (since we are considering only $L-1$ almost orthogonal vectors). But $1 - \min_\omega U(\omega)$ tends to $\tau_L > \tau_{L-1}$, implying that for all sufficiently small θ, the minimizer of $U(\omega)$ is in the interior of Ω_L. At the optimal point ω^* the gradient of U is proportional to a vector of all ones $\mathbf{1}$, from where we find

$$\omega^* = c(I_L + \Delta)^{-1} \mathbf{1}, \quad (28)$$

where $(I_L + \Delta)$ is the matrix of U, and the normalizing constant c is found from $\langle \omega^*, \mathbf{1} \rangle = 1$ yielding $c = \langle (I_L + \Delta)^{-1} \mathbf{1}, \mathbf{1} \rangle^{-1}$. Altogether, we get

$$\begin{aligned} U(\omega^*) &= \langle (I_L + \Delta)\omega^*, \omega^* \rangle = c \\ &= \langle I_L \mathbf{1}, \mathbf{1} \rangle + \langle \Delta \mathbf{1}, \mathbf{1} \rangle + O(\theta^2). \end{aligned}$$

Finally, since $\mathrm{rad}(x')^2 = 1 - U(\omega^*)$ we get (27).

To complete the proof of the Lemma, note that from (26), (27) and (25) we have

$$\mathrm{rad}(x)^2 \le \tau_L^2 - \frac{1}{L^2}\big(2\gamma - (L^2 - L - 2)\tau\big) + O(\gamma^2) + O(\tau^2).$$

Consequently, for appropriately defined $\tau(\gamma)$, if $\tau < \tau(\gamma)$ we should have $\mathrm{rad}(x) < \tau_L$. Furthermore, as $\gamma \to 0$ we have that $\tau(\gamma) = \frac{2\gamma}{L^2-L-2} + O_L(\gamma^2)$, as claimed. □

Proof of Theorem 16: Consider a regular $(M-1)$-simplex of unit vectors in $\mathcal{H}$. The pairwise distances are equal $\sqrt{\frac{2M}{M-1}}$ and thus from (21) we have that the radius of any L-tuple is at least $\tau_L \sqrt{\frac{M}{M-1}} = \tau_L + \Omega(1/M)$, proving the lower bound in (23).

We proceed to the upper bound. Fix a code C with $\rho_L(C) \ge \tau_L + \varepsilon$. The main idea is again essentially due to Konyagin: fixing one point $c \in C$ and considering its close neighbors, we notice that the radius constraint (cf. Lemma 18) introduces repulsion between these neighbors (that is they should be widely separated among themselves) and consequently, neighborhoods can not be too large.

We proceed with the argument. First, by (21) any L-tuple with $\mathrm{rad}(x) \ge \tau_L + \varepsilon$ also satisfies $\mathrm{diam}(x) \ge \sqrt{2} + \sqrt{\frac{2}{\tau_L}}\varepsilon$, and thus the code C is also an (L, ϵ')-system, with $\epsilon' = \frac{2}{\sqrt{\tau_L}}\varepsilon$.

Next, let $\epsilon_1 = \varepsilon^{\frac{L-1}{L}}$ and $\epsilon_2 = \tau(\epsilon_1)$, where $\tau(\cdot)$ is from Lemma 18. We consider two types of neighbors c of each point $c_i \in C$, depending on

$$-\epsilon' \le \langle c, c_i\rangle \le \epsilon_1, \text{ or } \langle c, c_i\rangle > \epsilon_1. \tag{29}$$

Let $\mathcal{N}'(c_i)$ and $\mathcal{N}''(c_i)$ be the two respective sets of neighbors. The rest of the points are "far away" from c_i and satisfy

$$\langle c, c_i\rangle < -\epsilon'. \tag{30}$$

First, notice that since C is an (L, ϵ')-system, we have that $\mathcal{N}'(c_i) \cup \mathcal{N}''(c_i)$ is an (m, ϵ')-system with $m \stackrel{\text{def}}{=} L-1$. Thus from Lemma 17

$$|\mathcal{N}'(c_i)| \le |\mathcal{N}'(c_i) \cup \mathcal{N}''(c_i)| \le C_m \epsilon'^{-\frac{m}{2}}. \tag{31}$$

Second, take any $m = (L-1)$ distinct points in $\mathcal{N}''(c_i)$. Adding c_i to this m-tuple and applying Lemma 18 to the resulting L-tuple, we conclude that $\mathcal{N}''(c_i)$ is an (m, ϵ_2)-system. Therefore, from Lemma 17 we have

$$\Big\| \sum_{c\in\mathcal{N}''(c_i)} c \Big\| \le C_m \epsilon_2^{-\frac{m-1}{2}}. \tag{32}$$

Consider

$$\begin{aligned}\Big\langle c_i, \sum_{c\in C} c\Big\rangle &= 1 + \Big\langle c_i, \sum_{c\in\mathcal{N}''(c_i)} c\Big\rangle + \Big\langle c_i, \sum_{c\in\mathcal{N}'(c_i)} c\Big\rangle \\ &\quad + \Big\langle c_i, \sum_{c\notin\mathcal{N}'\cup\mathcal{N}''\cup\{c_i\}} c\Big\rangle \qquad (33)\\ &\le 1 + C_m \epsilon_2^{-\frac{m-1}{2}} + C_m \epsilon_1 \epsilon'^{-\frac{m}{2}} \\ &\quad - \epsilon'(|C| - 1 - C_m \epsilon'^{-\frac{m}{2}}) \qquad (34)\end{aligned}$$

where the second term is estimated by Cauchy–Schwarz and (32), the third term is by the definition of $\mathcal{N}'(c_i)$ and (31), and the fourth term is the combination of (30) and the bound in (31).

Summing (34) over all $c_i \in C$ and using $\sum_{c_i, c\in C}\langle c_i, c\rangle \ge 0$ we get

$$\epsilon'|C| \le 1 + \epsilon' + C_m \epsilon_2^{-\frac{m-1}{2}} + C_m \epsilon_1 \epsilon'^{-\frac{m}{2}} + C_m \epsilon'^{1-\frac{m}{2}},$$

from where, recalling that $\epsilon_1 \asymp \epsilon_2 \asymp \varepsilon^{\frac{L-1}{L}}$ and $\epsilon' \asymp \varepsilon$ we get that the first two terms and the last are negligible compared to the third and fourth, which are comparable and $\asymp \varepsilon^{-\frac{(L-1)(L-2)}{2L}}$. Canceling ϵ' we get an upper bound in (23). □

X. Remarks and Open Problems

- The $L/2$ in Proposition 4 can be improved to $O(\sqrt{L})$ using a combination of the Beck–Fiala floating colors argument with Spencer's six deviations theorem. However, even with this improvement, we do not see a way to prove Theorem 2 with a good value of the implicit constant.
- For odd $L \ge 5$, the best upper bound we have is a tower of exponentials of height L. To that end, one colors L-tuples of codewords according to the measure ω for which $\mathrm{rad}(x) = \mathrm{mrad}_\omega(x)$, uses Ramsey's theorem to get a monochromatic set, and then proceeds similarly to the proof of Theorem 2.
- In the $<L$-list-decodable code in Theorem 1, the code length is exponential in ε^{-1}. One can restrict that code to a random subset of $O(\varepsilon^{-2}\log\varepsilon^{-1})$ coordinates, and obtain a code of asymptotically the same size $c_L\varepsilon^{-1} + O(1)$.
  For $L = 2$ and $L = 4$, the Levenshtein's code has length $O(\varepsilon^{-1})$ and size $c_L\varepsilon^{-1} + O(1)$. We do not know if one can make the code in Theorem 1 of linear size for general L.
- It should be possible to improve Theorem 16. We conjecture that for spherical codes, for all L we have $\mathrm{maxcode}_\varepsilon(L) = O(1/\varepsilon)$ with simplex being the optimal code.

Acknowledgment

The authors would like to thank Alan Frieze for providing the reference [14].

References

[1] N. Alon, "Explicit Ramsey graphs and orthonormal labelings," *Electron. J. Combinat.*, vol. 1, no. 1, p. 12, 1994.
[2] N. M. Blachman and L. Few, "Multiple packing of spherical caps," *Mathematika*, vol. 10, no. 1, pp. 84–88, 1963.
[3] V. M. Blinovsky, "Bounds for codes in decoding by a list of finite length," *Problemy Peredachi Inform.*, vol. 22, no. 1, pp. 11–25, 1986.
[4] V. Blinovsky, "Multiple packing of the Euclidean sphere," *IEEE Trans. Inf. Theory*, vol. 45, no. 4, pp. 1334–1337, May 1999.
[5] V. Blinovsky, "Generalization of Plotkin bound to the case of multiple packing," in *Proc. IEEE Int. Symp. Inf. Theory*, Jun. 2009, pp. 2048–2050.
[6] L. Danzer, B. Grünbaum, and V. Klee, "'Helly's theorem and its relatives,'" in *Convexity: Proc. Symp. Pure Math.*, vol. 7, 1963, pp. 101–180.
[7] R. L. Graham, B. L. Rothschild, and J. H. Spencer, *Ramsey Theory* (Wiley Interscience Series in Discrete Mathematics). New York, NY, USA: Wiley, 1980.
[8] S. V. Konyagin, "Systems of vectors in Euclidean space and an extremal problem for polynomials," *Math. Notes Acad. Sci. USSR*, vol. 29, no. 1, pp. 33–40, 1981. [Online]. Available: http://mi.mathnet.ru/mz10048, doi: 10.1007/BF01142512.
[9] V. I. Levenshtein, "Application of Hadamard matrices on coding problem," *Problems Cybern.*, vol. 5, pp. 123–136, 1961.
[10] Y. Polyanskiy, "Upper bound on list-decoding radius of binary codes," *IEEE Trans. Inf. Theory*, vol. 62, no. 3, pp. 1119–1128, Mar. 2016.
[11] R. A. Rankin, "The closest packing of spherical caps in n dimensions," *Glasgow Math. Assoc.*, vol. 2, no. 3, pp. 139–144, 1955.
[12] C. E. Shannon, "Probability of error for optimal codes in a Gaussian channel," *Bell Syst. Tech. J.*, vol. 38, no. 3, pp. 611–656, 1959.
[13] J. V. Neumann, "Uber ein okonomsiches gleichungssystem und eine verallgemeinering des browerschen fixpunktsatzes," *Erg. Math. Kolloquiums*, vol. 8, pp. 73–83, 1937.
[14] E. W. Weisstein. *Random Walk–1-Dimensional—From MathWorld-A Wolfram Web Resource*. Accessed: Nov. 21, 2016. [Online]. Available: http://mathworld.wolfram.com/RandomWalk1-Dimensional.html

Noga Alon is a Professor of Mathematics in Princeton University and a Baumritter Professor of Mathematics and Computer Science (on leave) in Tel Aviv University, Israel. He received his Ph. D. in Mathematics at the Hebrew University of Jerusalem in 1983.

His research interests are mainly in Combinatorics, Graph Theory and their applications in Theoretical Computer Science. His main contributions include the study of expander graphs and their applications, the investigation of derandomization techniques, the foundation of streaming algorithms, the development and applications of algebraic and probabilistic methods in Discrete Mathematics and the study of problems in Information Theory, Combinatorial Geometry and Combinatorial Number Theory.

He is an ACM Fellow and an AMS Fellow, a member of the Israel Academy of Sciences and Humanities and of the Academia Europaea, and received the Erdős Prize, the Feher Prize, the Polya Prize, the Bruno Memorial Award, the Landau Prize, the Gödel Prize, the Israel Prize, the EMET Prize and the Dijkstra Prize.

Boris Bukh studied computer science at the City College of San Francisco, leaving it in 2003 without a degree. Seeking life with less debugging, he then enrolled into UC Berkeley from which he graduated in 2005 with a bachelor's degree in mathematics. He then obtained a doctorate from Princeton in 2009. The next three years he spent in the University of Cambridge, leaving it with a few publications and a wife. He has been at Pittsburgh since 2012, working at Carnegie Mellon University, and growing two lovely daughters.

Prof. Bukh interests are in pure mathematics. He is easily seduced by combinatorial problems with analytic and algebraic flavours. His main contributions have been in the areas of additive combinatorics, and extremal combinatorics.

Prof. Bukh has received Sloan Fellowship (2015) and NSF Career award (2016).

Yury Polyanskiy (S'08–M'10–SM'14) is an Associate Professor of Electrical Engineering and Computer Science and a member of LIDS at MIT. Yury received M.S. degree in applied mathematics and physics from the Moscow Institute of Physics and Technology, Moscow, Russia in 2005 and Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ in 2010. Currently, his research focuses on basic questions in information theory, error-correcting codes, wireless communication and fault-tolerant and defect-tolerant circuits. Dr. Polyanskiy won the 2013 NSF CAREER award and 2011 IEEE Information Theory Society Paper Award.