

Lower Bounds for Approximate LDCs

Jop Briët* Zeev Dvir† Guangda Hu‡ Shubhangi Saraf§

Abstract

We study an approximate version of q -query LDCs (Locally Decodable Codes) over the real numbers and prove lower bounds on the encoding length of such codes. A q -query (α, δ) -approximate LDC is a set V of n points in \mathbb{R}^d so that, for each $i \in [d]$ there are $\Omega(\delta n)$ disjoint q -tuples $(\mathbf{u}_1, \dots, \mathbf{u}_q)$ in V so that $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_q)$ contains a unit vector whose i 'th coordinate is at least α . We prove exponential lower bounds of the form $n \geq 2^{\Omega(\alpha\delta\sqrt{d})}$ for the case $q = 2$ and, in some cases, stronger bounds (exponential in d).

1 Introduction

Error Correcting Codes (ECCs) have always played an important part in the development of theoretical computer science. In particular, many of the foundational results of computational complexity rely in some way or another on constructions and analysis of ECCs (e.g., hardness of approximation, hardness-randomness tradeoffs). The study of ECCs from the perspective of complexity theorists sometimes has different a focus than the traditional information theory viewpoint. One such difference is the study of special kinds of codes that are useful for theory (i.e., for proving theorems such as the PCP theorem) but were not studied previously.

One such example are Locally-Decodable-Codes (LDCs) which were formally defined in the seminal work of Katz and Trevisan [KT00] (but were implicit in several prior works [BK95, Lip90, BF90]). These are codes that allow the receiver of a (possibly corrupted) encoding $y = C(x) \in \{0, 1\}^n$ of a message $x \in \{0, 1\}^d$ to probabilistically decode w.h.p a single message bit x_i by reading only q positions in y (which might contain at most δn errors). We usually think of q as either a small constant or a very slow growing function of n and of δ as a constant.

The only case of LDCs which is mostly well understood is that of 2-query codes (it is easy to see that 1-query codes do not exist). The Hadamard code $C(x) = (\langle x, a \rangle)_{a \in \{0, 1\}^d}$ is a 2-query code with exponential encoding length. In [GKST06, KdW04] it was shown that this is tight, that is, we always have $n \geq 2^{\Omega(\delta d)}$ for 2-query codes. For $q > 2$ there are huge gaps between the known lower and upper bounds. The best known lower bound is $n = \tilde{\Omega}(d^{1+1/(\lceil q/2 \rceil - 1)})$ for $q > 4$ [Woo07]

*Courant Institute of Mathematical Sciences, New York University. Email: jop.briet@cims.nyu.edu. Supported by a Rubicon grant from the Netherlands Organisation for Scientific Research (NWO).

†Department of Computer Science and Department of Mathematics, Princeton University, Princeton NJ. Email: zeev.dvir@gmail.com. Research partially supported by NSF grants CCF-0832797, CCF-1217416 and by the Sloan fellowship.

‡Department of Computer Science Princeton University, Princeton NJ. Email: guangdah@cs.princeton.edu. Research partially supported by NSF grants CCF-0832797, CCF-1217416 and by the Sloan fellowship.

§Department of Computer Science and Department of Mathematics, Rutgers University. Email: shubhangi.saraf@gmail.com.

and $n = \Omega(d^2)$ for $k = 3, 4$ [KdW04, Woo12]. The best constructions for $q > 2$ are given by Matching-Vector codes, which were introduced by Yekhanin in [Yek08] and further developed in [Efr09, Pra07, KY09, IS10, CFL⁺10, DGY11, BET10]. These codes have block-length of roughly $n \leq \exp \exp((\log d)^{O(\log \log q / \log q)} (\log \log d))$.

One important sub-case of LDCs is that of *linear codes* (all known constructions are linear as far as we know). That is, the encoding is a linear mapping $C : \mathbb{F}^d \mapsto \mathbb{F}^n$ over some field \mathbb{F} . In this case, one can show that w.l.o.g. the decoding is linear as well. More formally, if we let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{F}^d$ be the rows of the generating matrix of C (so that $C(\mathbf{x})_i = \langle \mathbf{x}, \mathbf{v}_i \rangle$) then we have that, for each $i \in [d]$ there must exist a matching M_i of at least $\Omega(\delta n)$ disjoint pairs $\mathbf{v}_{j_1}, \mathbf{v}_{j_2}$ that span \mathbf{e}_i (the i 'th standard basis vector). To locally decode x_i one can simply pick a random pair in the matching M_i and calculate:

$$x_i = \langle \mathbf{x}, \mathbf{e}_i \rangle = a \langle \mathbf{x}, \mathbf{v}_{j_1} \rangle + b \langle \mathbf{x}, \mathbf{v}_{j_2} \rangle$$

for some field elements a, b satisfying $a\mathbf{v}_{j_1} + b\mathbf{v}_{j_2} = \mathbf{e}_i$. In [DS05] it was shown that the lower bound of [GKST06] for binary linear codes can be extended to linear codes over any field and so, we know that the Hadamard code cannot be beaten even if we allow for a large alphabet.

In this work we consider a new notion of linear LDCs in which the underlying field is the real numbers and the decoding is ‘approximate’. Building on the above characterization of linear codes, we will consider arrangements of points $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^d$ in which, for every $i \in [d]$ there are many disjoint pairs that ‘almost span’ \mathbf{e}_i in some concrete way (we give exact definitions below). Overall, our results are negative and show that, even if we allow a very loose notion of approximation, the encoding length is still exponential (either in \sqrt{d} or in d , depending on the model). We prove several theorems for various settings of the parameters, using a wide variety of techniques.

Motivation and related works: Our motivation for studying this problem comes from several directions. Firstly, one could hope to use approximate codes in practice (if these had sufficiently good parameters). As long as the approximation parameter is not too large we could hope to recover some approximation of x_i using the two queries to the code (assuming x_i is some quantity we are interested in and we don’t mind some small error). Another motivation comes from trying to understand 3-query codes. Here, even if we restrict our attention to real codes over \mathbb{R} , there is still an exponential gap between lower and upper bounds. In a recent work, [DSW13], a subset of the current authors and Avi Wigderson proved an $n > d^{2+\epsilon}$ lower bound (for some positive ϵ) for a closely related notion of 2-query Locally Correctable Codes (LCCs) over \mathbb{R} , improving upon the known quadratic bound. Originally, the proof of [DSW13] used a reduction from (exact) 3-LCCs over \mathbb{R} to 2-query approximate LDCs (later, a different proof was found). This raises the possibility that, in the future, perhaps approximate codes will find more applications. We are also motivated by connections to well studied questions in combinatorial geometry. In [BDWY12, DSW12] it was shown that proving lower bounds on LCCs is closely related to questions in the spirit of the Sylvester-Gallai theorem. Here, one tries to take local information about a point configuration (say, many collinear triples) and convert this information to a global bound on the dimension spanned by the points. We can view some of the theorems in this work in this spirit. Approximate versions of Sylvester-Gallai type theorems and LCCs were recently explored in [ADSW12].

1.1 Definitions and results

We begin with some notations. A q -matching M in $[n]$ is defined to be a set of disjoint unordered q -tuples (i.e. disjoint subsets of size q) of $[n]$. We denote by \mathbf{e}_i the i 'th standard basis vector in \mathbb{R}^d . The standard inner product of two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ is given by $\langle \mathbf{x}, \mathbf{y} \rangle$ and the ℓ_2 norm of $\mathbf{x} \in \mathbb{R}^d$ is $\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$.

Definition 1.1 (weight_i). For a vector $\mathbf{u} \in \mathbb{R}^d$ we define $\text{weight}_i(\mathbf{u}) = |\langle \mathbf{u}, \mathbf{e}_i \rangle| / \|\mathbf{u}\|_2$ (i.e., the absolute value of the i 'th coordinate of the normalized vector $\mathbf{u} / \|\mathbf{u}\|_2$).

Clearly we have $\sum_{i \in [d]} \text{weight}_i(\mathbf{u})^2 = 1$. We now state our definition of approximate LDC.

Definition 1.2 (Approximate LDC). Let d, n, q be positive integers and $\alpha, \delta \in [0, 1]$ real numbers. A q -query (α, δ) -approximate LDC is a pair (V, M) with

1. $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ a multiset of vectors in \mathbb{R}^d . The parameter n is the size (or block length) of the code and the parameter d is the dimension (or message length) of the code.
2. $M = (M_1, \dots, M_d)$ with each M_i being a q -matching in $[n]$ so that, if $\{j_1, \dots, j_q\} \in M_i$, then there exists $\mathbf{u} \in \text{span}\{\mathbf{v}_{j_1}, \dots, \mathbf{v}_{j_q}\}$ with $\text{weight}_i(\mathbf{u}) \geq \alpha$.

The sizes of the matchings M_i must satisfy $|M_1| + |M_2| + \dots + |M_d| \geq \delta dn$ and the parameter δ is called the density of the code¹.

Our first theorem gives an exponential bound on the block length of approximate 2-LDCs for any $\alpha > 0$. Notice that the bound gets worse as α approaches $1/\sqrt{d}$, at which point we cannot expect any lower bound to hold (since a single vector \mathbf{u} can have $\text{weight}_i(\mathbf{u}) \geq 1/\sqrt{d}$ for all $i \in [d]$).

Theorem 1.3. [General lower bound] A 2-query (α, δ) -approximate LDC of size n and dimension d must satisfy $n \geq 2^{\Omega(\alpha\delta\sqrt{d})}$.

We could hope to replace the exponential dependence on \sqrt{d} with an exponential dependence on d (as is the case with exact 2-LDCs). In fact, we conjecture that a general bound of the form $n \geq \exp(\delta\alpha^2 d)$ should hold (the quadratic dependence on α is necessary to avoid hitting the $\alpha = 1/\sqrt{d}$ barrier). Currently, we are only able to prove this conjecture when α is sufficiently close to 1. This is stated in the next theorem.

Theorem 1.4. [Lower bound for large α] Let $\alpha_0 = \sqrt{1 - 1/(4\pi^2)} \approx 0.987$. A 2-query (α, δ) -approximate LDC of size n , dimension d and $\alpha > \alpha_0$ must satisfy $n \geq 2^{\Omega(\delta d)}$, where the hidden constant in the $\Omega(\cdot)$ depends on $\alpha - \alpha_0$.

There is another special case where we can get an exponential dependence on d instead of \sqrt{d} . It is a natural restriction of the general definition but it requires two new notions (that will be useful in their own right down the road). The first is that of a *simple* code (we will only care about 2-query codes).

Definition 1.5 (Simple code). Let (V, M) be a 2-query (α, δ) -approximate LDC. We say that (V, M) is a simple code if, for every $i \in [d]$ and $\{j_1, j_2\} \in M_i$ we have $\text{weight}_i(\mathbf{v}_{j_2} - \mathbf{v}_{j_1}) \geq \alpha$.

¹The traditional definition would ask for each M_i to be of size at least δn but our definition is more general, which makes our (negative) results stronger.

In other words, a simple code is an arrangements of points in \mathbb{R}^d so that, for any $i \in [d]$ there are $\approx \delta n$ (on average) disjoint pairs of points that ‘point’ in a direction that has projection at least α on the i ’th axis. An example of such an arrangement is the boolean cube $\{0, 1\}^d \subset \mathbb{R}^d$ (all zero/one vectors), where M_i consists of all $n/2$ pairs that differ only in the i ’th entry (so $\alpha = 1$).

Another feature of the hypercube is that all the distances between pairs in M_1, \dots, M_d are equal (they all equal one), motivating the following definition.

Definition 1.6 (*c*-bounded). *Let $c \geq 1$ and let (V, M) be a 2-query (α, δ) -approximate LDC. We say that (V, M) is c -bounded if, for every $i \in [d]$ and $\{j_1, j_2\} \in M_i$ we have $\|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\| \in [1, c]$.*

The fact that the hypercube is both c -bounded (with $c = 1$) and simple motivates the study of structures that satisfy these two conditions. In particular, we ask whether there exists a point arrangement in \mathbb{R}^d which is ‘roughly’ like the hypercube but has far fewer than 2^d points. Here, the notion of ‘roughly’ is captured by allowing pairwise distance to be ‘close’ to 1 and the differences between adjacent vertices to be only somewhat axis parallel. The following theorem shows that such configurations do not exist (that is, you cannot beat the hypercube by much).

Theorem 1.7. *[Lower bound for simple c -bounded LDCs] A 2-query c -bounded simple (α, δ) -approximate LDC of size n and dimension d must satisfy $n \geq 2^{\Omega(\alpha^2 \delta^2 d / (\log c)^2)}$.*

Finally, we consider arbitrary q -query approximate codes and observe that the lower bound proof of [KT00] can be made to work also for approximate LDCs (with some additional work). This gives the following theorem.

Theorem 1.8. *Let $q \geq 1$ be an integer constant. A q -query (α, δ) -approximate LDC of size n and dimension d must satisfy $n \geq \Omega((\alpha^2 \delta^{1/q} d)^{\frac{q}{q-1}})$.*

1.2 Techniques

We briefly outline the techniques that appear in the proofs of our theorems.

Simple codes: An important ingredient in the proofs of Theorem 1.3 and 1.4 is a general reduction from any approximate 2-LDC to a simple code. The reduction follows by first normalizing the lengths of all vectors and then observing that if some linear combination $a\mathbf{v}_{j_1} + b\mathbf{v}_{j_2}$ has large weight_i , then either one of the vectors $\mathbf{v}_{j_1}, \mathbf{v}_{j_2}$ has large weight_i or the coefficients a, b are close to $1, -1$. We can thus throw away all pairs in the matching M_i in which one of the vectors has large weight_i and get a simple code (we do not throw away too many pairs since each vector has only a few large coordinates). Since this reduction does not preserve c -boundedness, we can unfortunately not use it to argue that Theorem 1.7 works for non-simple c -bounded codes.

Proof of general bound: The proof of Theorem 1.3 (for simple codes w.l.o.g.) is via a recursive partitioning argument. In each step we pick a random $i \in [d]$ and partition V into two sets using a random shift of a hyperplane orthogonal to \mathbf{e}_i . We analyze the expected number of edges (pairs in some M_i) cut in this process and show that is bounded by $O(\sqrt{d}/\alpha) \cdot \min\{|S|, |\bar{S}|\}$ with S, \bar{S} representing the two parts of the cut. The same inequality holds also when partitioning any subset $V' \subset V$ and so we can proceed recursively and obtain a bound of $O((\sqrt{d}/\alpha)n \log_2 n)$ on the total number of edges. Since this number is at least δdn the theorem follows. This proof is inspired by the one appearing in [GKST06] for exact (simple) 2-LDCs.

Proof of bound for large α : Here we rely on a recent work of [KROW12] which gives a (randomized) tiling of \mathbb{R}^d with cells that have volume 1 and surface area $O(\sqrt{d})$ (same as a sphere up to a constant). This result gives a randomized rounding algorithm that we can leverage towards ‘rounding’ our approximate code to an exact code (very roughly speaking) when α is large. This step is then combined with a random partitioning argument as in the proof of Theorem 1.3.

Proof for simple c -bounded codes: For this setting we use the LDC to construct a function F from \mathbb{R}^d to the space of complex $n \times n$ matrices given by $F(\mathbf{x}) = (e^{-i(\mathbf{x}, \mathbf{v}_s - \mathbf{v}_t)})_{s,t=1}^n$. The crux of the proof applies an inequality relating the trace norms of the first level (matrix) Fourier-coefficients of a matrix-valued function to its average trace norm (see Lemma 5.3). The crucial observation is that the norms of the first level Fourier coefficients of the above defined F can be lower bounded using the LDC property. The result then follows by combining this with the trivial upper bound on the average norm of F . This proof loosely follows an argument of [BARdW08] used for binary (non linear) LDCs and is inspired by work of [BNR12] linking LDCs to geometry of Banach spaces.

Organization: We describe our reduction from general to simple 2-query codes in Section 2. In Section 3 we prove the bound for general codes (Theorem 1.3). In Section 4 we prove the bound for α close to 1 (Theorem 1.4). In Section 5 we prove the bound for c -bounded codes (Theorem 1.7). Finally, in Section 6 we prove the bound for general q -query approximate codes (Theorem 1.8).

Acknowledgments: The authors would like to thank Avi Wigderson for many helpful conversations.

2 Simple codes

In this section we prove the following theorem showing that any 2-query approximate LDC can be transformed into a simple code with similar parameters.

Theorem 2.1. *If there exists a 2-query (α, δ) -approximate LDC of size n dimension d , then, for any integer $k > 1/\alpha^2$, there exists a simple 2-query (α', δ') -approximate LDC of size n' and dimension d , where $\alpha' \geq \sqrt{\alpha^2 - 1/k}$, $\delta' \geq \delta - k/d$ and $n' \leq 2n$.*

The main idea behind the proof of this result is as follows. Suppose that we have a pair of unit vectors $\mathbf{u}, \mathbf{w} \in \mathbb{R}^d$ with $\text{weight}_i(a\mathbf{u} + b\mathbf{w}) \geq \alpha$. It will be convenient to think of α as being close to one (the proof will work for any α). So, after normalizing the coefficients a, b we have that the unit vector $v = a\mathbf{u} + b\mathbf{w}$ is close to \mathbf{e}_i . We separate into two cases. In the first case, both \mathbf{u} and \mathbf{w} are almost orthogonal to \mathbf{e}_i . In this case, we must have that $\mathbf{u} - \mathbf{w}$ ‘points’ in the direction of \mathbf{e}_i (see diagrams in the complete proof) and so we don’t really need the coefficients a, b . In the other case, at least one of \mathbf{u}, \mathbf{w} have significant inner product with \mathbf{e}_i . Notice, however, that, for each fixed \mathbf{u} , this can only happen with a small number of \mathbf{e}_i ’s when $i \in [n]$. These ‘bad’ pairs can be removed from the matchings without causing a big decrease in their average size.

It will be convenient to use the following corollary of Theorem 2.1 in which we set $k = \lceil 2/\alpha^2 \rceil$.

Corollary 2.2. *Suppose $d \geq 6/\alpha^2\delta$. If there exists a 2-query (α, δ) -approximate LDC of size n dimension d , then there exists a simple 2-query approximate (α', δ') -LDC of size n' and dimension d , where $\alpha' \geq \alpha/\sqrt{2}$, $\delta' \geq \delta/2$ and $n' \leq 2n$.*

We now move on to the formal proof of Theorem 2.1.

Proof of Theorem 2.1: We first modify the code in the following way. Let $k > \frac{1}{\alpha^2}$ be a fixed integer.

1. For every matching M_i , remove the pairs $\{j_1, j_2\}$ in which the i 'th entry of \mathbf{v}_{j_1} is among its $k - 1$ largest (in absolute value) entries or if this is the case for \mathbf{v}_{j_2} . Since every $\mathbf{v} \in V$ causes at most k pairs (from all d matchings) that contain it to be removed, there are at most kn pairs removed in this step altogether.
2. Normalize all vectors in V so that $\|\mathbf{v}_j\|_2 = 1$ for all $j \in [n]$ and discard all zero vectors.
3. For every $\mathbf{v}_j \in V$, add $-\mathbf{v}_j$ to V . For every original tuple $\{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$, we replace it with two tuples: either $\{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$, $\{-\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$ or $\{\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$, $\{-\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ (to be determined later).

Let $(V', M' = (M'_1, \dots, M'_d))$ be the vectors and matchings obtained from the above procedure.

Claim 2.3. *After the first step, if $\{j_1, j_2\} \in M_i$ is not deleted, then*

$$|v_{j_1 i}| \leq \frac{1}{\sqrt{k}} \|\mathbf{v}_{j_1}\|_2 \quad \text{and} \quad |v_{j_2 i}| \leq \frac{1}{\sqrt{k}} \|\mathbf{v}_{j_2}\|_2.$$

Proof: We only consider \mathbf{v}_{j_1} . According to the first step, the i 'th coordinate must not be among the maximum $k - 1$ ones. If this coordinate has absolute value greater than $\|\mathbf{v}_{j_1}\|_2/\sqrt{k}$, then there are at least k coordinates greater than $\|\mathbf{v}_{j_1}\|_2/\sqrt{k}$, which is impossible. \square

Claim 2.4. *After the first step, for any remaining pair $\{j_1, j_2\} \in M_i$, \mathbf{v}_{j_1} and \mathbf{v}_{j_2} are linearly independent. This implies that no remaining pair contains $\mathbf{0}$ and so discarding all zero vectors in step 2 above does not remove any additional pairs from the matchings.*

Proof: Assume the contrary and $\mathbf{v}_{j_1} \neq \mathbf{0}$. Then $\text{span}\{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ contains only multiples of \mathbf{v}_{j_1} . Thus, the i 'th coordinate of \mathbf{v}_{j_1} has magnitude at least $\alpha \|\mathbf{v}_{j_1}\|_2 > \|\mathbf{v}_{j_1}\|_2/\sqrt{k}$, violating Claim 2.3. \square


For the new code (V', M') the dimension is $d' = d$, the size is $n' \leq 2n$, and the number of tuples is at least $2(\delta dn - kn)$, which implies the density $\delta' \geq \delta - k/d$. Notice that we might have removed some of the matchings completely. We still, however, consider the dimension as $d' = d$ (since we only use the sum of sizes of all matchings).

Next we argue that the pairs remaining after step 1 above give a simple code (up to changing signs) and calculate the parameter α' . Fix a pair $\{j_1, j_2\} \in M_i$ that remains after the first step. We will show that either $\mathbf{v}_{j_1} - \mathbf{v}_{j_2}$ or $\mathbf{v}_{j_1} - (-\mathbf{v}_{j_2})$ has a large i 'th coordinate. Precisely, we show that either

$$\text{weight}_i(\mathbf{v}_{j_1} - \mathbf{v}_{j_2}) \geq \sqrt{\alpha^2 - \frac{1}{k}} \quad \text{or} \quad \text{weight}_i(\mathbf{v}_{j_1} - (-\mathbf{v}_{j_2})) \geq \sqrt{\alpha^2 - \frac{1}{k}}.$$

Then if the first one holds, we choose to use $\{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ and $\{-\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$ at the third step; otherwise if the second one holds, we select $\{\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$ and $\{-\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$. Thus the code is reduced to a simple code with $\alpha' \geq \sqrt{\alpha^2 - 1/k}$.

We consider the plane determined by \mathbf{v}_{j_1} and \mathbf{v}_{j_2} , and set up Cartesian axes at the origin. Let the projection of \mathbf{e}_i on this plane be in the direction of the y axis and choose one of two possible directions for the x axis arbitrarily. Let τ be the angle between the plane and \mathbf{e}_i . This setting is shown in figure (a).



Since $\{j_1, j_2\} \in M_i$, we see that

$$\cos \tau \geq \alpha. \quad (1)$$

Let $\theta_1 \in [0, 2\pi)$ be the angle between \mathbf{v}_{j_1} and the x axis, and $\theta_2 \in [0, 2\pi)$ be the angle between \mathbf{v}_{j_2} and the x axis. This is shown as in figure (b). \mathbf{v}_{j_1} and \mathbf{v}_{j_2} correspond to points $(\cos \theta_1, \sin \theta_1)$ and $(\cos \theta_2, \sin \theta_2)$ on the plane.

By Claim 2.3, the i 'th coordinates of \mathbf{v}_{j_1} and \mathbf{v}_{j_2} are at most $1/\sqrt{k}$. Therefore

$$|\sin \theta_1| \cos \tau \leq \frac{1}{\sqrt{k}} \quad \text{and} \quad |\sin \theta_2| \cos \tau \leq \frac{1}{\sqrt{k}}.$$

We can see that the angles of $\pm \mathbf{v}_{j_1}$ and $\pm \mathbf{v}_{j_2}$, which are $\pm \theta_1$ and $\pm \theta_2$, must fall into two regions

$$\left[-\arcsin \frac{1}{\sqrt{k} \cos \tau}, \arcsin \frac{1}{\sqrt{k} \cos \tau} \right] \quad \text{and} \quad \left[\pi - \arcsin \frac{1}{\sqrt{k} \cos \tau}, \pi + \arcsin \frac{1}{\sqrt{k} \cos \tau} \right].$$

These are shown as two gray circular sectors (the left one and the right one) in figure (c).

We pair \mathbf{v}_{j_1} to the one of $\pm \mathbf{v}_{j_2}$ that lies in the same sector with \mathbf{v}_{j_1} , and do the same for $-\mathbf{v}_{j_1}$. It is easy to see our pairing is either $\{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$, $\{-\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$ or $\{\mathbf{v}_{j_1}, -\mathbf{v}_{j_2}\}$, $\{-\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$. We now argue that subtraction of vectors in a one of the new pairs (which belong to the same sector) must have a large i 'th coordinate. Without loss of generality, we assume \mathbf{v}_{j_1} and \mathbf{v}_{j_2} are paired and only consider this pair.

The vector $\mathbf{v}_{j_2} - \mathbf{v}_{j_1}$ is parallel to the tangent line to the unit circle at angle $(\theta_1 + \theta_2)/2$. One can verify that a unit vector parallel to $\mathbf{v}_{j_2} - \mathbf{v}_{j_1}$ must have y coordinate $\pm \cos(\theta_1 + \theta_2)/2$. Therefore

$$\frac{|\langle \mathbf{v}_{j_1} - \mathbf{v}_{j_2}, \mathbf{e}_i \rangle|}{\|\mathbf{v}_{j_1} - \mathbf{v}_{j_2}\|_2} = \cos \frac{\theta_1 + \theta_2}{2} \cdot \cos \tau = \sqrt{1 - \sin^2 \frac{\theta_1 + \theta_2}{2}} \cdot \cos \tau.$$

Since \mathbf{v}_{j_1} and \mathbf{v}_{j_2} are in the same circular sector, $|\sin(\theta_1 + \theta_2)/2| \leq 1/(\sqrt{k} \cos \tau)$. It follows that

$$\frac{|\langle \mathbf{v}_{j_1} - \mathbf{v}_{j_2}, \mathbf{e}_i \rangle|}{\|\mathbf{v}_{j_1} - \mathbf{v}_{j_2}\|_2} \geq \sqrt{\cos^2 \tau - \frac{1}{k}} \geq \sqrt{\alpha^2 - \frac{1}{k}}.$$

Here we used $\cos \tau \geq \alpha$ (Inequality (1)). This completes the proof. \square

3 Lower Bound for General Simple Codes

We associate with a simple code $C = (V, M)$ a labeled graph G_C on vertex set V with edges given by all pairs in M_1, \dots, M_d . We label each edge in M_i with the label i and allow for parallel edges (with different labels). We refer to the label of an edge e as the *direction* of the edge and denote it by $\text{dir}(e) \in [d]$. The proof will follow by analyzing cuts in the graph G_C , which we assume contains at least δdn edges.

For $S \subseteq V$, let $\text{Edge}(S)$ be the set of edges of G_C with both end points in S . We say that (S_1, S_2) is a *cut* if $S_1 \cup S_2 = S$ and $S_1 \cap S_2 = \emptyset$. The cut is *non-trivial* if $S_1, S_2 \neq \emptyset$. We use $\text{Edge}(S_1, S_2)$ to denote the set of edges with one endpoint in S_1 and the other in S_2 .

The next lemma of [GKST06, Appendix] relates the sizes of cuts in the graph with the total number of edges (the lemma holds for any graph). We include its proof for completeness.

Lemma 3.1. *Suppose that for every $S \subseteq V$ with $|S| \geq 2$, there exists a non-trivial cut (S_1, S_2) satisfying $|\text{Edge}(S_1, S_2)| \leq c \cdot \min\{|S_1|, |S_2|\}$, then G_C has at most $\frac{c}{2}|V| \log_2 |V|$ edges.*

Proof: The given condition is equivalent to

$$|\text{Edge}(S)| \leq c \cdot \min\{|S_1|, |S_2|\} + |\text{Edge}(S_1)| + |\text{Edge}(S_2)|. \quad (2)$$

We induct on $|S|$ to show $|\text{Edge}(S)| \leq \frac{c}{2}|S| \log_2 |S|$ for every non-empty $S \subseteq V$. For $|S| = 1$, this is trivial. Assume this is true for $|S| < k$, $k \geq 2$, and consider a subset S of size $|S| = k$. Let $S_1 \subseteq S$ be a proper and nonempty subset of S and let $S_2 = S \setminus S_1$ be its complement in S . Then, by the above condition 2 and the induction hypothesis,

$$|\text{Edge}(S)| \leq c \cdot \min\{|S_1|, |S_2|\} + \frac{c}{2}|S_1| \log_2 |S_1| + \frac{c}{2}|S_2| \log_2 |S_2|$$

Assume (w.l.o.g) that $\min\{|S_1|, |S_2|\} = |S_1|$ and let $\eta = |S_1|/k$. Notice that η belongs to $[0, 1/2]$. Then the above right-hand side can be re-written as $(ck/2)(2\eta - H_2(\eta)) + (ck/2) \log k$, where $H_2(\tau) = -\tau \log_2 \tau - (1 - \tau) \log_2 \tau$ is the binary entropy function. The function H_2 is concave on $[0, 1/2]$ and satisfies $H_2(0) = 0$ and $H_2(1/2) = 1$. Hence, the term $2\eta - H_2(\eta)$ is non-positive and we get the result $|\text{Edge}(S)| \leq (ck/2) \log k$, as claimed. \square

We now proceed to prove Theorem 1.3. We will show $n = 2^{\alpha\delta\sqrt{d}}$ for any (α, δ) simple code (the general case will follow using Corollary 2.2). This will follow by combining the following lemma and Lemma 3.1.

Lemma 3.2. *Let $C = (V, M)$ be an (α, δ) simple code and let G_C be the associated graph described above. Then, for any $S \subseteq V$ with $|S| \geq 2$, there exists a non-trivial cut (S_1, S_2) such that*

$$|\text{Edge}(S_1, S_2)| \leq \frac{2\sqrt{d}}{\alpha} \cdot \min\{|S_1|, |S_2|\}.$$

Proof: If S contains no edges, an arbitrary cut will satisfy the requirement. We thus assume that S contains at least one edge. We now analyze the size of a random cut chosen in a specific way.

Assume all points in V are in a $(d$ -dimensional) box of edge length L . We pick a random direction $i \in [d]$ and then pick a plane perpendicular to e_i at a random position intersecting the box. The plane cuts the box into two parts. We define S_1 to be the set of points in one part and S_2

to be the set of points in the other part (the probability of having a point on the hyperplane is zero). We analyze the edges in this cut (S_1, S_2) . We say that an edge $e \in \text{Edge}(S_1, S_2)$ is cut in the *right* direction if the plane is perpendicular to the direction of e , i.e. $\text{dir}(e) = i$.

We consider a specific edge. Let $e_0 = \{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ with $\{j_1, j_2\} \in M_{i_0}$ be an edge in direction $\text{dir}(e_0) = i_0$ and denote $\mathbf{v}_{j_1} - \mathbf{v}_{j_2} = (u_1, u_2, \dots, u_d)$.

For each $i' \in [d]$ the probability that e_0 is cut by a plane perpendicular to $\mathbf{e}_{i'}$ is

$$\Pr[i = i'] \cdot \Pr[\text{the plane falls between } v_{j_1 i'} \text{ and } v_{j_2 i'}] = \frac{1}{d} \cdot \frac{|u_{i'}|}{L}.$$

Therefore,

$$\Pr[e_0 \in \text{Edge}(S_1, S_2)] = \sum_{i'=1}^d \frac{1}{d} \cdot \frac{|u_{i'}|}{L} = \frac{|u_1| + |u_2| + \dots + |u_d|}{dL}.$$

Moreover, by the definition of an approximate code ($|u_{i_0}| \geq \alpha \|\mathbf{v}_{j_1} - \mathbf{v}_{j_2}\|_2$) and the Cauchy-Schwarz inequality, edge e_0 is cut in the right direction with probability

$$\begin{aligned} \frac{|u_{i_0}|}{dL} &\geq \frac{1}{dL} \cdot \alpha \sqrt{u_1^2 + u_2^2 + \dots + u_d^2} \\ &\geq \frac{1}{dL} \cdot \frac{\alpha}{\sqrt{d}} (|u_1| + |u_2| + \dots + |u_d|) \\ &= \frac{\alpha}{\sqrt{d}} \Pr[e_0 \in \text{Edge}(S_1, S_2)]. \end{aligned}$$

Since $\mathbf{v}_{j_1} - \mathbf{v}_{j_2}$ has at least one non-zero coordinate, $\Pr[e_0 \in \text{Edge}(S_1, S_2)]$ must be strictly positive. It follows that edge e_0 is cut in the right direction with probability strictly greater than

$$\frac{\alpha}{2\sqrt{d}} \Pr[e_0 \in \text{Edge}(S_1, S_2)].$$

Hence, the expected number of edges that are cut in the right direction is strictly greater than $\alpha \mathbb{E}[|\text{Edge}(S_1, S_2)|] / (2\sqrt{d})$. There must therefore exist an $i \in [d]$, a plane perpendicular to \mathbf{e}_i and a corresponding cut (S_1, S_2) which cuts strictly more than $\alpha |\text{Edge}(S_1, S_2)| / (2\sqrt{d})$ in the right direction. Since this number is non-negative, there must be at least one edge cut in the right direction. This implies that S_1 and S_2 are not empty.

All edges cut in the right direction must have the same direction i . Hence, these edges are disjoint (they form a matching in V), implying that the total number of cut edges is at most $\min\{|S_1|, |S_2|\}$. It follows immediately that

$$|\text{Edge}(S_1, S_2)| \leq \frac{2\sqrt{d}}{\alpha} \min\{|S_1|, |S_2|\}.$$

Therefore the cut (S_1, S_2) satisfies the requirement. \square

Now using Lemma 3.1 we can conclude that $\delta dn \leq (2\sqrt{d}/\alpha)n \cdot \log_2 n$, which gives $n \geq 2^{\alpha\delta\sqrt{d}}$ as required. This completes the proof of Theorem 1.3.

4 Lower Bound for Simple Codes with Large α

In this section we prove Theorem 1.7. By Theorem 2.1 it is enough to consider simple codes (the general case will follow by applying Theorem 2.1 with k a sufficiently large constant). We will use the definition and terminology of the graph G_C defined in the last section for simple codes. Hence, we think of pairs in M_i as edges in ‘direction’ $\text{dir}(e) = i$. We define the *length* of an edge $e = \{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ to be $\|\mathbf{v}_{j_1} - \mathbf{v}_{j_2}\|_2$.

We will use a recent result of [KROW12] concerning a partitioning (or tiling) of \mathbb{R}^d . Let $G = \{g\mathbf{z} \mid \mathbf{z} \in \mathbb{Z}^d\}$ be the set of grid points with grid distance $g \in \mathbb{R}^+$. Suppose we have a cell containing the origin and no other points of G . We can attempt to tile the space by taking all the shifts of this cell by all vectors in G . Clearly, one can do this using square tiles. However, it was an open problem to find the ‘most efficient’ way of tiling \mathbb{R}^d (in some well defined geometric sense of ‘efficient’). [KROW12] gives a randomized algorithm outputting the shape of the cell so that the entire space is fully covered and no two cells overlap (thus, it is a tiling) and each cell corresponds to one grid point. Let $C(\mathbf{x}) \in G$ ($\mathbf{x} \in \mathbb{R}^d$) denote the grid point in the cell containing \mathbf{x} (so we can think of $C(\mathbf{x})$ as a ‘rounding’ of \mathbf{x}). [KROW12] proved the following²:

Theorem 4.1 ([KROW12]). *There is a randomized algorithm partitioning the whole space \mathbb{R}^d into cells such that*

1. For every $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{s} \in G$, $C(\mathbf{x} + \mathbf{s}) = C(\mathbf{x}) + \mathbf{s}$.
2. For every two points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $\Pr[C(\mathbf{x}) \neq C(\mathbf{y})] \leq 2\pi\|\mathbf{y} - \mathbf{x}\|_2/g$.

Let $\epsilon \in (0, 1)$ and $t \in \mathbb{Z}^+$ be two parameters to be determined later. We partition \mathbb{R}^+ into sets $\mathbb{R}^+ = I_0 \cup I_1 \cup \dots \cup I_{t-1}$, where

$$I_j = \bigcup_{k \in \mathbb{Z}} \left[(1 + \epsilon)^{kt+j}, (1 + \epsilon)^{kt+j+1} \right).$$

For $I \subseteq \mathbb{R}^+$, we say an edge is *contained* in I if its length falls in I . Without loss of generality we assume I_0 is the one among $\{I_0, I_1, \dots, I_{t-1}\}$ that contains the most edges. We remove all edges not contained in I_0 . The density δ is decreased by a factor of at most t .

Recall that $I_0 = \bigcup_{k \in \mathbb{Z}} [(1 + \epsilon)^{kt}, (1 + \epsilon)^{kt+1})$. We say that the *level* of an edge is k if it is contained in $[(1 + \epsilon)^{kt}, (1 + \epsilon)^{kt+1})$. For an edge e , we use $\text{lev}(e)$ to denote its level. Let k_{\min} and k_{\max} to be the minimum level and the maximum level of all edges respectively.

For every integer $k \in [k_{\min}, k_{\max}]$ we use Theorem 4.1 to generate an (independent) random partition with grid distance

$$g_k = \frac{(1 + \epsilon)^{kt} + (1 + \epsilon)^{kt+1}}{2\alpha} = \frac{(2 + \epsilon)(1 + \epsilon)^{kt}}{2\alpha},$$

and use $C_k(\mathbf{x})$ to denote the corresponding rounding function.

Consider an edge $e = \{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ and say $\text{dir}(e) = i_0$. We assume $\langle \mathbf{v}_{j_2} - \mathbf{v}_{j_1}, \mathbf{e}_{i_0} \rangle > 0$. (Otherwise we switch the order of \mathbf{v}_{j_1} and \mathbf{v}_{j_2} .) We say the edge is *good* if the following properties are satisfied:

1. For $k = \text{lev}(e)$, $C_k(\mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0}) = C_k(\mathbf{v}_{j_2})$. Since $C_k(\mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0}) = C_k(\mathbf{v}_{j_1}) + g_k \mathbf{e}_{i_0}$, this means that the two cells containing \mathbf{v}_{j_1} and \mathbf{v}_{j_2} are adjacent along the direction \mathbf{e}_{i_0} .

²[KROW12] only considered $g = 1$ but the general result follows by simple scaling.

2. For $k > \text{lev}(e)$, $C_k(\mathbf{v}_{j_1}) = C_k(\mathbf{v}_{j_2})$. In other words, the two ends are in the same cell.

$$\mathbf{v}_{j_1} \xrightarrow{\quad \mathbf{v}_{j_2} \quad} \mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0} \xrightarrow{\quad \mathbf{e}_{i_0} \quad}$$

Claim 4.2. *Every edge is good with probability at least*

$$1 - \left(2\pi \sqrt{1 - \alpha^2 + \left(\frac{\alpha\epsilon}{2 + \epsilon} \right)^2} + \frac{4\pi\alpha(1 + \epsilon)}{(2 + \epsilon)((1 + \epsilon)^t - 1)} \right).$$

Proof: We consider the edge $e = \{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$ with direction i_0 , and assume $\langle \mathbf{v}_{j_2} - \mathbf{v}_{j_1}, \mathbf{e}_{i_0} \rangle > 0$. Then $\langle \mathbf{v}_{j_2} - \mathbf{v}_{j_1}, \mathbf{e}_{i_0} \rangle \geq \alpha \|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2$ and

$$\frac{\|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2}{g_{\text{lev}(e)}} \in \left[\frac{2\alpha}{2 + \epsilon}, \frac{2\alpha(1 + \epsilon)}{2 + \epsilon} \right) = \left[\alpha - \frac{\alpha\epsilon}{2 + \epsilon}, \alpha + \frac{\alpha\epsilon}{2 + \epsilon} \right),$$

We consider the probability that e is not a good edge.

1. For $k = \text{lev}(e)$, we have

$$\begin{aligned} \Pr [C_k(\mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0}) \neq C_k(\mathbf{v}_{j_2})] &\leq 2\pi \|\mathbf{v}_{j_2} - (\mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0})\|_2 / g_k \\ &\leq 2\pi / g_k \cdot \sqrt{\|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2^2 + g_k^2 - 2\alpha g_k \|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2} \\ &= 2\pi \sqrt{(\|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2 / g_k - \alpha)^2 + 1 - \alpha^2} \\ &\leq 2\pi \sqrt{1 - \alpha^2 + \left(\frac{\alpha\epsilon}{2 + \epsilon} \right)^2}. \end{aligned}$$

2. For $k > \text{lev}(e)$, we have

$$\begin{aligned} \Pr [C_k(\mathbf{v}_{j_1}) \neq C_k(\mathbf{v}_{j_2})] &\leq 2\pi \|\mathbf{v}_{j_2} - \mathbf{v}_{j_1}\|_2 / g_k \\ &\leq 2\pi \cdot \frac{2\alpha(1 + \epsilon)}{2 + \epsilon} \cdot \frac{g_{\text{lev}(e)}}{g_k} \\ &= \frac{4\pi\alpha(1 + \epsilon)}{2 + \epsilon} \cdot \frac{1}{(1 + \epsilon)^{(k - \text{lev}(e))t}}. \end{aligned}$$

By union bound, the probability that e is not a good edge is at most

$$\begin{aligned} &2\pi \sqrt{1 - \alpha^2 + \left(\frac{\alpha\epsilon}{2 + \epsilon} \right)^2} + \sum_{k=\text{lev}(e)+1}^{k_{\max}} \left(\frac{4\pi\alpha(1 + \epsilon)}{2 + \epsilon} \cdot \frac{1}{(1 + \epsilon)^{(k - \text{lev}(e))t}} \right) \\ &< 2\pi \sqrt{1 - \alpha^2 + \left(\frac{\alpha\epsilon}{2 + \epsilon} \right)^2} + \frac{4\pi\alpha(1 + \epsilon)}{(2 + \epsilon)((1 + \epsilon)^t - 1)}. \end{aligned}$$

Thus the claim is proved. □

For any $\alpha > \sqrt{1 - 1/(4\pi^2)}$, we can always pick ϵ sufficiently small and t sufficiently large so that each edge is good with positive probability. For example, if $\alpha = 0.99$, we can take $\epsilon = 0.01$ and $t = 500$, in which case each edge is good with probability at least 0.069. For simplicity, we use $O(\cdot)$ and $\Omega(\cdot)$ to suppress the exact values of constants α , ϵ and t . The above claim tells us every edge is good with probability $\Omega(1)$. By a simple expectation argument, there exists a series of space partitions (for every $k \in [k_{\min}, k_{\max}]$) such that $\Omega(1)$ fraction of all edges are good. We fix these partitions and remove all edges that are not good. In the remaining code the density is $\Omega(\delta)$.

Next, we prove the lower bound $n = 2^{\Omega(\delta d)}$. This follows immediately from the following lemma and Lemma 3.1.

Lemma 4.3. *For any $S \subseteq V$ with $|S| \geq 2$, there exists a non-trivial cut (S_1, S_2) such that $|\text{Edge}(S_1, S_2)| \leq \min\{|S_1|, |S_2|\}$.*

Proof: If S contains no edges, an arbitrary partition will satisfy the requirement. Otherwise, we consider the edges in S and pick an edge with the maximum level. Say this edge is $e = \{\mathbf{v}_{j_1}, \mathbf{v}_{j_2}\}$, and $\text{dir}(e) = i_0$, $\text{lev}(e) = k$. We assume $\langle \mathbf{v}_{j_2} - \mathbf{v}_{j_1}, \mathbf{e}_{i_0} \rangle > 0$. Then since this edge is good, $C_k(\mathbf{v}_{j_1})$ and $C_k(\mathbf{v}_{j_2})$ are adjacent grid points,

$$C_k(\mathbf{v}_{j_1}) + g_k \mathbf{e}_{i_0} = C_k(\mathbf{v}_{j_1} + g_k \mathbf{e}_{i_0}) = C_k(\mathbf{v}_{j_2}).$$

For any point $\mathbf{v} \in \mathbb{R}$ and $i \in [d]$, we use $C_k(\mathbf{v})_i$ to denote the i 'th coordinate of $C_k(\mathbf{v})$. Let $h = [C_k(\mathbf{v}_{j_1})_{i_0} + C_k(\mathbf{v}_{j_2})_{i_0}] / 2$. We define S_1 and S_2 as follows.

$$\begin{aligned} S_1 &= \{\mathbf{v} \in S \mid C_k(\mathbf{v})_{i_0} < h\}, \\ S_2 &= \{\mathbf{v} \in S \mid C_k(\mathbf{v})_{i_0} > h\}. \end{aligned}$$

We can see that S_1 and S_2 are not empty because $\mathbf{v}_{j_1} \in S_1$ and $\mathbf{v}_{j_2} \in S_2$. There is no point \mathbf{v} satisfying $C_k(\mathbf{v})_{i_0} = h$, because $C_k(\mathbf{v})$ is a grid point and h is not a multiple of g_k . Hence (S_1, S_2) is a non-trivial cut of S .

We consider the edges in $\text{Edge}(S_1, S_2)$, and show that every edge in $\text{Edge}(S_1, S_2)$ must have direction i_0 . Assume this is not true, and let $e' = \{\mathbf{v}'_{j_1}, \mathbf{v}'_{j_2}\}$ be such an edge. Say $\text{dir}(e') = i'$ ($i' \neq i_0$). There are two cases.

1. $\text{lev}(e') = k$. By the first requirement in the definition of good edges,

$$C_k(\mathbf{v}'_{j_1}) + g_k \mathbf{e}_{i'} = C_k(\mathbf{v}'_{j_1} + g_k \mathbf{e}_{i'}) = C_k(\mathbf{v}'_{j_2}).$$

2. $\text{lev}(e') < k$. By the second requirement in the definition of good edges,

$$C_k(\mathbf{v}'_{j_1}) = C_k(\mathbf{v}'_{j_2}).$$

In both cases we have $C_k(\mathbf{v}'_{j_1})_{i_0} = C_k(\mathbf{v}'_{j_2})_{i_0}$. Hence the edge $e' \notin \text{Edge}(S_1, S_2)$.

Therefore all edges in $\text{Edge}(S_1, S_2)$ have direction i_0 . Since the edges of the same direction are disjoint, we have $|\text{Edge}(S_1, S_2)| \leq \min\{|S_1|, |S_2|\}$. \square

5 Lower Bound for c -bounded Simple Codes

In this section we prove $n = 2^{\Omega(\alpha^2 \delta^2 d / (\log c)^2)}$ for c -bounded (α, δ) simple codes. The following simple lemma shows that it suffices to consider the 2-bounded codes.

Lemma 5.1. *A c -bounded (α, δ) simple code is a 2-bounded (α, δ') simple code for $\delta' \geq \delta / \lceil \log_2 c \rceil$.*

Proof: We partition the interval $[1, c]$ into $\lceil \log_2 c \rceil$ intervals

$$[1, c] = [2^0, 2^1) \cup [2^1, 2^2) \cup \dots \cup [2^{\lceil \log_2 c \rceil - 1}, c].$$

By the Pigeonhole Principle, there is an interval that $1/\lceil \log_2 c \rceil$ fraction of the edges have lengths in it. We only consider these edges, and scale the points in V to make all these edge lengths in $[1, 2]$. The code becomes 2-bounded and the density is at least $\delta/\lceil \log_2 c \rceil$. \square

5.1 Preliminaries and warm-up

We let $\mathbb{N} = \{0, 1, 2, \dots\}$ and for a vector $\boldsymbol{\sigma} \in \mathbb{N}^d$ we write $|\boldsymbol{\sigma}| = \sigma_1 + \dots + \sigma_d$. We denote by $\mathbb{E}_{\mathbf{x} \in_{\gamma} \mathbb{R}^d}$ the expectation with respect to a random d -dimensional vector \mathbf{x} whose entries are independent standard Gaussian random variables.

We collect some basic facts of the Hermite polynomials (see e.g., [AAR99, Section 6.1]). These polynomials form a complete orthonormal basis for the Hilbert space of square integrable functions $f : \mathbb{R}^n \rightarrow \mathbb{C}$ endowed with the inner product $(f, g) = \mathbb{E}_{\mathbf{x} \in_{\gamma} \mathbb{R}^d} [\overline{f(\mathbf{x})} g(\mathbf{x})]$. The polynomials can be obtained by Gram-Schmidt orthogonalization on the monomials $x_1^{\sigma_1} \dots x_d^{\sigma_d}$ for $\boldsymbol{\sigma} \in \mathbb{N}^d$. Each Hermite polynomial $h_{\boldsymbol{\sigma}} \in \mathbb{R}[x_1, \dots, x_d]$ can thus be uniquely represented by a nonnegative integer vector $\boldsymbol{\sigma}$ and the linear ones are $h_{\mathbf{e}_i}(\mathbf{x}) = \langle \mathbf{e}_i, \mathbf{x} \rangle = x_i$. We define the Fourier-Hermite coefficients of a function f by $\widehat{f}(\boldsymbol{\sigma}) = (h_{\boldsymbol{\sigma}}, f)$. Orthonormality easily gives *Parseval's identity*:

$$\sum_{\boldsymbol{\sigma} \in \mathbb{N}^d} \widehat{f}(\boldsymbol{\sigma})^2 = \mathbb{E}_{\mathbf{x} \in_{\gamma} \mathbb{R}^d} [|f(\mathbf{x})|^2]. \quad (3)$$

We exploit a connection between particular functions related to the Hermite polynomials and the Fourier transform over \mathbb{R}^d . Recall that the Fourier transform of a function f at a point $\mathbf{y} \in \mathbb{R}^d$ is given by

$$(\mathcal{F}(f))(\mathbf{y}) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-i\langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x}.$$

The fact we use is that the functions $H_{\boldsymbol{\sigma}}(\mathbf{x}) = e^{-\|\mathbf{x}\|_2^2/2} h_{\boldsymbol{\sigma}}(\mathbf{x})$ (known as the Hermite functions) are eigenfunctions of the Fourier transform: they satisfy $\mathcal{F}(H_{\boldsymbol{\sigma}}) = (-i)^{|\boldsymbol{\sigma}|} H_{\boldsymbol{\sigma}}$. In particular this gives the useful identity

$$\mathbb{E}_{\mathbf{x} \in_{\gamma} \mathbb{R}^d} [h_{\mathbf{e}_i}(\mathbf{x}) e^{-i\langle \mathbf{x}, \mathbf{y} \rangle}] = (\mathcal{F}(H_{\mathbf{e}_i}))(\mathbf{y}) = -i H_{\mathbf{e}_i}(\mathbf{y}) = \frac{-i \langle \mathbf{e}_i, \mathbf{y} \rangle}{e^{\|\mathbf{y}\|_2^2/2}}. \quad (4)$$

As a ‘warm-up’ to the 2-query case, we show how one can prove the following simple bound on 1-query LDCs using these tools.

Lemma 5.2. *Let $(\mathbf{v}_1, \dots, \mathbf{v}_n) \in (\mathbb{R}^d)^n$ be a 1-query (α, δ) -approximate LDC. Then, $d \leq e/(\alpha^2 \delta)$.*

Proof: Without loss of generality we may assume that the vectors \mathbf{v}_s , $s \in [n]$, have unit 2-norm. Define the *vector-valued* function $f : \mathbb{R}^d \rightarrow \mathbb{C}^n$ by $f(\mathbf{x}) = (e^{-i\langle \mathbf{v}_s, \mathbf{x} \rangle})_{s=1}^n$. Define the (vector-valued) Fourier-Hermite coefficients of f in the obvious way by $\widehat{f}(\boldsymbol{\sigma}) = \mathbb{E}_{\mathbf{x} \in \gamma \mathbb{R}^d} [h_{\boldsymbol{\sigma}}(\mathbf{x}) f(\mathbf{x})]$. Parseval's identity 3 applied to the coordinates of f gives

$$\sum_{i=1}^d \|\widehat{f}(\mathbf{e}_i)\|_2^2 \leq \sum_{\boldsymbol{\sigma} \in \mathbb{N}^d} \|\widehat{f}(\boldsymbol{\sigma})\|_2^2 = \mathbb{E}_{\mathbf{x} \in \gamma \mathbb{R}^d} [\|f(\mathbf{x})\|_2^2]. \quad (5)$$

The right-hand side of (5) clearly equals n . By (4) the left-hand side is at least

$$\sum_{i=1}^d \|\widehat{f}(\mathbf{e}_i)\|_2^2 = \sum_{i=1}^d \sum_{s=1}^n \left| \mathbb{E}_{\mathbf{x} \in \gamma \mathbb{R}^d} [h_{\mathbf{e}_i}(\mathbf{x}) e^{-i\langle \mathbf{v}_s, \mathbf{x} \rangle}] \right|^2 = \sum_{i=1}^d \sum_{s=1}^n \frac{|\langle \mathbf{e}_i, \mathbf{v}_s \rangle|^2}{e} \geq \frac{\delta d n \alpha^2}{e}, \quad (6)$$

where the last inequality follows from the definition of a 1-query approximate LDC. Putting things together gives $\alpha^2 \delta d n / e \leq n$, which implies the result. \square

5.2 A matrix valued function from a 2-query code

Let $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ be a (2-query) 2-bounded (α, δ) simple code. We define the vector-valued function $f : \mathbb{R}^d \rightarrow \mathbb{C}^n$ given by $f(\mathbf{x}) = (e^{-i\langle \mathbf{x}, \mathbf{v}_s \rangle})_{s=1}^n$. And from f we define the matrix-valued function $F(\mathbf{x}) = f(\mathbf{x}) f(\mathbf{x})^*$, so

$$F(\mathbf{x}) = \left(e^{-i\langle \mathbf{x}, \mathbf{v}_s - \mathbf{v}_t \rangle} \right)_{s,t=1}^n.$$

Note that each $F(\mathbf{x})$ is a Hermitian matrix with rank 1. We define the matrix-valued Fourier-Hermite coefficients $\widehat{F}(\boldsymbol{\sigma})$ in the obvious way by $\widehat{F}(\boldsymbol{\sigma})_{s,t} = (h_{\boldsymbol{\sigma}}, F_{s,t})$, where $F_{s,t}$ is the function corresponding to the (s, t) -coordinate of F .

By (4) $\widehat{F}(\mathbf{e}_i)$ therefore has as (s, t) -entry given by

$$\widehat{F}(\mathbf{e}_i)_{s,t} = \mathbb{E}_{\mathbf{x} \in \gamma \mathbb{R}^d} \left[h_{\mathbf{e}_i}(\mathbf{x}) e^{-i\langle \mathbf{x}, \mathbf{v}_s - \mathbf{v}_t \rangle} \right] \stackrel{(4)}{=} \frac{-i\langle \mathbf{e}_i, \mathbf{v}_s - \mathbf{v}_t \rangle}{e^{\|\mathbf{v}_s - \mathbf{v}_t\|_2^2/2}}.$$

Since the code is simple and 2-bounded there are $|M_i|$ disjoint $\{s, t\}$ -pairs such that the (s, t) -entry of $\widehat{F}(\mathbf{e}_i)$ has magnitude

$$|\widehat{F}(\mathbf{e}_i)_{s,t}| \geq \frac{|\langle \mathbf{e}_i, \mathbf{v}_s - \mathbf{v}_t \rangle|}{e^{\|\mathbf{v}_s - \mathbf{v}_t\|_2^2/2}} \geq \frac{\alpha}{e}. \quad (7)$$

The matrix $\widehat{F}(\mathbf{e}_i)$ thus has large entries (in absolute values) on the coordinates corresponding to the matching M_i .

5.3 A Fourier inequality for the trace norm

We now establish a matrix analog of (5) (Lemma 5.3 below), which is expressed in terms of the Schatten-1 norm (or trace norm). The *Schatten-1 norm* $\|A\|_{S_1}$ of a complex matrix A is defined as the sum of its singular values. We also use the following dual characterization of this norm. For a pair of matrices $A, X \in \mathbb{C}^{n \times n}$ let $\langle A, X \rangle = \text{tr}[A^* X]$ denote their trace inner product, where A^*

denotes the conjugate transpose of A . The *spectral norm* $\|X\|_{S_\infty}$ of a matrix X is defined as its maximum singular value. We have the well-known duality characterization

$$\|A\|_{S_1} = \max\{|\langle A, X \rangle| : \|X\|_{S_\infty} \leq 1\}. \quad (8)$$

The analog of (5) that allows us to prove the lower bound on 2-bounded simple codes is as follows.

Lemma 5.3. *Let $F : \mathbb{R}^d \rightarrow \mathbb{C}^{n \times n}$ be a Hermitian matrix-valued function. Then,*

$$\left(\sum_{i=1}^d \|\widehat{F}(\mathbf{e}_i)\|_{S_1}^2 \right)^{1/2} \leq \sqrt{2 \log(2en)} \left(\mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^n} [\|F(\mathbf{x})\|_{S_1}^2] \right)^{1/2}.$$

The proof of this lemma relies on the following non-commutative version of the Khintchine inequality [Tro12, Section 4.4].

Theorem 5.4 (Non-commutative Khintchine inequality [Tro12]). *For any collection of Hermitian matrices $A_1, \dots, A_d \in \mathbb{C}^{n \times n}$ and i.i.d. standard Gaussian random variables x_1, \dots, x_d , we have*

$$\left(\mathbb{E} \left[\left\| \sum_{i=1}^d x_i A_i \right\|_{S_\infty}^2 \right] \right)^{1/2} \leq \sqrt{2 \log(2en)} \left\| \sum_{i=1}^d A_i^2 \right\|_{S_\infty}^{1/2}.$$

Proof of Lemma 5.3: By homogeneity we may assume that $\|\widehat{F}(\mathbf{e}_1)\|_{S_1}^2 + \dots + \|\widehat{F}(\mathbf{e}_d)\|_{S_1}^2 = 1$. Let $X_1, \dots, X_d \in \mathbb{C}^{n \times n}$ be such that $\|X_i\|_{S_\infty} \leq 1$ and $\langle \widehat{F}(\mathbf{e}_i), X_i \rangle = \|\widehat{F}(\mathbf{e}_i)\|_{S_1}$ for every $i \in [d]$. Let $Y_i = \|\widehat{F}(\mathbf{e}_i)\|_{S_1} X_i$ and notice that

$$\sum_{i=1}^d \|Y_i\|_{S_\infty}^2 \leq \sum_{i=1}^d \|\widehat{F}(\mathbf{e}_i)\|_{S_1}^2 = 1. \quad (9)$$

We consider the quantity

$$\mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^d} \left[\left\langle F(\mathbf{x}), \sum_{i=1}^d x_i Y_i \right\rangle \right]. \quad (10)$$

First, by linearity of the trace function, it equals

$$\sum_{i=1}^d \left\langle \mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^d} [x_i F(\mathbf{x})], Y_i \right\rangle = \sum_{i=1}^d \langle \widehat{F}(\mathbf{e}_i), Y_i \rangle = \sum_{i=1}^d \|\widehat{F}(\mathbf{e}_i)\|_{S_1}^2 = 1.$$

Second, by Hölder's inequality for the trace and spectral norms [Bha97] (which follows from (8)) and the Cauchy-Schwarz inequality, (10) is at most

$$\mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^d} \left[\|F(\mathbf{x})\|_{S_1} \left\| \sum_{i=1}^d x_i Y_i \right\|_{S_\infty} \right] \leq \left(\mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^d} [\|F(\mathbf{x})\|_{S_1}^2] \right)^{1/2} \left(\mathbb{E}_{\mathbf{x} \in_\gamma \mathbb{R}^d} \left[\left\| \sum_{i=1}^d x_i Y_i \right\|_{S_\infty}^2 \right] \right)^{1/2}.$$

By Theorem 5.4, the triangle inequality and the fact $\|Y_i^2\|_{S_\infty} \leq \|Y_i\|_{S_\infty}^2$, the last factor is at most

$$\sqrt{2 \log(2en)} \left\| \sum_{i=1}^d Y_i^2 \right\|_{S_\infty}^{1/2} \leq \sqrt{2 \log(2en)} \left(\sum_{i=1}^d \|Y_i\|_{S_\infty}^2 \right)^{1/2} \leq \sqrt{2 \log(2en)}.$$

□

5.4 Lower bound on 2-bounded codes

We now combine fact (7) and Lemma 5.3 to lower bound the length of 2-bounded simple codes. Recall that we defined the matrix-valued function $F(\mathbf{x}) = f(\mathbf{x})f(\mathbf{x})^*$ where $f(\mathbf{x}) = (e^{-i\langle \mathbf{x}, \mathbf{v}_s \rangle})_{s=1}^n$. Then $F(\mathbf{x})$ has $f(\mathbf{x})$ as an eigenvector with eigenvalue n (its other eigenvalues being zero). Hence,

$$\|F(\mathbf{x})\|_{S_1} = n. \quad (11)$$

Recall from (7) that for every $i \in [d]$ there are $|M_i|$ disjoint $\{s, t\}$ -pairs such that $|\widehat{F}(\mathbf{e}_i)_{s,t}| \geq \alpha/e$. From [BARdW08, p. 14–15] it directly follows that

$$\|\widehat{F}(\mathbf{e}_i)\|_{S_1} \geq \frac{\alpha}{e}|M_i|. \quad (12)$$

Putting the above facts together gives

$$\begin{aligned} \frac{\alpha\delta\sqrt{dn}}{e} &\leq \frac{\alpha}{e} \frac{1}{\sqrt{d}} \sum_{i=1}^d |M_i| \leq \left(\frac{\alpha^2}{e^2} \sum_{i=1}^d |M_i|^2 \right)^{1/2} \stackrel{(12)}{\leq} \left(\sum_{i=1}^d \|\widehat{F}(\mathbf{e}_i)\|_{S_1}^2 \right)^{1/2} \\ &\stackrel{\text{Lemma 5.3}}{\leq} \sqrt{2\log(2en)} \left(\mathbb{E}_{\mathbf{x} \in \gamma \mathbb{R}^d} [\|F(\mathbf{x})\|_{S_1}^2] \right)^{1/2} \stackrel{(11)}{=} \sqrt{2\log(2en)} n, \end{aligned}$$

where the second inequality follows from Cauchy-Schwarz. Hence, $n = 2^{\Omega(\alpha^2\delta^2d)}$.

6 Approximate q -query Code for General q

In this section we prove Theorem 1.8 by showing that $n = \Omega((\alpha^2\delta^{1/q}d)^{\frac{q}{q-1}})$ for general q -query approximate code. The proof is similar to [KT00]: we select a random subset of V with size $\Theta(\delta^{-\frac{1}{q}n^{\frac{q-1}{q}}})$, and show that w.h.p it contains a q -tuple from at least $\Omega(d)$ matchings (or directions). This will imply the size of the subset is $\Omega(\alpha^2d)$. The lower bound of n follows immediately.

We first note that a subset containing tuples from many different matchings must be large.

Lemma 6.1. *If a set $S \subseteq V$ contains at least one tuple from k different matchings ($k \leq d$), then $|S| \geq \alpha^2k$.*

Proof: By the definition of an approximate LDC, for every tuple $\{j_1, \dots, j_k\} \in M_i$, there exists a unit vector $\mathbf{u} \in \text{span}\{\mathbf{v}_{j_1}, \dots, \mathbf{v}_{j_k}\}$ with the i 'th coordinate at least α in absolute value. We also assume w.l.o.g. that $u_i \geq \alpha$ (otherwise take $-\mathbf{u}$). Therefore, there exists unit vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k \in \text{span}\{S\}$ such that each of them has a different coordinate at least α . Without loss of generality we assume $u_{11}, u_{22}, \dots, u_{kk} \geq \alpha$.

To show $|S| \geq \alpha^2k$, it suffices to show $\text{rank}\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\} \geq \alpha^2k$. Let U be the matrix consisting of $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k$ as its row vectors. For simplicity we remove the last $n - k$ columns in U if $n > k$. Now U is a square matrix with diagonal elements at least α . Let r be the rank of U . We only need to show $r \geq \alpha^2k$. This is a variant of the well-known theorem saying that the rank of a matrix is large if its diagonal elements are larger than the off-diagonal ones. We give the following proof, which is based on the idea of Lemma 3.5 in [BDWY12] (similar lemmas can be found in most standard texts on matrix analysis).

Let $U = Q\Sigma P^*$ be the singular value decomposition of U , where Q, P are unitary matrices. Let $\sigma_1, \sigma_2, \dots, \sigma_r > 0$ be the non-zero singular values of U , i.e. $\Sigma = \text{diag}\{\sigma_1, \sigma_2, \dots, \sigma_r, 0, 0, \dots, 0\}$. We have

$$(\alpha k)^2 = \text{tr}(U)^2 = \text{tr}(Q\Sigma P^*)^2 = \text{tr}((P^*Q)\Sigma)^2 \leq \text{tr}(\Sigma)^2.$$

The last inequality holds since P^*Q is a unitary matrix and every element has absolute value at most 1. On the other hand,

$$\text{tr}(\Sigma)^2 = (\sigma_1 + \sigma_2 + \dots + \sigma_r)^2 \leq r \cdot (\sigma_1^2 + \sigma_2^2 + \dots + \sigma_r^2) = r \cdot \|U\|_F^2 \leq r \cdot k.$$

Combine these two inequalities we have $r \geq \alpha^2 k$. □

Now we can prove the theorem in the same way as in [KT00]. We pick a random set $S \subseteq V$ of size $\Theta(\delta^{-\frac{1}{q}} n^{\frac{q-1}{q}})$. By Lemma 5 in [KT00], S contains tuples in $\Omega(d)$ different directions in expectation. We fix an S that contains tuples in $\Omega(d)$ directions. By Lemma 6.1, we have

$$\delta^{-\frac{1}{q}} n^{\frac{q-1}{q}} = \Omega(\alpha^2 d).$$

The lower bound $n = \Omega((\alpha^2 \delta^{1/q} d)^{\frac{q}{q-1}})$ follows immediately.

References

- [AAR99] G. E. Andrews, R. Askey, and R. Roy. *Special Functions*, volume 71 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1999.
- [ADSW12] A. Ai, Z. Dvir, S. Saraf, and A. Wigderson. Sylvester-Gallai theorems for approximate collinearity. *Forum of mathematics - Sigma* (to appear), 2012.
- [BARdW08] A. Ben-Aroya, O. Regev, and R. de Wolf. A hypercontractive inequality for matrix-valued functions with applications to quantum computing and LDCs. In *FOCS'08*, pages 477–486, 2008.
- [BDWY12] B. Barak, Z. Dvir, A. Wigderson, and A. Yehudayoff. Fractional Sylvester-Gallai theorems. *Proceedings of the National Academy of Sciences*, 2012.
- [BET10] A. Ben-Aroya, K. Efremenko, and A. Ta-Shma. Local list decoding with a constant number of queries. In *FOCS'10*, pages 715–722, 2010.
- [BF90] D. Beaver and J. Feigenbaum. Hiding instances in multioracle queries. In *STACS'90*, pages 37–48, 1990.
- [Bha97] R. Bhatia. *Matrix Analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer, 1997.
- [BK95] M. Blum and S. Kannan. Designing programs that check their work. *J. ACM*, 42(1):269–291, January 1995.
- [BNR12] J. Briët, A. Naor, and O. Regev. Locally decodable codes and the failure of cotype for projective tensor products. *Electronic Research Announcements in Mathematical Sciences (ERA-MS)*, 19:120–130, 2012.

- [CFL⁺10] Y. M. Chee, T. Feng, S. Ling, H. Wang, and L. F. Zhang. Query-efficient locally decodable codes of subexponential length. *Electronic Colloquium on Computational Complexity (ECCC)*, TR10-173, 2010.
- [DGY11] Z. Dvir, P. Gopalan, and S. Yekhanin. Matching vector codes. *SIAM J. Comput.*, 40(4):1154–1178, 2011.
- [DS05] Z. Dvir and A. Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. In *STOC’05*, pages 592–601, 2005.
- [DSW12] Z. Dvir, S. Saraf, and A. Wigderson. Improved rank bounds for design matrices and a new proof of Kelly’s theorem. *Forum of mathematics - Sigma* (to appear), 2012.
- [DSW13] Z. Dvir, S. Saraf, and A. Wigderson. Breaking the quadratic barrier for 3-LCCs over the reals. Manuscript, 2013.
- [Efr09] K. Efremenko. 3-query locally decodable codes of subexponential length. In *STOC’09*, pages 39–44, 2009.
- [GKST06] O. Goldreich, H. Karloff, L. J. Schulman, and L. Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. *Computational Complexity*, 15(3):263–296, 2006.
- [IS10] T. Itoh and Y. Suzuki. Improved constructions for query-efficient locally decodable codes of subexponential length. *IEICE Transactions on Information and Systems*, E93-D(2):263–270, 2010.
- [KdW04] I. Kerenidis and R. de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004.
- [KROW12] G. Kindler, A. Rao, R. O’Donnell, and A. Wigdersons. Spherical cubes: optimal foams from computational hardness amplification. *Commun. ACM*, 55(10):90–97, October 2012.
- [KT00] J. Katz and L. Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *STOC’00*, pages 80–86, New York, NY, USA, 2000. ACM.
- [KY09] K. S. Kedlaya and S. Yekhanin. Locally decodable codes from nice subsets of finite fields and prime factors of Mersenne numbers. *SIAM J. Comput.*, 38(5):1952–1969, 2009.
- [Lip90] R. J. Lipton. Efficient checking of computations. In *STACS’90*, pages 207–215, 1990.
- [Pra07] R. Prasad. A note on Yekhanin’s locally decodable codes. *Electronic Colloquium on Computational Complexity (ECCC)*, TR07-016, 2007.
- [Tro12] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

- [Woo07] D. P. Woodruff. New lower bounds for general locally decodable codes. *Electronic Colloquium on Computational Complexity (ECCC)*, TR07-006, 2007.
- [Woo12] D. P. Woodruff. A quadratic lower bound for three-query linear locally decodable codes over any field. *Journal of Computer Science and Technology*, 27(4):678–686, 2012.
- [Yek08] S. Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM*, 55(1):1–16, 2008.