Electronic Colloquium on Computational Complexity, Report No. 86 (2015)

On Embeddings of ℓ_1^k from Locally Decodable Codes

Jop Briët*

Abstract

We show that any q-query locally decodable code (LDC) gives a copy of ℓ_1^k with small distortion in the Banach space of q-linear forms on $\ell_{p_1}^N \times \cdots \times \ell_{p_q}^N$, provided $1/p_1 + \cdots + 1/p_q \leq 1$ and where k, N, and the distortion are simple functions of the code parameters. We exhibit the copies of ℓ_1^k by constructing a basis directly from "smooth" LDC decoders, thus bypassing a matching lemma often used in the LDC literature. Based on this, we give alternative proofs for known lower bounds on the length of 2-query LDCs. Using similar techniques, we reprove known lower bounds for larger q. We also discuss the relation with an alternative proof, due to Pisier, of a result of Naor, Regev, and the author on cotype properties of projective tensor products of ℓ_p spaces.

1 Introduction

A locally decodable code (LDC) is an error correcting code that maps a message string into a codeword such that, even if part of the codeword is changed adversarially, any single message symbol can be retrieved by querying only a small number of randomly selected codeword coordinates. More formally, for positive integers k, N, and q, real numbers $\delta, \varepsilon \in (0, 1/2]$, and a finite alphabet Γ, a map $C : \{0,1\}^k \to \Gamma^N$ is a (q, δ, ε)-*locally decodable code* if there exists a *decoder* (a probabilistic algorithm) $\mathcal{A}$ such that:

- For every message $x \in \{0,1\}^k$, index $i \in [k]$, and string $y \in \Gamma^N$ that differs from the codeword $C(x)$ in at most δN coordinates,
$$\Pr[\mathcal{A}(i, y) = x_i] \geq \frac{1}{2} + \varepsilon.$$
- $\mathcal{A}$ (non-adaptively) queries at most q coordinates of y.

The most general decoder first samples a set $S \subseteq [N]$ of at most q codeword coordinates from a probability distribution that depends on i only. Then, it outputs a random bit whose distribution depends only on i, S, and the sequence $(y_s)_{s \in S}$ of (possibly corrupted) codeword entries at S.[1]

The central problem regarding LDCs is to determine the smallest possible codeword length N as a function of the message length k for various ranges of the query complexity q and alphabet size $|\Gamma|$ when δ and ε are fixed constants.

*Center for Mathematics & Computer Science (CWI), The Netherlands. Funded by a Rubicon grant from the Netherlands Organisation for Scientific Research (NWO). E-mail: j.briet@cwi.nl

[1] Adaptive decoders, whose queries depend on the values of previously queried coordinates, can be made non-adaptive at the cost of a factor $1/|\Gamma|^{q-1}$ in the decoding bias ε.

ISSN 1433-8092

The topic of this paper (summarized in Theorem 1.1) is a link between LDCs and a geometric property of certain Banach spaces of q-linear forms: the property of containing copies of ℓ_1^k with small distortion. This link follows implicitly from work of Pisier [Pis73] and Naor, Regev, and the author [BNR12]. But while presenting an alternative route to the main result of [BNR12] in a workshop, Pisier [Pis12] made the link explicit using basic properties of an LDC construction of Efremenko [Efr09], which unfortunately appears to be unpublished. Here we present a somewhat more general and direct version of this result, but this paper may be seen as partly expository.

In particular, we explicitly construct copies of ℓ_1^k from general LDCs, while bypassing the "matching lemma" [KT00, BARdW08] (also see Appendix A), which is a handy and often-used tool in the analysis of LDCs, but which also appears to necessitate Elton's theorem (also see Appendix A), which was used by Pisier. The matching lemma identifies for each $i \in [k]$ a q-matching of codeword coordinates from which, *on average* over k-bit messages, the ith message bit can be decoded with good probability from an uncorrupted codeword. Based on the q-matchings, one can find q-linear forms $A_1, \ldots, A_k$ such that the average norm of a randomly signed sum $\pm A_1 + \cdots + \pm A_k$ is close to k. Elton's theorem then asserts that for some $\eta \in (0, 1)$, an η-fraction of these forms span a (distorted) copy of $\ell_1^{\eta k}$. We avoid these two steps and construct copies of ℓ_1^k using only a prerequisite for the matching lemma, namely the fact that any decoder must sample from "smooth" distribution in which no one coordinate is favored.

The main complexity-theoretic message of this paper is that due to a known upper bound on the dimension k for which certain spaces of bilinear forms can accommodate ℓ_1^k, we obtain new proofs for known lower bounds on the length of 2-query LDCs (over binary alphabets and alphabets of non-constant size). More generally, the above-mentioned link suggests a new avenue to explore for proving such bounds when $q \geq 3$, for which techniques are currently in short supply. In similar geometric spirit, but inspired by techniques used by Kerenidis and de Wolf [KdW04], we reprove known lower bounds for LDCs with a larger number of queries in the Appendix.

Origins and applications. The notion of LDCs originated from works on probabilistically checkable proofs [BFLS91, Sud92] and private information retrieval (PIR) [CGKS98], though they where first formally defined by Katz and Trevisan [KT00] in the context of noisy data transmission. Since then, the range of areas where these codes turn out to play a role has grown steadily. Applications in theoretical computer science now include polynomial identity testing [DS07], data structures [Wol09, CGW13], and complexity theory [Dvi10]. In pure mathematics, they recently found applications in discrete geometry [BDYW11, BDHS14] and Banach spaces [BNR12].

Constructions. Currently four constructions roughly cover the best-known trade-offs between codeword length, query complexity, and alphabet size. The family of Reed-Muller codes, which work based on polynomial interpolation, give LDCs for a large range of parameters [BFLS91]. For example, the Hadamard code is a binary 2-query LDC of length $N = 2^k$ and for constant-sized alphabets the Reed-Muller family gives LDCs with query complexity $q = \text{poly}(\log k)$ and length $\text{poly}(k)$. Great strides were made recently with the discovery of Matching Vector codes (MV-codes) [Yek07, Efr09, DGY10] and Multiplicity codes [KSY11], which outperform Reed-Muller codes in constant and $\text{poly}(k)$ query-complexity regimes, respectively. See [Yek12] for a detailed survey, and for very recent work on high query complexity expander-based codes, see [HOW14, KMRZS15]. Since our focus will be on the constant query complexity regime, we highlight that for constant $q \geq 3$ there are q-query MV-codes with constant alphabet size and length $\exp\exp(o(\log k))$.

The best constructions of LDCs over large alphabets come from PIR schemes.[2] A PIR scheme replicates a k-bit database among $q \geq 2$ non-communicating servers that interact with a user wishing to know some entry $i \in [k]$ of the database that he/she wants to keep hidden from the servers. The goal is to find a scheme that achieves the above with minimal communication between the user and the servers. Katz and Trevisan [KT00] observed (and [GKST06] showed formally) that q-query LDCs are essentially equivalent to q-server PIRs where communication proceeds in two rounds and the total number of communicated bits per index $i \in [k]$ is given by $2\log(|\Gamma|N)$. A recent breakthrough of Dvir and Gopi [DG14] gave two-round q-server PIRs with communication cost $\exp(o(\log k))$; these schemes rely on the same combinatorial objects, called "matching vector families," as MV-codes. Most remarkably, their construction shows that there are 2-query LDCs whose alphabet size *and* length is $\exp\exp(o(\log k))$.

Lower bounds. What we know about the *necessary* length of LDCs has changed little during the last decade and most of the best-known lower bounds are far from the parameters of the best-known constructions. There are currently two general cases where optimal bounds are known. First, it was shown in [KT00] that independent of the code length, 1-query LDCs can only encode a constant number of message bits once we fix δ, ε, and $|\Gamma|$. The second case concerns binary 2-query LDCs. Those turn out to require exponential length, as is achieved by the Hadamard code. The original proof of the exponential bound due to Kerenidis and de Wolf [KdW04], which is based on quantum-information-theoretic arguments, gives the bound

$$N \geq 2^{\Omega(\delta\varepsilon^2 k)}.$$

Ben-Aroya, Regev and de Wolf [BARdW08] obtained a similar bound using a Fourier-analytic inequality for matrix-valued functions, which they derived from a deep result from Banach space theory on uniform convexity of Schatten-1 [BCL94]. These proofs also form the basis for the best-known lower bounds for $q \geq 3$. For even integers $q \geq 4$ and constant δ and ε, [KdW04] used a reduction to maps akin to 2-query LDCs to prove that binary q-query LDCs have length $\Omega((k/\log k)^{q/(q-2)})$. A similar reduction gives the same bound based on [BARdW08]. Later, Woodruff [Woo07] slightly improved this bound to $\Omega(k^{q/(q-2)}/\log k)$ using a more careful reduction. Oddly, for odd $q \geq 3$, we do not know how to prove better lower bounds other than by using the ones for $q+1$ queries.

For 2-query LDCs over large alphabets, Wehner and de Wolf [WdW05] proved the bound $|\Gamma|^2 \log N \geq \Omega(k)$, which implies an $\Omega(\log k)$ bound on the communication required in any (two-round) 2-server PIR scheme.[3] Their proof also used quantum information theory. Here too, there thus remains a large gap with the best construction. Slightly better bounds are known when the alphabet is $\{0,1\}^n$ and the decoder, after sampling a set of codeword coordinates, returns a random bit whose distribution depends on at most $m \leq n$ predetermined bits at those coordinates. Such codes may be seen as PIR schemes where the servers send the user an n-bit string of which the user only reads at most m bits. This happens, for example, in [DG14], where the best-known constructions of matching vectors [Gro00] give $m \approx \sqrt{n}$. In [WdW05] it is proved that

$$2^m \sum_{l=0}^{m} \binom{n}{l} \log N \geq \Omega(k).$$

[2] With information-theoretic security.

[3] The current best constant is obtained by combining their result with the bound $\log N \geq 2\log k - 2\log|\Gamma| - O(1)$ due to [KT00], which gives $(5 - o(1))\log k$.

For example, if $m = n^\eta$ for some constant $\eta \in (0,1)$, this implies a bound of $\Omega((\log k)^{1/\eta - o(1)})$ on the communication for two-round two-server PIRs.

Banach space geometry. Different results from Banach space theory were used on several occasions to prove lower bounds on LDCs or similar objects [BARdW08, DSW14, BDHS14]. In the opposite direction, the aforementioned 3-query MV-codes were used in [BNR12] to *solve* an open problem on Banach spaces. The following basic definitions and facts will allow us to elaborate. For $p \in [1,\infty]$, a distortion parameter $K \geq 1$, and a positive integer d, a Banach space X is said to contain a K-isomorphic copy of ℓ_p^d if there exist $A_1, \ldots, A_d \in X$ such for any vector $\alpha \in \mathbf{R}^d$,

$$\|\alpha\|_{\ell_p} \leq \Big\| \sum_{i=1}^d \alpha_i A_i \Big\|_X \leq K \|\alpha\|_{\ell_p}.$$

The containment of copies of certain finite-dimensional ℓ_p spaces is strongly linked with the notions of (Rademacher) type and cotype, which are defined as follows. The space X has *type* $p > 0$ if there exists a constant $T < \infty$ such that for every positive integer d and $A_1, \ldots, A_d \in X$, we have

$$\mathbb{E}_{x \in \{-1,1\}^d} \Big\| \sum_{i=1}^d x_i A_i \Big\|_X \leq T \Big(\sum_{i=1}^d \|A_i\|_X^p \Big)^{1/p}. \tag{1}$$

Observe that the right-hand side decreases as p increases and that by the triangle inequality, any space has type 1. We say that a space *fails* nontrivial type if there is no $p > 1$ for which it has type p. The infimum over T satisfying (1) for any $d \in \mathbb{N}$ and $A_1, \ldots, A_d \in X$ is denoted by $T_p(X)$.

A space X has *cotype* $r > 0$ if there exists a constant $C < \infty$ such that for every positive integer d and $A_1, \ldots, A_d \in X$, we have

$$\mathbb{E}_{x \in \{-1,1\}^d} \Big\| \sum_{i=1}^d x_i A_i \Big\|_X \geq \frac{1}{C} \Big(\sum_{i=1}^d \|A_i\|_X^r \Big)^{1/r}. \tag{2}$$

By convexity of norms and Jensen's inequality, any space has cotype ∞ and we say that a space *fails* finite cotype if there is no $r < \infty$ such that it has cotype r. The infimum over C satisfying (2) for any $d \in \mathbb{N}$ and $A_1, \ldots, A_d \in X$ is denoted by $C_r(X)$.

As a well-behaved example, Hilbert space has type 2 and cotype 2. Two important examples that fail one or the other are ℓ_1, which fails nontrivial type, and ℓ_∞, which fails finite cotype; both failures are easily seen by setting the A_i to be distinct standard basis vectors. It turns out that these are not just some examples that fail either nontrivial type or cotype, but in the sense alluded to above, they are the only examples. Indeed, Pisier [Pis73] showed that a Banach space X fails nontrivial type if and only if there exists a $K < \infty$ such that X contains a K-isomorphic copy of ℓ_1^d for every positive integer d. Complementing this, Maurey and Pisier [MP73] equated failure of finite cotype with containment of a K-isomorphic copy of ℓ_∞^d for every d.

LDCs and copies of ℓ_p^d. The following Banach spaces are relevant to LDCs. For positive integers N and $q \geq 2$, and a vector $\mathbf{p} = (p_1, \ldots, p_q) \in (1,\infty)^q$ such that $1/p_1 + \cdots + 1/p_q \leq 1$, we shall consider the real N^q-dimensional vector space of q-linear forms on $\mathbf{R}^N$ endowed with the norm

$$\|A\|_{\mathbf{p}} = \sup \Big\{ \frac{A(z[1], \ldots, z[q])}{\|z[1]\|_{\ell_{p_1}} \cdots \|z[q]\|_{\ell_{p_q}}} : z[1], \ldots, z[q] \in \mathbf{R}^N \setminus \{\mathbf{0}\} \Big\}.$$

We denote this Banach space by $\mathcal{L}(N;\mathbf{p})$. Note that $\mathcal{L}(N;(2,2))$ can be identified with the space of matrices endowed with the Schatten-∞ norm and that the spaces $(\mathcal{L}(N;\mathbf{p}))_{N\in\mathbb{N}}$ are subspaces of the Banach space of bounded q-linear forms on $\ell_{p_1}\times\cdots\times\ell_{p_q}$, which we denote by $\mathcal{B}(\mathbf{p})$.

In [BNR12] it is shown that for fixed q, δ, ε, and vector $\mathbf{p}$ as above, the existence of an infinite family of binary q-query LDCs with sub-exponential length implies that for any $r\in[2,\infty)$, the cotype-r constant of the *dual* of $\mathcal{L}(N;\mathbf{p})$ satisfies

$$\lim_{N\to\infty} C_r\big(\mathcal{L}(N;\mathbf{p})^*\big)=\infty. \tag{3}$$

Since the MV-codes of [Efr09] have sub-exponential length, the above holds for $q\geq 3$. It follows that the infinite-dimensional space $\mathcal{B}(\mathbf{p})^*$ fails finite cotype, which allowed [BNR12] to answer in the negative a question of [DFS03] on the permanence of finite cotype under the projective tensor product.[4] This in turn has implications for the space $\mathcal{B}(\mathbf{p})$ itself. For any Banach space X and any $p,r\in(1,\infty)$ such that $1/p+1/r=1$, it holds that $T_p(X)\geq C_r(X^*)$ [Pis99, Proposition 3.2]. It thus follows from (3) that for any $p>1$, the type-p constants of $\mathcal{L}(N;\mathbf{p})$ are unbounded and hence $\mathcal{B}(\mathbf{p})$ fails nontrivial type. The LDCs therefore imply that there exists a $K<\infty$ such that for every $d\in\mathbb{N}$, the space $\mathcal{B}(\mathbf{p})$ contains a K-isomorphic copy of ℓ_1^d.

1.1 Main result

Given that LDCs imply the *existence* of copies of ℓ_1^d in $\mathcal{B}(\mathbf{p})$, it is natural to ask what these copies look like. Here we give explicit constructions of those copies. After stating the main theorem we shall elaborate on its implications for LDC lower bounds and Banach space geometry.

Theorem 1.1. *Let k, N, and $q\geq 2$ be positive integers, $\delta,\varepsilon\in(0,1/2]$, and let Γ be a finite set. Assume there exists a (q,δ,ε)-LDC from $\{0,1\}^k$ to Γ^N. Then, for any $\mathbf{p}\in(1,\infty)^q$ such that $1/p_1+\cdots+1/p_q\leq 1$, every integer $N'\geq 2|\Gamma|N$, and any real number $K\geq 2^q|\Gamma|^{(q+2)/2}/(\delta\varepsilon)$, the space $\mathcal{L}(N';\mathbf{p})$ contains a K-isomorphic copy of ℓ_1^k.*

That is, there exist q-linear forms $A_1,\ldots,A_k$ on $\ell_{p_1}^{N'}\times\cdots\times\ell_{p_q}^{N'}$ (that we give explicitly) such that for any vector $\alpha\in\mathbf{R}^k$,

$$\|\alpha\|_{\ell_1}\leq\Big\|\sum_{i=1}^k\alpha_iA_i\Big\|_{\mathbf{p}}\leq K\|\alpha\|_{\ell_1}. \tag{4}$$

Moreover, if for positive integers $m\leq n$, we have $\Gamma=\{0,1\}^n$ and the LDC decoder's output depends on at most m predetermined bits of each queried codeword symbol, then the above holds for

$$N'\geq\binom{n}{\leq m}N\qquad\text{and}\qquad K\geq q\binom{n}{\leq m}^{(q+2)/2}/\delta\varepsilon,$$

where $\binom{n}{\leq m}=\binom{n}{0}+\binom{n}{1}+\cdots+\binom{n}{m}$.

[4] The space $\mathcal{B}(\mathbf{p})^*$ is precisely the projective tensor product of the spaces $\ell_{p_1},\ldots,\ell_{p_q}$, denoted $\ell_{p_1}\widehat{\otimes}\cdots\widehat{\otimes}\ell_{p_q}$ [Rya02, Chapter 2, Section 2.2]. The result is stated only for the case $q=3$ but the same proof works when $q\geq 3$.

Application for LDC lower bounds. If the known LDC lower bounds leave any room for improvement, then the near lack thereof in the last decade could indicate that new techniques are needed to make progress. Alternative techniques with which the known bounds can be reproved were already asked for by Trevisan [Tre04, Question 3]. Theorem 1.1 gives a method based on showing that $\mathcal{L}(N;\mathbf{p})$ contains no copies of ℓ_1^d for large dimension d and small distortion. As we show in Section 4, via this method we immediately recover the above-mentioned lower bounds for 2-query LDCs (up-to slightly poorer dependence on δ and $|\Gamma|$). Indeed, previous results easily imply that any $O(1)$-isomorphic copy of ℓ_1^d in $\mathcal{L}(N;(2,2))$ must satisfy $d \le O(\log N)$.

Application for cotype. As observed by Pisier [Pis12], Theorem 1.1 also gives an alternative route from LDCs to the result (3) of [BNR12]. Indeed, for $q \ge 3$, the theorem combined with the parameters of q-query MV-codes of [Efr09] implies that $\mathcal{L}(N;\mathbf{p})$ contains an $O(1)$-isomorphic copy of ℓ_1^d for $d \ge (\log N)^{\omega(1)}$—in stark contrast with the case $\mathcal{L}(N;(2,2))$ mentioned above. If we now let the vector α in Theorem 1.1 be random and uniformly distributed over $\{-1,1\}^k$, then averaging (4) gives that for any $p > 1$, we have

$$T_p(\mathcal{L}(N;\mathbf{p})) \ge (\log N)^{\omega(1)}. \tag{5}$$

A celebrated result of Pisier [Pis80] (which bounds the K-convexity constant of finite-dimensional Banach spaces; see also [Mau03, Lemma 7, Theorem 13]) implies that there exists an absolute constant $c \in (0,\infty)$ such that for any finite-dimensional Banach space X and any $p, r \in (1,\infty)$ such that $1/p + 1/r = 1$, we have

$$C_r(X^*) \ge \frac{c\, T_p(X)}{1 + \log \dim(X)}. \tag{6}$$

Combining (5) and (6) with $\log \dim(\mathcal{L}(N;\mathbf{p})) = q \log N$, we thus obtain (3).

Open questions. For proving LDC lower bounds it is of interest to know what is the largest d such that $\mathcal{L}(N;\mathbf{p})$ contains an $O(1)$-isomorphic copy of ℓ_1^d when $q \ge 3$. For this purpose it in fact suffices to restrict to copies of ℓ_1^d spanned by the type of forms appearing in the proof of Theorem 1.1 below, which may be seen as lying in a generalization of the Birkhoff polytope (the set of doubly stochastic matrices). Another question is if there is a converse to Theorem 1.1: Can a copy of ℓ_1^k inside $\mathcal{L}(N;\mathbf{p})$ be turned into an LDC-like object?

Outline. In Section 2 we set a few notational conventions and gather some basic facts of normed spaces and Fourier analysis over the boolean hypercube. In Section 3 we prove the main result, Theorem 1.1. In Section 4 we give alternative proofs for lower bounds on 2-query LDCs. In the Appendix we combine similar ideas with a reduction inspired by [KdW04] to give alternative proofs for lower bounds on LDCs with more queries.

Acknowledgements. I thank Oded Regev for inspiring conversations and useful comments on an earlier version of this manuscript, and I thank Mark Kim for helpful discussions early on.

2 Preliminaries

Notation. For a positive integer n denote $[n] = \{1,\dots,n\}$. Denote by $B_{n,d} \subseteq \{0,1\}^n$ the Hamming ball of radius d around the origin. For a finite set S denote by $\mathbb{E}_{x\in S}$ the expectation with

respect to a uniformly distributed random element x in S. For a probability distribution μ denote by $\mathbb{E}_{x\sim\mu}$ the expectation with respect to a random variable with distribution μ. For sets Γ and S, a positive integer r, and a pair of ordered tuples $z \in \Gamma^S$ and $\mathbf{S} = (s_1, \ldots, s_r) \in S^r$, we denote by $z_{\mathbf{S}} \in \Gamma^r$ the ordered tuple $(z_{s_1}, \ldots, z_{s_r})$. With some abuse of notation we will apply set operations to ordered tuples: for $S, \mathbf{S}$ as above write $s \in \mathbf{S}$ if $s = s_j$ for some $j \in [r]$ and write $S \cap \mathbf{S}$ for the set $\{s \in S : s_j = s \text{ for some } j \in [r]\}$.

Norms and spaces. For $1 \le p < \infty$, the ℓ_p-norm of a vector $u \in \mathbf{R}^N$ is defined by

$$\|u\|_{\ell_p} = \left(\sum_{i=1}^{N} |u_i|^p\right)^{1/p}.$$

Moreover, $\|u\|_{\ell_\infty} = \max_{i\in[N]}\{|u_i|\}$. For $p \in [1, \infty]$ denote by ℓ_p^N the Banach space $(\mathbf{R}^N, \| \ \|_{\ell_p})$. For a finite set S we denote by $\ell_q(S) = (\mathbf{R}^S, \| \ \|_{\ell_p})$ the space of vectors indexed by S endowed with the ℓ_p norm.

Fourier analysis over the boolean hypercube. For a positive integer n, the n-dimensional boolean hypercube, denoted H_n, is the group formed by the set $\{0,1\}^n$ endowed with entry-wise addition modulo 2. The character group of H_n is formed by the functions $\chi_u : H_n \to \mathbf{R}$ given by $\chi_u(x) = (-1)^{u\cdot x}$ for each $u \in \{0,1\}^n$, where $u \cdot x = u_1x_1 + \cdots + u_nx_n$. A character χ_u has *degree* d if the string u has Hamming weight d. The character functions form a complete orthonormal basis for the Hilbert space of functions $f : H_n \to \mathbf{R}$ endowed with the inner product

$$\langle f, g\rangle = \mathbb{E}_{x\in H_n}\big[f(x)g(x)\big]. \tag{7}$$

The Fourier transform $\widehat{f} : H_n \to \mathbf{R}$ of a function $f : H_n \to \mathbf{R}$ is given by $\widehat{f}(u) = \langle f, \chi_u\rangle$. A function f has *degree* d if its Fourier transform is supported by $B_{n,d}$. Orthogonality of the character functions with respect to the inner product (7) easily gives the *Fourier inversion formula*

$$f(x) = \sum_{u\in H_n} \widehat{f}(u)\chi_u(x)$$

and *Parseval's identity*

$$\sum_{u\in H_n} \widehat{f}(u)^2 = \mathbb{E}_{x\in H_n}\big[f(x)^2\big].$$

It also follows easily from the above that a function f depends only on a subset $S \subseteq [n]$ of its variables if and only if $\widehat{f}(u) = 0$ for every $u \in H_n$ such that $u_j = 1$ for some $j \notin S$. In particular, such a function has degree $|S|$.

The above extends to Cartesian products of H_n, since for positive integers q, we have $H_n^q \cong H_{qn}$. The characters of H_n^q are given by $\chi_{\mathbf{u}} = \chi_{u[1]} \cdots \chi_{u[q]}$ for every $\mathbf{u} = (u[1], \ldots, u[q]) \in H_n^q$ and a function $f : H_n^q \to \mathbf{R}$ has *degree* d if its Fourier transform is supported by $(B_{n,d})^q$.

3 Copies of ℓ_1^k from LDCs

In this section we prove Theorem 1.1. In the restatement below, we use the fact that at a loss of at most a factor of 2 in $|\Gamma|$, we may assume that $\Gamma = H_n$ for some positive integer n. Also, for convenience later on, we will switch the message alphabet from $\{0,1\}$ to $\{-1,1\}$.

Theorem 3.1. *Let $\delta, \varepsilon \in (0, 1/2]$ and k, N, q, m, n be positive integers such that $q \geq 2$ and $m \leq n$. Suppose that there exists a (q, δ, ε)-LDC from $\{-1, 1\}^k$ to H_n^N that has a decoder that uses at most m predetermined bits of each queried codeword symbol. Then, for any vector $\mathbf{p} \in (1, \infty)^q$ such that $1/p_1 + \cdots + 1/p_q \leq 1$, integer $N' \geq \binom{n}{\leq m} N$, and real number*

$$K \geq \frac{q\binom{n}{\leq m}^{(q+2)/2}}{2\delta\varepsilon},$$

the space $\mathcal{L}(N'; \mathbf{p})$ contains a K-isomorphic copy of ℓ_1^k.

For the rest of this section, let $k, N, q, m, n, \delta, \varepsilon, \mathbf{p}$ be as in Theorem 3.1.

3.1 Smooth decoding

The proof of Theorem 3.1 relies on a variant of a result of [KT00]. Qualitatively the result says that an LDC allows us to retrieve any message bit with high probability from an uncorrupted codeword by sampling q-tuples of codeword coordinates from a "smooth" distribution, in which the marginal distribution over single coordinates is roughly uniform.

Lemma 3.2. *Let $C : \{-1, 1\}^k \to H_n^N$ be a (q, δ, ε)-LDC. Then, for each $i \in [k]$ there exists a probability distribution μ_i over $[N]^q$ and for each $\mathbf{S} \in [N]^q$ there exists a function $f_{\mathbf{S}}^i : H_n^q \to [-1, 1]$ such that:*

- *For every $x \in \{-1, 1\}^k$, we have $x_i \, \mathbb{E}_{\mathbf{S} \sim \mu_i}\big[f_{\mathbf{S}}^i\big(C(x)_{\mathbf{S}}\big)\big] \geq 2\varepsilon$.*
- *For every $s \in [N]$, we have $\Pr_{\mathbf{S} \sim \mu_i}[s \in \mathbf{S}] \leq 2q/(\delta N)$.*

Moreover, if the LDC decoder's output depends on at most m predetermined bits of each queried codeword symbol, then $f_{\mathbf{S}}^i$ has degree at most m.

Proof. Fix an $i \in [k]$. Let ν_i be a probability distribution over sets $S \subseteq [N]$ of cardinality at most q and for every set S in the support of ν_i let ϕ_S be a map from H_n^S to the set of $\{-1, 1\}$-valued random variables. Suppose that upon receiving the index i and a string $y \in H_n^N$, the decoder samples a set S from ν_i and outputs the random variable $\phi_S(y_S)$.

Let S be a random set with distribution ν_i. Let $B \subseteq [N]$ be the set of *bad* coordinates $s \in [N]$ satisfying $\Pr[s \in S] \geq q/(\delta N)$. Since ν_i is supported only on sets of size at most q, it follows that $|B| \leq \delta N$. Let $\widetilde{\mathbf{S}} = (\tilde{s}_s)_{s \in S}$ be the random sequence such that for each bad coordinate $s \in S$, the entry $\tilde{s}_s$ is independent and uniformly distributed over $[N]$ and for the other coordinates, we have $\tilde{s}_s = s$. We claim that, similar to the second item in the lemma, for every $s \in [N]$, we have

$$\Pr[s \in \widetilde{\mathbf{S}}] \leq \frac{2q}{\delta N}. \tag{8}$$

For $s \in [N] \setminus B$, the probability in (8) is at most $\Pr[s \in S] \leq q/(\delta N)$ plus the probability that s appears in a bad coordinate of $\widetilde{\mathbf{S}}$. By independence, the latter probability is at most q/N, showing (8) for $[N] \setminus B$. Bad elements $s \in B$ only appear at bad coordinates of $\widetilde{\mathbf{S}}$. By independence, such elements therefore appear with probability at most q/N, giving the claim.

Let $\widetilde{S} = \{\tilde{s}_s : s \in S\}$ be the random set of distinct entries of $\widetilde{\mathbf{S}}$ and let $\tilde{\nu}_i$ be the distribution of $\widetilde{S}$. Let $T \subseteq [N]$ be a set of cardinality at most q. Recall from our notational convention (see

Section 2) that for a vector $z \in H_n^T$, conditioned on the event $\widetilde{S} = T$, the vector $z_{\widetilde{\mathbf{S}}} = (z_{\tilde{s}_s})_{s \in S}$ is well-defined as one lying in H_n^S. This allows us to define a function $g_T^i : H_n^T \to [-1, 1]$ by

$$g_T^i(z) = \mathbb{E}\big[\phi_S^i(z_{\widetilde{\mathbf{S}}}) \mid \widetilde{S} = T\big], \tag{9}$$

where the expectation is taken over the set S, the sequence $\widetilde{\mathbf{S}}$, and the random value in $\{-1, 1\}$ assumed by the function ϕ_S^i. We show that these functions g_T^i satisfy an inequality similar to the first item in the lemma, namely, we show that for every $x \in \{-1, 1\}^k$ and random set T with distribution $\tilde{\nu}_i$, we have

$$x_i \, \mathbb{E}_{T \sim \tilde{\nu}_i}\big[g_T^i\big(C(x)_T\big)\big] \geq 2\varepsilon. \tag{10}$$

To show this, consider the random string $y \in H_n^N$ where for every $s \in [N] \setminus B$, we have $y_s = C(x)_s$ and for every $s \in B$, we set $y_s = C(x)_{t_s}$ where t_s is independent and uniformly distributed over $[N]$. As such, y is thus a random "corrupted" version of $C(x)$ in which at most $|B| \leq \delta N$ entries are replaced with other entries of the codeword. We claim that the random sequences $C(x)_{\widetilde{\mathbf{S}}}$ and y_S have the same distribution. Indeed, observe that we get the first sequence if we sample S and then corrupt the sequence $C(x)_S$ by replacing its entries at bad coordinates s by the random value $C(x)_{t_s}$ for t_s as above. The second sequence y_S corresponds to doing things in reverse order: first corrupt $C(x)$, giving y, and then sample S. The claim follows since the corrupted entries in S are independent of S. It follows that the random variables $\phi_S^i\big(C(x)_{\widetilde{\mathbf{S}}}\big)$ and $\phi_S^i(y_S)$ also have the same distribution and, since y differs from $C(x)$ in at most δN coordinates,

$$\Pr\big[\phi_S^i\big(C(x)_{\widetilde{\mathbf{S}}}\big) = x_i\big] = \Pr\big[\phi_S^i(y_S) = x_i\big] \geq \frac{1}{2} + \varepsilon. \tag{11}$$

Hence, since $\widetilde{S}$ has the distribution $\tilde{\nu}_i$, we have

$$\begin{aligned}
x_i \, \mathbb{E}_{T \sim \tilde{\nu}_i}\big[g_T^i\big(C(x)_T\big)\big] &\overset{(9)}{=} x_i \, \mathbb{E}_{T \sim \tilde{\nu}_i}\Big[\mathbb{E}\big[\phi_S^i\big(C(x)_{\widetilde{\mathbf{S}}}\big) \mid \widetilde{S} = T\big]\Big] \\
&= x_i \, \mathbb{E}\big[\phi_S^i\big(C(x)_{\widetilde{\mathbf{S}}}\big)\big] \\
&\overset{(11)}{=} x_i \, \mathbb{E}\big[\phi_S^i(y_S)\big] \\
&\geq 2\varepsilon,
\end{aligned}$$

where the first inner expectation and the second and third expectations are taken over the set S, the sequence $\widetilde{\mathbf{S}}$, the set $\widetilde{S}$, and the random value of the function ϕ_S^i. This shows (10).

Define the probability distribution $\mu_i : [N]^q \to [0, 1]$ as follows. For a set $T \subseteq [N]$ with cardinality at most q, let $\mathcal{F}(T) \subseteq T^q$ be the family of ordered sequences that contain each element of T at least once. For each $\mathbf{T} \in \mathcal{F}(T)$ set $\mu_i(\mathbf{T}) = \tilde{\nu}_i(T)/|\mathcal{F}(T)|$. Then, by (8) we have

$$\Pr_{\mathbf{T} \sim \mu_i}[s \in \mathbf{T}] = \Pr_{T \sim \tilde{\nu}_i}[s \in T] \leq \frac{2q}{\delta N}$$

for each $s \in [N]$. For each set T in the support of $\tilde{\nu}_i$ and every $\mathbf{T} \in \mathcal{F}(T)$, there exists a function $f_{\mathbf{T}}^i : H_n^q \to [-1, 1]$ such that $f_{\mathbf{T}}^i\big(z_{\mathbf{T}}\big) = g_T^i(z)$ holds for each $z \in H_n^T$ (as $z_{\mathbf{T}}$ and z have entries from the same set). Pick one such function arbitrarily. For all remaining $\mathbf{T} \in [N]^q$ let $f_{\mathbf{T}}^i$ be identically zero. By (10), these functions satisfy the first item of the lemma.

Finally, observe that if the decoder's output depends on at most m predetermined bits of each queried codeword symbol, then for each set T in the support of ν_i, the function $h_T^i : H_n^T \to [-1,1]$ defined by $h_T^i(z) = \mathbb{E}[\phi_T^i(z)]$, where the expectation is taken over the randomness in ϕ_T^i, has degree at most m. Since the functions g_T^i in (9) are linear combinations of these h_T^i, they also have degree at most m. It follows that the functions $f_{\mathbf{T}}^i$ can be chosen to satisfy the same. □

3.2 Norms of some forms

The proof of Theorem 1.1 uses the functions $f_{\mathbf{S}}^i$ and distributions μ_i of Lemma 3.2 to construct a basis $A_1, \dots, A_k \in \mathcal{L}(N'; \mathbf{p})$ for a copy of ℓ_1^k as in (4). Viewed as a q-tensor, the form A_i will consist of blocks, one block for each q-tuple $\mathbf{S} \in [N]^q$, and the entries of each block will contain the Fourier coefficients of the function $f_{\mathbf{S}}^i$ scaled by the probability $\mu_i(\mathbf{S})$. We use the following facts to show that these forms have the desired properties.

Proposition 3.3. *Let $f : H_n^q \to [-1,1]$ be a function of degree at most m. Define the q-linear form F on $\mathbf{R}^{B_{n,m}}$ by*

$$F(\mathbf{z}) = \sum_{\mathbf{u} \in B_{n,m}^q} \widehat{f}(\mathbf{u})\, z[1]_{u[1]} \cdots z[q]_{u[q]}$$

for $\mathbf{z} = (z[1], \dots, z[q]) \in \mathbf{R}^{B_{n,m}} \times \cdots \times \mathbf{R}^{B_{n,m}}$. Then, $\|F\|_{\mathbf{p}} \le |B_{n,m}|^{q/2}$.

Proof. Hölder's inequality implies that a vector in the unit ball of ℓ_p^t has ℓ_2-norm at most $t^{1/2-1/p}$. Hence, by the Cauchy-Schwarz inequality and Parseval's identity,

$$\begin{aligned}
|F(\mathbf{z})| &\le \left(\sum_{\mathbf{u} \in B_{n,m}^q} \widehat{f}(\mathbf{u})^2 \right)^{1/2} \prod_{j=1}^{q} \|z[j]\|_{\ell_2} \\
&= \prod_{j=1}^{q} \|z[j]\|_{\ell_2} \\
&\le \prod_{j=1}^{q} |B_{n,m}|^{1/2-1/p_j} \|z[j]\|_{\ell_{p_j}} \\
&\le |B_{n,m}|^{q/2} \prod_{j=1}^{q} \|z[j]\|_{\ell_{p_j}}.
\end{aligned}$$

□

We use a generalization of a doubly-stochastic matrix. Let $\mathbf{1} \in \mathbf{R}^N$ denote the all-ones vector.

Definition 3.4 (Plane sub-stochastic form)**.** *A q-linear form A on $\mathbf{R}^N$ is* plane sub-stochastic *if the tensor $T = \big(A(e_{s_1}, \dots, e_{s_q})\big)_{\mathbf{S} \in [N]^q}$ is nonnegative and for every $s \in [N]$, we have*

$$\begin{aligned}
A(e_s, \mathbf{1}, \mathbf{1}, \dots, \mathbf{1}) &\le 1 \\
A(\mathbf{1}, e_s, \mathbf{1}, \dots, \mathbf{1}) &\le 1 \\
&\vdots \\
A(\mathbf{1}, \mathbf{1}, \dots, \mathbf{1}, e_s) &\le 1.
\end{aligned} \tag{12}$$

Remark 1. *The above definition gives the Birkhoff polytope when we set $q = 2$ and we change the inequalities in (12) to equalities. Recall that the Birkhoff–von Neumann Theorem states that the Birkhoff polytope is the convex hull of the set of $N \times N$ permutation matrices. Interestingly, Linial and Luria [LL14] showed that for $q \geq 3$, the polytope of q-linear "plane stochastic forms" corresponding to equalities in (12) is* not *contained in the convex hull of the set of "permutation tensors" defined as 0/1 tensors satisfying equality in (12).*

Proposition 3.5. *If $A \in \mathcal{L}(N; \mathbf{p})$ is plane sub-stochastic, then $\|A\|_{\mathbf{p}} \leq 1$.*

The proof uses the following result of Carlen, Loss, and Lieb [CLL06].

Theorem 3.6 (Multi-linear Riesz–Thorin Interpolation Theorem)**.** *Let A be a q-linear form on $\mathbf{R}^N$. Let $\psi : [0,1]^q \to \mathbf{R}_+$ be the function defined by $\psi(1/r_1, \ldots, 1/r_q) = \|A\|_{\mathbf{r}}$, for any $\mathbf{r} \in [1,\infty]^q$. Then, $\ln(\psi)$ is a convex function on $[0,1]^q$.*

Proof of Proposition 3.5. We first show that $\|A\|_{\mathbf{r}} \leq 1$ for any $\mathbf{r}$ consisting of one 1 entry and all the others ∞. Indeed, by Hölder's inequality and the assumption that $|A|$ is plane sub-stochastic,

$$\begin{aligned} |A(z[1], \ldots, z[q])| &\leq \sum_{\mathbf{S} \in [N]^q} |A(e_{s_1}, \ldots e_{s_q}) z[1]_{s_1} \cdots z[q]_{s_q}| \\ &\leq \Big(\prod_{j=1}^{q-1} \|z[j]\|_{\ell_\infty} \Big) \sum_{s=1}^{N} A(\mathbf{1}, \ldots, \mathbf{1}, s) |z[q]_s| \\ &\leq \Big(\prod_{j=1}^{q-1} \|z[j]\|_{\ell_\infty} \Big) \|z[q]\|_{\ell_1}. \end{aligned}$$

Hence, if $\mathbf{r} = (\infty, \ldots, \infty, 1)$, we have $\|A\|_{\mathbf{r}} \leq 1$. The cases for the other positions of the 1-entry are proved in the same way. Since for these choices of $\mathbf{r}$, the vectors $(1/r_1, \ldots, 1/r_q)$ are the q standard basis vectors, the vector $\mathbf{p}$ lies in their convex hull, and the result follows from Theorem 3.6. □

3.3 Proof of the main result

With this, the proof of Theorem 3.1 is straightforward.

Proof of Theorem 3.1. Let $C : \{-1,1\}^k \to H_n^N$ be a (q, δ, ε)-LDC as in the theorem. For each index $i \in [k]$ and $\mathbf{S} = (s_1, \ldots, s_q) \in [N]^q$ let μ_i and $f_{\mathbf{S}}^i$ be a distribution and function as in Lemma 3.2. Note that the Fourier transform of each $f_{\mathbf{S}}^i$ is supported by the Cartesian product of Hamming balls $(B_{n,m})^q$. Define a q-linear form $F_{\mathbf{S}}^i$ on $\mathbf{R}^{B_{n,m}}$ based on the Fourier coefficients of the function $f_{\mathbf{S}}^i$ as in Proposition 3.3. For a vector $z \in \mathbf{R}^{[N] \times B_{n,m}}$ and $s \in [N]$ write z_s for the projection of z onto the coordinates (s,u) with $u \in B_{n,m}$, that is, $z_s = (z_{(s,u)})_{u \in B_{n,m}} \in \mathbf{R}^{B_{n,m}}$. For a tuple $\mathbf{z} = (z[1], \ldots, z[q])$ with each $z[j] \in \mathbf{R}^{[N] \times B_{n,m}}$, write $\mathbf{z}_{\mathbf{S}} = (z[1]_{s_1}, \ldots, z[q]_{s_q})$. Let A_i be the q-linear form on $\mathbf{R}^{[N] \times B_{n,m}}$ defined by

$$A_i(\mathbf{z}) = \mathbb{E}_{\mathbf{S} \sim \mu_i} \big[F_{\mathbf{S}}^i(\mathbf{z}_{\mathbf{S}}) \big] = \mathbb{E}_{\mathbf{S} \sim \mu_i} \Big[\sum_{\mathbf{u} \in (B_{n,m})^q} \widehat{f_{\mathbf{S}}^i}(\mathbf{u}) \, z[1]_{(s_1, u[1])} \cdots z[q]_{(s_q, u[q])} \Big]. \tag{13}$$

Fix a vector $\alpha \in \mathbf{R}^k$, let $x = (\mathrm{sign}(\alpha_i))_{i=1}^k$ and let $y = C(x)$ be the codeword in H_n^N corresponding to the message x. Define the sign vector $\widehat{y} = (\chi_u(y_s))_{s\in[N],u\in B_{n,m}}$. Let $\widehat{\mathbf{y}} = (\widehat{y},\dots,\widehat{y})$ (q times). By the Fourier Inversion Formula,

$$F^i_{\mathbf{S}}(\widehat{\mathbf{y}}_{\mathbf{S}}) \;=\; \sum_{\mathbf{u}\in(B_{n,m})^q} \widehat{f^i_{\mathbf{S}}}(\mathbf{u})\,\chi_{\mathbf{u}}(y_{\mathbf{S}}) \;=\; f^i_{\mathbf{S}}(y_{\mathbf{S}}) \;=\; f^i_{\mathbf{S}}(C(x)_{\mathbf{S}}). \tag{14}$$

Combining (14) with the first property of the $f^i_{\mathbf{S}}$ in Lemma 3.2 then gives

$$\alpha_i\,\mathbb{E}_{\mathbf{S}\sim\mu_i}\big[F^i_{\mathbf{S}}(\widehat{\mathbf{y}}_{\mathbf{S}})\big] \ge 2|\alpha_i|\varepsilon. \tag{15}$$

Since $\widehat{y}$ is a sign vector of dimension $N|B_{n,m}|$, it has ℓ_p-norm $(N|B_{n,m}|)^{1/p}$. Normalizing accordingly and using q-linearity of the A_i, we get

$$\begin{aligned}\Big\|\sum_{i=1}^k \alpha_i A_i\Big\|_{\mathbf{p}} &\ge \frac{1}{(N|B_{n,m}|)^{1/p_1+\cdots+1/p_q}}\left(\sum_{i=1}^k \alpha_i A_i\right)(\widehat{\mathbf{y}})\\ &\overset{(13)}{\ge} \frac{1}{N|B_{n,m}|}\sum_{i=1}^k \alpha_i \mathbb{E}_{\mathbf{S}\sim\mu_i}\big[F^i_{\mathbf{S}}(\widehat{\mathbf{y}}_{\mathbf{S}})\big]\\ &\overset{(14),(15)}{\ge} \frac{2\varepsilon}{N|B_{n,m}|}\|\alpha\|_{\ell_1}.\end{aligned} \tag{16}$$

Next, we bound the norms of the forms A_i themselves. Let $\mathbf{z} = (z[1],\dots,z[q])$ be a q-tuple of nonzero vectors in $\mathbf{R}^{[N]\times B_{n,m}}$. Recall from Proposition 3.3 that each $F^i_{\mathbf{S}}$ has norm at most $|B_{n,m}|^{q/2}$. This implies

$$\begin{aligned}\big|A_i(\mathbf{z})\big| &\overset{(13)}{\le} \mathbb{E}_{\mathbf{S}\sim\mu_i}\Big[\big|F^i_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}})\big|\Big]\\ &\le |B_{n,m}|^{q/2}\,\mathbb{E}_{\mathbf{S}\sim\mu_i}\Big[\|z[1]_{s_1}\|_{\ell_{p_1}}\cdots\|z[q]_{s_q}\|_{\ell_{p_q}}\Big].\end{aligned} \tag{17}$$

To bound the above expectation define the q-tuple $\mathbf{a} = (a[1],\dots,a[q])$ of (nonnegative) vectors

$$a[j] = \Big(\|z[j]_s\|_{\ell_{p_j}}\Big)_{s=1}^N, \quad j\in[q]. \tag{18}$$

Define the q-linear form M on $\mathbf{R}^N$ given by $M(\mathbf{b}) = \mathbb{E}_{\mathbf{S}\sim\mu_i}[b[1]_{s_1}\cdots b[q]_{s_q}]$. Then, the expectation in (17) equals $M(\mathbf{a})$. The form M is clearly nonnegative and by the second item in Lemma 3.2, the scaled version $(\delta N/q)M$ is plane sub-stochastic since for each $t\in[N]$ and $j\in[q]$, we have

$$M[\underbrace{\mathbf{1},\dots,\mathbf{1}}_{1,\dots,j-1},\underbrace{e_t}_{j},\underbrace{\mathbf{1},\dots,\mathbf{1}}_{j+1,\dots,q}] = \mathbb{E}_{\mathbf{S}\sim\mu_i}\big[(e_t)_{s_j}\big] = \Pr_{\mathbf{S}\sim\mu_i}[s_j = t] \le \frac{q}{\delta N}.$$

By Proposition 3.5, the form M therefore has norm at most $\|M\|_{\mathbf{p}} \le q/(\delta N)$. Since each $a[j]$ as in (18) has norm $\|a[j]\|_{\ell_{p_j}} = \|z[j]\|_{\ell_{p_j}}$, we conclude that each A_i has norm $\|A_i\|_{\mathbf{p}} \le q|B_{n,m}|^{q/2}/(\delta N)$.

Hence, for any vector $\alpha\in\mathbf{R}^k$, by (16) and the triangle inequality,

$$\frac{2\varepsilon}{N|B_{n,m}|}\|\alpha\|_{\ell_1} \le \Big\|\sum_{i=1}^k \alpha_i A_i\Big\|_{\mathbf{p}} \le \frac{q|B_{n,m}|^{q/2}}{\delta N}\|\alpha\|_{\ell_1}.$$

Scaling the A_i by $N|B_{n,m}|/(2\varepsilon)$ then gives the copy of ℓ_1^k as desired. □

4 Lower bounds on 2-query LDCs

In this section we use Theorem 1.1 to prove the 2-query LDC lower bounds mentioned in the Introduction (up-to slightly poorer dependence on δ and $|\Gamma|$). The key is the following bound on the dimension k for which $\mathcal{L}(N;(2,2))$ can accommodate a copy of ℓ_1^k with distortion K. The bound is surely well-known, but it does not appear to be published in the form below.

Lemma 4.1. *There exists an absolute constant $C \in (0,\infty)$ such that the following holds. Suppose that for $K < \infty$ the space $\mathcal{L}(N;(2,2))$ contains a K-isomorphic copy of ℓ_1^k. Then $k \le CK^2 \log(2N)$.*

Lemma 4.1 follows easily from a random-matrix inequality belonging to a family of "non-commutative Khintchine inequalities" due to Tomczak-Jaegermann [TJ74] (not to be confused with the stronger non-commutative Khintchine inequalities of Lust-Piquard and Pisier [LPP91]). Recall that the Schatten-∞ norm $\|A\|_{S_\infty}$ of a matrix $A \in \mathbf{R}^{N\times N}$ is the supremum of $|u^{\mathsf{T}} A v|/\|u\|_2\|v\|_2$ over nonzero vectors $u, v \in \mathbf{R}^N$.

Theorem 4.2 (Tomczak-Jaegermann)**.** *There exists an absolute constant $C \in (0,\infty)$ such that the following holds. Let N and k be positive integers, let $A_1,\dots,A_k \in \mathbf{R}^{N\times N}$, and let $\epsilon_1,\dots,\epsilon_k$ be independent uniformly distributed $\{-1,1\}$-valued random variables. Then,*

$$\mathbb{E}\Big[\Big\|\sum_{i=1}^k \epsilon_i A_i\Big\|_{S_\infty}\Big] \le C\sqrt{\log(2N)}\,\Big(\sum_{i=1}^k \|A_i\|_{S_\infty}^2\Big)^{1/2}. \tag{19}$$

Remark 2. *To extract the above from [TJ74, Theorem 3.1] we used the standard and easy fact that for $p = \log N$, the Schatten-p norm of a matrix $A \in \mathbf{R}^{N\times N}$, defined as the ℓ_p-norm of the vector of singular values of A, satisfies $\|A\|_{S_\infty} \le \|A\|_p \le C\|A\|_{S_\infty}$ for some absolute constant $C \in [1,\infty)$.*

Remark 3. *Similar (stronger) estimates were proved in [LPP91, Buc05, Oli10, Tro12].*

Proof of Lemma 4.1. Identify the space $\mathcal{L}(N;(2,2))$ with $(\mathbf{R}^{N\times N}, \|\ \|_{S_\infty})$. Let $A_1,\dots,A_k \in \mathbf{R}^{N\times N}$ be matrices such that (4) holds (with $X = S_\infty$). Setting the vector α in (4) to be a standard basis vector we see that $\|A_i\|_{S_\infty} \le K$ for each $i \in [k]$. Hence, by (4) and Theorem 4.2,

$$k \le \mathbb{E}\Big[\Big\|\sum_{i=1}^k \epsilon_i A_i\Big\|_{S_\infty}\Big] \le CK\sqrt{k\log(2N)}. \qquad \square$$

Theorem 1.1 asserts that a $(2,\delta,\varepsilon)$-LDC from $\{0,1\}^k$ to Γ^N gives a $4|\Gamma|^2/(\delta\varepsilon)$-isomorphic copy of ℓ_1^k in the space $\mathcal{L}(2|\Gamma|N;(2,2))$. Combining this with Lemma 4.1 immediately gives the following exponential lower bounds on binary 2-query LDCs.

Corollary 4.3. *Any binary $(2,\delta,\varepsilon)$-LDC satisfies $N \ge 2^{\Omega(\delta^2\varepsilon^2 k)}$.*

For LDCs over larger alphabets we obtain the following bound.

Corollary 4.4. *Any $(2,\delta,\varepsilon)$-LDC with $\Gamma = H_n$ and a decoder that uses at most m out of n predetermined bits of each queried codeword symbol satisfies*

$$\binom{n}{\le m}^3 \left(\log N + \log\binom{n}{\le m}\right) \ge \Omega(\delta^2\varepsilon^2 k). \tag{20}$$

Remark 4. *A more careful analysis in Section 3 for the case $1/p_1 + \cdots + 1/p_q = 1$ allows one to replace the third power in the left-hand side of (20) with a square.*

References

[BARdW08] Avraham Ben-Aroya, Oded Regev, and Ronald de Wolf. A hypercontractive inequality for matrix-valued functions with applications to quantum computing and LDCs. In *Proceedings of the 49th Annual IEEE Annual Symposium on Foundations of Computer Science (FOCS 2008)*, pages 477–486, 2008.

[BCL94] Keith Ball, Eric Carlen, and Elliot Lieb. Sharp uniform convexity and smoothness inequalities for trace norms. *Invent. Math.*, 115:463–482, 1994.

[BDHS14] Jop Briët, Zeev Dvir, Guangda Hu, and Shubhangi Saraf. Lower bounds for approximate LDCs. In *Automata, Languages, and Programming*, volume 8572 of *Lecture Notes in Computer Science*, pages 259–270. Springer Berlin Heidelberg, 2014. Full version available at arXiv:1402.6952.

[BDYW11] Boaz Barak, Zeev Dvir, Amir Yehudayoff, and Avi Wigderson. Rank bounds for design matrices with applications to combinatorial geometry and locally correctable codes. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC 2011)*, pages 519–528. ACM, 2011.

[BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *Proceedings of the 23rd annual ACM symposium on Theory of computing (STOC 1991)*, pages 21–32. ACM, 1991.

[BNR12] Jop Briët, Assaf Naor, and Oded 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.

[Buc05] Artur Buchholz. Optimal constants in Khintchine type inequalities for Fermions, Rademachers and q-Gaussian operators. *Bulletin of the Polish Academy of Sciences. Mathematics*, 53(3):315–321, 2005.

[CGKS98] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. *J. ACM*, 45:965–981, 1998.

[CGW13] Victor Chen, Elena Grigorescu, and Ronald de Wolf. Error-correcting data structures. *SIAM Journal on Computing*, 42(1):84–111, 2013.

[CLL06] Eric Carlen, Elliott H. Lieb, and Michael Loss. An inequality of Hadamard type for permanents. *Methods Appl. Anal.*, 13(1):1–17, 2006.

[DFS03] Joe Diestel, Jan Fourie, and Johan Swart. The projective tensor product. I. In *Trends in Banach spaces and operator theory (Memphis, TN, 2001)*, volume 321 of *Contemp. Math.*, pages 37–65. Amer. Math. Soc., Providence, RI, 2003.

[DG14] Zeev Dvir and Sivakanth Gopi. 2-Server PIR with sub-polynomial communication. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC 2015)*, 2014. (To appear) Available at arXiv:1407.6692.

[DGY10] Zeev Dvir, Parakshit Gopalan, and Sergey Yekhanin. Matching vector codes. In *Proceedings of the 51st Annual IEEE Annual Symposium on Foundations of Computer Science (FOCS 2010)*, 2010.

[DS07] Zeev Dvir and Amir Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM Journal on Computing*, 36(5):1404–1434, 2007.

[DSW14] Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. Breaking the quadratic barrier for 3-LCC's over the Reals. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 784–793. ACM, 2014.

[Dvi10] Zeev Dvir. On matrix rigidity and locally self-correctable codes. In *Twenty-Fifth Annual IEEE Conference on Computational Complexity (CCC 2010)*, pages 291–298, 2010.

[Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the 41st annual ACM symposium on Theory of computing (STOC 2009)*, pages 39–44, 2009.

[GKST06] Oded Goldreich, Howard Karloff, Leonard J Schulman, and Luca Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. *Computational Complexity*, 15(3):263–296, 2006. Preliminary version appeared in CCC'02.

[Gro00] Vince Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit Ramsey graphs. *Combinatorica*, 20(1):71–86, 2000.

[HOW14] Brett Hemenway, Rafail Ostrovsky, and Mary Wootters. Local correctability of expander codes. *Information and Computation*, 2014.

[KdW04] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *J. of Computer and System Sciences*, 69:395–420, 2004. Preliminary version appeared in STOC'03.

[KMRZS15] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High rate locally-correctable and locally-testable codes with sub-polynomial query complexity. *arXiv preprint arXiv:1504.05653*, 2015.

[KSY11] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC 2011)*, pages 167–176. ACM, 2011.

[KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC 2000)*, pages 80–86. ACM Press, 2000.

[LL14] Nathan Linial and Zur Luria. On the vertices of the d-dimensional Birkhoff polytope. *Discrete Comput. Geom.*, 51(1):161–170, 2014.

[LPP91] Françoise Lust-Piquard and Gilles Pisier. Non commutative Khintchine and Paley inequalities. *Arkiv för Matematik*, 29(1):241–260, 1991.

[Mau03] Bernard Maurey. Type, cotype and K-convexity. In *Handbook of the geometry of Banach spaces, Vol. 2*, pages 1299–1332. North-Holland, Amsterdam, 2003.

[MP73] Bernard Maurey and Gilles Pisier. Caractérisation d'une classe d'espaces de Banach par des propriétés de séries aléatoires vectorielles. *C. R. Acad. Sci. Paris Sér A*, 277:687—690, 1973.

[MU05] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.

[MV03] Shahar Mendelson and Roman Vershynin. Entropy and the combinatorial dimension. *Invent. Math.*, 152(1):37–55, 2003.

[Oli10] Roberto Imbuzeiro Oliveira. Sums of random Hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab*, 15(203-212):26, 2010.

[Pis73] Gilles Pisier. Sur les espaces de banach qui ne contiennent pas uniformément de ℓ_n^1. *C. R. Acad. Sci. Paris Sér A*, 277:991–994, 1973.

[Pis80] Gilles Pisier. Un théorème sur les opérateurs linéaires entre espaces de Banach qui se factorisent par un espace de Hilbert. *Ann. Sci. École Norm. Sup. (4)*, 13(1):23–43, 1980.

[Pis99] Gilles Pisier. *The volume of convex bodies and Banach space geometry*. Cambridge University Press, 1999.

[Pis12] Gilles Pisier. 15th workshop on non-commutative harmonic analysis, Będlewo, Poland, 2012.

[Rya02] Raymond A. Ryan. *Introduction to Tensor Products of Banach Spaces*. Springer Monographs in Mathematics. Springer, London, 2002.

[Sud92] Madhu Sudan. *Efficient checking of polynomials and proofs and the hardness of approximation problems*. PhD thesis, University of California at Berkeley, 1992.

[TJ74] Nicole Tomczak-Jaegermann. The moduli of smoothness and convexity and the Rademacher averages of trace classes S_p ($1 \leq p < \infty$). *Studia Math.*, 50:163–182, 1974.

[Tre04] Luca Trevisan. Some applications of coding theory in computational complexity. *Quaderni di Matematica*, pages 347–424, 2004.

[Tro12] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[WdW05] Stephanie Wehner and Ronald de Wolf. Improved lower bounds for locally decodable codes and private information retrieval. In *Proc. 32nd Intern. Colloquium on Automata, Languages and Programming (ICALP'05)*, volume 3580 of *Lecture Notes in Computer Science*, pages 1424–1436. Springer, 2005.

[Wol09] Ronald de Wolf. Error-correcting data structures. In *26th International Symposium on Theoretical Aspects of Computer Science STACS 2009*, pages 313–324, 2009.

[Woo07] David Woodruff. New lower bounds for general locally decodable codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(006), 2007.

[Yek07] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. In *Proceedings of the 39th annual ACM symposium on Theory of computing (STOC 2007)*, pages 266–274, 2007.

[Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3):139–255, 2012.

A Lower bounds on LDCs with more queries

Here we prove lower bounds on binary q-query LDCs for $q \geq 3$ using a method similar to the reductions to the two-query case used in [KdW04], but in the spirit of Theorem 1.1.

Theorem A.1. *Let $\delta, \varepsilon \in (0, 1/2]$ and let k, l, N, r be positive integers such that $r \geq 2$ and*

$$\frac{\delta\varepsilon}{(2r)^3} l^r \geq (2rN)^{r-1}. \tag{21}$$

Suppose that there exists a $(2r, \delta, \varepsilon)$-LDC from $\{-1,1\}^k$ to $\{-1,1\}^N$. Then, there exist matrices $A_1, \ldots, A_k \in \mathbf{R}^{N^l \times N^l}$ such that $\|A_i\|_{S_\infty} \leq 1$ for each $i \in [k]$, and for independent uniformly distributed $\{-1,1\}$-valued random variables $\epsilon_1, \ldots, \epsilon_k$, we have

$$\mathbb{E}\Big[\Big\|\sum_{i=1}^k \epsilon_i A_i\Big\|_{S_\infty}\Big] \geq \frac{\varepsilon k}{8^r}. \tag{22}$$

Elton's Theorem asserts that the above is sufficient to find a finite-dimensional copy of ℓ_1 (see Vershynin and Mendelson [MV03, Theorem 3] for the quantitatively optimal form stated below).

Theorem A.2 (Elton's Theorem)**.** *There exists a absolute constant $c > 0$ such that the following holds. Let X be a Banach space, let $A_1, \ldots, A_k$ be vectors in the unit ball of X, and let $\eta > 0$ be such that for independent uniformly distributed $\{-1,1\}$-valued random variables $\epsilon_1, \ldots, \epsilon_k$, we have*

$$\mathbb{E}\Big[\Big\|\sum_{i=1}^k \epsilon_i A_i\Big\|_X\Big] \geq \eta k.$$

Then, there exists a set $I \subseteq [k]$ of cardinality $|I| \geq c\eta^2 k$ such that for any $\alpha \in \mathbf{R}^I$, we have

$$c\eta\|\alpha\|_{\ell_1} \leq \Big\|\sum_{i\in I} \alpha_i A_i\Big\|_X \leq \|\alpha\|_{\ell_1}.$$

Combining Theorem A.1 with Elton's Theorem shows that, for a positive integer $r \geq 2$, a $(2r, \delta, \varepsilon)$-LDC gives a K-isomorphic copy of ℓ_1^d inside $\mathcal{L}(N^l; (2,2))$ for $K = \delta(\varepsilon/q)^2$ and $d \geq ck/K^2$. Through Lemma 4.1 this leads to a lower bound on $(2r)$-LDCs similar to the one stated in the Introduction. However, combining Theorem A.1 with Theorem 4.2 gives the following lower bound that has slightly better dependence on δ, ε, and r.

Corollary A.3. *For every integer $r \geq 2$ there exists a $c > 0$ such that the following holds. Suppose that for positive integers k and N and $\delta, \varepsilon \in (0, 1/2]$, there exists a $(2r, \delta, \varepsilon)$-LDC from $\{-1,1\}^k$ to $\{-1,1\}^N$. Then, $N \geq c(\delta\varepsilon^3 k/\log k)^{r/(r-1)}$.*

Proof. Let l be the smallest integer satisfying (21) and let $A_1, \ldots, A_k$ be matrices as in Theorem A.1. Theorem 4.2 then gives

$$\frac{\varepsilon k}{8^r} \leq \mathbb{E}\Big[\Big\|\sum_{i=1}^{k} \epsilon_i A_i\Big\|_{S_\infty}\Big] \leq \sqrt{2k\log(2eN^l)} \leq c\sqrt{\frac{kN^{(r-1)/r}\log N}{(\delta\varepsilon)^{1/r}}},$$

where $c < \infty$ is a constant depending on r only. Rearranging gives $N^{(r-1)/r}\log N \geq c'(\delta\varepsilon)^{2/r}\varepsilon k$, where $c' > 0$ is a constant depending on r only, which implies the claim. □

The above bound is slightly poorer than the one stated in [Woo07], albeit only by a $\mathrm{poly}(\log k)$ factor. It would be interesting to see if Elton's Theorem can be avoided in creating a copy of ℓ_1^d inside $\mathcal{L}(N; (2,2))$.

We proceed with the proof of Theorem A.1, for which we use the following slight variant of a standard "matching lemma" of [BARdW08, Appendix B], shown in [BNR12] for $q = 3$. We omit the proof, which is a straightforward modification of [BNR12, Lemma 3.1].

Lemma A.4 (Ben-Aroya–Regev–de Wolf)**.** *Let $C : \{-1,1\}^k \to \{-1,1\}^N$ be a (q, δ, ε)-LDC. Then, there exists a function $C' : \{-1,1\}^k \to \{-1,1\}^{qN}$ such that the following holds. For every $i \in [k]$ there exists a family $\mathcal{M}_i$ of at least $\delta\varepsilon N/q^2$ pairwise disjoint sets $S \subseteq [qN]$ of q elements each, such that for a uniformly distributed random string $x \in \{-1,1\}^k$, we have*

$$\Big|\mathbb{E}\Big[x_i \prod_{s\in S} C'(x)_s\Big]\Big| \geq \frac{\varepsilon}{2^q}. \tag{23}$$

We also use the following proposition, which may be interpreted as a generalization of the Birthday Paradox.

Proposition A.5. *For $\eta > 0$ and positive integers N and $r \geq 2$, let $\mathcal{F}$ be a family of ηN pairwise disjoint subsets $S \subseteq [N]$, each of cardinality $|S| = 2r$. Let l be a positive integer such that*

$$\eta l^r \geq N^{r-1}. \tag{24}$$

Then, there exists a set $\mathcal{I} \subseteq [N]^l$ of cardinality at least $N^l/4^r$ such that for each sequence $\mathbf{S} \in \mathcal{I}$, there exists an $S \in \mathcal{F}$ for which $|S \cap \mathbf{S}| \geq r$.

The proof of Proposition A.5 uses a standard Poisson approximation result for "balls and bins" problems [MU05, Theorem 5.10]. A discrete Poisson random variable Y with expectation μ is nonnegative, integer valued, and has probability density function

$$\Pr[Y = m] = \frac{e^{-\mu}\mu^m}{m!}, \qquad \forall m = 0, 1, 2, \ldots \tag{25}$$

Theorem A.6 (Poisson approximation)**.** *For positive integers l and N, suppose we toss l balls into N bins independently and uniformly at random. For each $s \in [N]$ let X_s be the random variable counting the number of balls in bin number s. Let $Y_1, \ldots, Y_N$ be independent Poisson random variables with expectation l/N. Then, for any function $f : \{0, \ldots, l\}^N \to \mathbf{R}$ such that $\mathbb{E}[f(X_1, \ldots, X_N)]$ increases or decreases monotonically with l, we have $\mathbb{E}[f(X_1, \ldots, X_N)] \leq 2\mathbb{E}[f(Y_1, \ldots, Y_N)]$.*

Proof of Proposition A.5. Let us assume for simplicity that N is a multiple of $2r$. Partition the elements in $[N]$ not covered by any set in $\mathcal{F}$ into disjoint sets of size $2r$. With $\mathcal{F}$, this gives a partition $\mathcal{P}$ of $[N]$ into $M = N/(2r)$ sets of size $2r$. Label the sets in $\mathcal{F}$ with distinct numbers in $[\eta N]$ and label the remaining partitions with distinct numbers in $\{\eta N + 1, \dots, M\}$.

Let $b_1, \dots, b_l$ be independent uniformly distributed random variables over $[N]$, the balls. Say that ball b_j lands in bin $S \in \mathcal{P}$ if $b_j \in S$ and notice that the balls land in a uniformly random bin. For each $s \in [M]$ let X_s be the random variable counting the number of balls in bin number s and let Y_s be a discrete Poisson random variable with expectation $\mu = l/M$.

Let $f : \{0, 1, 2, \dots\}^M \to \{0, 1\}$ be the function that assumes the value 1 if and only if its first ηN variables have value strictly less than r. Clearly $\mathbb{E}[f(X_1, \dots, X_M)]$ decreases monotonically with l, the number of balls we toss, since this expectation equals the probability that all bins in $\mathcal{F}$ have strictly less than r balls. Therefore, by Theorem A.6, we have

$$\Pr[f(X_1, \dots, X_M) = 1] \leq 2\Pr[f(Y_1, \dots, Y_M) = 1]. \tag{26}$$

Independence of the Y_j gives

$$\begin{aligned}
\Pr[f(Y_1, \dots, Y_M) = 1] &= \prod_{j=1}^{\eta N} \Pr[Y_j < r] \\
&= \left(\sum_{m=0}^{r-1} \frac{e^{-\mu}\mu^m}{m!} \right)^{\eta N} \\
&= \left(1 - e^{-\mu} \sum_{m=r}^{\infty} \frac{\mu^m}{m!} \right)^{\eta N} \\
&\leq \left(1 - \frac{\mu^r}{r!} \right)^{\eta N}, \qquad (27)
\end{aligned}$$

where in the third line we used the Taylor expansion of the exponential function at zero. By our assumption (24) on l and the easy bound $r^r/r! \geq 1$, we have

$$\frac{\mu^r}{r!} = \frac{1}{r!}\left(\frac{l}{M}\right)^r = \frac{1}{r!}\left(\frac{2rl}{N}\right)^r \geq \frac{1}{\eta N}.$$

Hence, (27) is at most $1/e$ and it follows from (26) that with probability $1 - 2/e \geq 1/4$, one of the first ηN bins has at least r balls. In other words, for at least $N^l/4$ sequences $\mathbf{S} \in [N]^l$ there exists an $S \in \mathcal{F}$ such that r entries of $\mathbf{S}$ belong to S. Of those sequences, a $2r(2r-1)\cdots r/(2r)^r \geq (1/2)^r$ fraction has those entries distinct. Since we assumed that $r \geq 2$, at least $N^l/(2^{2+r}) \geq N^l/4^r$ of the sequences have the desired property. □

Proof of Theorem A.1. Let $C : \{-1, 1\}^k \to \{-1, 1\}^N$ be a $(2r, \delta, \varepsilon)$-LDC. Let $C' : \{-1, 1\}^k \to \{-1, 1\}^{2rN}$ be a map and $\mathcal{M}_i$ be families as in Lemma A.4. Let $N' = 2rN$ and recall that each $\mathcal{M}_i$ consists of at least $\delta\varepsilon N'/(2r)^3$ pairwise disjoint sets $S \subseteq [N']$ of cardinality $2r$ each. Let l be an integer such that (21) holds. Fix an $i \in [k]$. By proposition A.5, there is a set $\mathcal{I}_i \subseteq [N']^l$ of at least $(N')^l/4^r$ sequences $\mathbf{S} \in [N']^l$ such that $|S \cap \mathbf{S}| \geq r$ for some $S \in \mathcal{M}_i$.

We define a partial matching $\mathcal{P}_i$ in $[N']^l$. For each $\mathbf{S} \in \mathcal{I}_i$ and associated $S \in \mathcal{M}_i$, pick a set $T \subseteq [l]$ of r coordinates such that $|\mathbf{S}_T \cap S| = r$. Let $\mathbf{S}' \in [N']^l$ be a sequence such that $\mathbf{S}'_t = \mathbf{S}_t$

for every $t \notin T$ and such that (with slight abuse of notation) $\mathbf{S}_T \cup \mathbf{S}'_T = S$. There are $r!$ choices of $\mathbf{S}'_T$. Choose one arbitrarily but uniquely. Let $\mathcal{P}_i$ be the family of said $\{\mathbf{S}, \mathbf{S}'\}$ pairs and observe that $|\mathcal{P}_i| = |\mathcal{I}_i|/2 \geq (N')^l)/4^r$.

Define a matrix $A_i : [N']^l \times [N']^l \to \{-1, 0, 1\}$ as follows. Let x be a uniformly distributed random string over $\{-1, 1\}^k$ and notice that $x_1, \ldots, x_k$ are independent and uniformly distributed over $\{-1, 1\}$. For each pair $\{\mathbf{S}, \mathbf{S}'\} \in \mathcal{P}_i$ with associated set $S \in \mathcal{M}_i$ as above, set

$$A_i(\mathbf{S}, \mathbf{S}') = \operatorname{sign}\Big(\mathbb{E}\Big[x_i \prod_{s \in S} C'(x)_s\Big]\Big). \tag{28}$$

Set the other entries of A_i to zero. Moreover, since $\mathcal{P}_i$ is a (partial) matching, each row and column of A_i has at most one nonzero element and it follows that $\|A_i\|_{S_\infty} \leq 1$. For each $x \in \{-1, 1\}^n$ let $D(x) = C'(x)^{\otimes l}$. Then, for each $\{\mathbf{S} = (s_1, \ldots, s_l), \mathbf{S}' = (s'_1, \ldots, s'_l)\} \in \mathcal{P}_i$, with associated sets $S \in \mathcal{M}_i$ and $T \subseteq [l]$ as above,

$$D(x)_{\mathbf{S}} D(x)_{\mathbf{S}'} = \Big(\prod_{t \notin T} C'(x)^2_{s_t}\Big) \prod_{t \in T} C'(x)_{s_t} C'(x)_{s'_t} = \prod_{s \in S} C'(x)_s. \tag{29}$$

Hence, by Lemma A.4, for a uniformly distributed $x \in \{-1, 1\}^k$, we have

$$\begin{aligned}
\mathbb{E}\Big[\Big\|\sum_{i=1}^k x_i A_i\Big\|_{S_\infty}\Big] &\geq \mathbb{E}\Big[\frac{D(x)^{\mathsf{T}}}{\sqrt{(N')^l}}\Big(\sum_{i=1}^k x_i A_i\Big)\frac{D(x)}{\sqrt{(N')^l}}\Big] \\
&\overset{(28)}{=} \mathbb{E}\Big[\frac{2}{(N')^l}\sum_{i=1}^k \sum_{\{\mathbf{S},\mathbf{S}'\} \in \mathcal{P}_i} x_i\, D(x)_{\mathbf{S}} D(x)_{\mathbf{S}'}\, A_i(\mathbf{S}, \mathbf{S}')\Big] \\
&\overset{(29)}{=} \frac{2}{(N')^l}\sum_{i=1}^k \sum_{\{\mathbf{S},\mathbf{S}'\} \in \mathcal{P}_i} \Big|\mathbb{E}_x\Big[x_i \prod_{s \in S} C(x)_s\Big]\Big| \overset{(23)}{\geq} \frac{\varepsilon k}{8^r}.
\end{aligned}$$

□

ECCC ISSN 1433-8092
http://eccc.hpi-web.de