

Outlaw distributions and locally decodable codes*

Jop Briët[†]

Zeev Dvir[‡]

Sivakanth Gopi[§]

Abstract

Locally decodable codes (LDCs) are error correcting codes that allow for decoding of a single message bit using a small number of queries to a corrupted encoding. Despite decades of study, the optimal trade-off between query complexity and codeword length is far from understood. In this work, we give a new characterization of LDCs using distributions over Boolean functions whose expectation is hard to approximate (in L_∞ norm) with a small number of samples. We coin the term ‘outlaw distributions’ for such distributions since they ‘defy’ the Law of Large Numbers. We show that the existence of outlaw distributions over sufficiently ‘smooth’ functions implies the existence of constant query LDCs and vice versa. We give several candidates for outlaw distributions over smooth functions coming from finite field incidence geometry, additive combinatorics and from hypergraph (non)expanders.

We also prove a useful lemma showing that (smooth) LDCs which are only required to work on average over a random message and a random message index can be turned into true LDCs at the cost of only constant factors in the parameters.

1 Introduction

Error correcting codes (ECCs) solve the basic problem of communication over noisy channels. They encode a message into a codeword from which, even if the channel partially corrupts it, the message can later be retrieved. With one of the earliest applications of the *probabilistic method*, formally introduced by Erdős in 1947, pioneering work of Shannon [Sha48] showed the existence of optimal (capacity-achieving) ECCs. The problem of explicitly constructing such codes has fueled the development of coding theory ever since. Similarly, the exploration of many other fascinating structures, such as Ramsey graphs, expander graphs, two source

*A preliminary version of this paper appeared in the proceedings of ITCS 2017 [BDG17]

[†]CWI, Science Park 123, 1098 XG Amsterdam, the Netherlands. Supported by a VENI grant and the Gravitation-grant NETWORKS-024.002.003 from the Netherlands Organisation for Scientific Research (NWO). Email: j.briet@cw.nl.

[‡]Department of Computer Science and Department of Mathematics, Princeton University, Princeton, NJ 08540, USA. Research supported by NSF grant CCF-1523816 and by the Sloan foundation. Email: zeev.dvir@gmail.com.

[§]Department of Computer Science, Princeton University, Princeton, NJ 08540, USA. Research supported by NSF grant CCF-1523816 and by the Sloan foundation. Email: sgopi@cs.princeton.edu.

extractors, etc., began with a striking existence proof via the probabilistic method, only to be followed by decades of catch-up work on explicit constructions. Locally decodable codes (LDCs) are a special class of error correcting codes whose development has not followed this line. The defining feature of LDCs is that they allow for ultra fast decoding of single message bits, a property that typical ECCs lack, as their decoders must read an entire (possibly corrupted) codeword to achieve the same. They were first formally defined in the context of channel coding in [KT00], although they (and the closely related locally correctable codes) implicitly appeared in several previous works in other settings, such as program checking [BK95], probabilistically checkable proofs [AS98, ALM⁺98] and private information retrieval schemes (PIRs) [CKGS98]. More recently, LDCs have even found applications in Banach-space geometry [BNR12] and LDC-inspired objects called local reconstruction codes found applications in fault tolerant distributed storage systems [GHSY12]. See [Yek12] for a survey of LDCs and some of the applications.

Despite their many applications, our knowledge of LDCs is very limited; the best-known constructions are far from what is currently known about their limits. Although standard random (linear) ECCs do allow for some weak local-decodability, they are outperformed by even the earliest explicit constructions [KS07]. All the known constructions of LDCs were obtained by explicitly designing such codes using some algebraic objects like low-degree polynomials or matching vectors [Yek12].

In this paper, we give a characterization of LDCs in probabilistic and geometric terms, making them amenable to probabilistic constructions. On the flip side, these characterizations might also be easier to work with for the purpose of showing lower bounds. We will make this precise in the next section. Let us first give the formal definition of an LDC.

Definition 1.1 (Locally decodable code). For positive integers k, n, q and parameters $\eta, \delta \in (0, 1/2]$, a map $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a (q, δ, η) -*locally decodable code* if, for every $i \in [k]$, there exists a randomized *decoder* (a probabilistic algorithm) \mathcal{A}_i such that:

- For every message $x \in \{0, 1\}^k$ and string $y \in \{0, 1\}^n$ that differs from the codeword $C(x)$ in at most δn coordinates,

$$\Pr[\mathcal{A}_i(y) = x_i] \geq \frac{1}{2} + \eta. \quad (1)$$

- The decoder \mathcal{A}_i (non-adaptively) queries at most q coordinates of y .¹

Known results The main parameters of LDCs are the number of queries q and the length of the encoding n as a function of k and q , typically the parameters δ, η are fixed constants. The simplest example is the Hadamard code, which is a 2-query LDC with $n = 2^k$. The 2-query regime is the only nontrivial case where optimal lower bounds are known: it was

¹ We can assume that on input $y \in \{0, 1\}^n$, the decoder \mathcal{A}_i first samples a set $S \subseteq [n]$ of at most q coordinates according to a probability distribution depending on i only and then returns a random bit depending only on i, S and the values of y at S .

shown in [KW04, GKST02] that exponential length is necessary. In general, Reed-Muller codes of degree $q - 1$ are q -query LDCs of length $n = \exp(O(k^{1/q-1}))$. For a long time, these were the best constructions for constant q , until in a breakthrough work by [Yek07, Efr09], 3-query LDCs were constructed with subexponential length $n = \exp(\exp(O(\sqrt{\log k})))$. More generally they constructed 2^r -query LDCs with length $n = \exp(\exp(O(\log^{1/r} k)))$. For $q \geq 3$, the best-known lower bounds leave huge gaps, giving only polynomial bounds. A 3-query LDC must have length $n \geq \tilde{\Omega}(k^2)$, and more generally for a q -query LDC, $n \geq \tilde{\Omega}(k^{1+1/(\lceil q/2 \rceil - 1)})$ [KW04, Woo07]². LDCs where the codewords are over a large alphabet are studied because of their relation to private information retrieval schemes [CKGS98, KT00]. In [DG15], 2-query LDCs of length $n = \exp(k^{o(1)})$ over an alphabet of size $\exp(k^{o(1)})$ were constructed. There is also some exciting recent work on LDCs where the number of queries grows with k , in which case there are explicit constructions with constant-rate (that is, $n = O(k)$) and query complexity $q = \exp(O(\sqrt{\log n}))$; in fact we can even achieve the optimal rate-distance tradeoff of traditional error correcting codes [KSY14, KMRS16, GKO⁺17]. We cannot yet rule out the exciting possibility that constant rate LDCs with polylogarithmic query complexity exist.

1.1 LDCs from distributions over smooth Boolean functions

Our main result shows that LDCs can be obtained from “outlaw” distributions over “smooth” functions. The term outlaw refers to the Law of Large Numbers, which says that the average of independent samples tends to the expectation of the distribution from which they are drawn. Roughly speaking, a probability distribution is an outlaw if many samples are needed for a good estimation of the expectation and a smooth function over the n -dimensional Boolean hypercube is one that has no influential variables. Paradoxically, while many instances of the probabilistic method use the fact that sample means of a small number of independent random variables tend to concentrate around the true mean, as captured for example by the Chernoff bound, our main result requires precisely the opposite. We show that if *at least* k samples from a distribution over smooth functions are needed to approximate the mean, then there exists an $O(1)$ -query LDC sending $\{0, 1\}^{\Omega(k)}$ to $\{0, 1\}^n$, where the hidden constants depend only the smoothness and mean-estimation parameters.

To make this precise, we now formally define smooth functions and outlaw distributions. Given a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, its *spectral norm* is defined as

$$\|f\|_{\text{sp}} = \sum_{S \subseteq [n]} |\hat{f}(S)|,$$

where $\hat{f}(S)$ are the Fourier coefficients of f (see Section 2 for basics on Fourier analysis). We also consider the supremum norm, $\|f\|_{L_\infty} = \sup\{|f(x)| : x \in \{-1, 1\}^n\}$. It follows from the Fourier inversion formula that $\|f\|_{L_\infty} \leq \|f\|_{\text{sp}}$. The *i th discrete derivative* of f is the function $(D_i f)(x) = (f(x) - f(x^i))/2$, where x^i is the point that differs from x on the i th

²We use $\tilde{\Omega}(\cdot)$, $\tilde{O}(\cdot)$, $\tilde{\omega}(\cdot)$, $\tilde{o}(\cdot)$ to hide polylogarithmic factors through out this paper.

coordinate. Smooth functions are functions whose discrete derivatives have small spectral norms.

Definition 1.2 (σ -smooth functions). For $\sigma > 0$, a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is σ -smooth, if for every $i \in [n]$, we have $\|D_i f\|_{\text{sp}} \leq \sigma/n$.

Intuition for the above definition may be gained from the fact that smooth functions have no influential variables. The influences, $(\mathbb{E}_{x \in \{-1, 1\}^n} [(D_i f)(x)^2])^{1/2}$, measure the extent to which changing the i th coordinate of a randomly chosen point changes the value of f . Since $\|D_i f\|_{L^\infty} \leq \|D_i f\|_{\text{sp}}$, the directional derivatives of σ -smooth functions are uniformly bounded by σ/n , which is a much stronger condition than saying that the derivatives are small on average. Outlaws are defined as follows.

Definition 1.3 (Outlaw). Let n be a positive integer and μ be a probability distribution over real-valued functions on $\{-1, 1\}^n$. For a positive integer k and $\varepsilon > 0$, say that μ is a (k, ε) -outlaw if for independent random μ -distributed functions f_1, \dots, f_k and $\bar{f} = \mathbb{E}_\mu[f]$,

$$\mathbb{E} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - \bar{f}) \right\|_{L^\infty} \right] \geq \varepsilon.$$

Denote by $\kappa_\mu(\varepsilon)$ the largest integer k such that μ is a (k, ε) -outlaw.

To approximate the true mean of an outlaw μ to within ε on average in the L_∞ -distance, one thus needs $\kappa_\mu(\varepsilon) + 1$ samples. Note that if μ is a distribution over σ -smooth functions, then the distribution $\tilde{\mu}$ obtained by scaling functions in the support of μ by $1/\sigma$ is a distribution over 1-smooth functions and $\kappa_{\tilde{\mu}}(\varepsilon/\sigma) = \kappa_\mu(\varepsilon)$.

Our main result is then as follows.

Theorem 1.4 (Main theorem). *Let n be a positive integer and $\varepsilon > 0$. Let μ be a probability distribution over 1-smooth functions on $\{-1, 1\}^n$ and $k = \kappa_\mu(\varepsilon)$. Then, there exists a (q, δ, η) -LDC sending $\{0, 1\}^l$ to $\{0, 1\}^n$ where $l = \Omega(\varepsilon^2 k / \log(1/\varepsilon))$, $q = O(1/\varepsilon)$, $\delta = \Omega(\varepsilon)$ and $\eta = \Omega(\varepsilon)$. Additionally, if μ is supported by degree- d functions, then we can take $q = d$.*

Note that the smoothness requirement is essential. For example the uniform distribution over the n dictator functions $f_i(x) = x_i$ for $i \in [n]$ is an $(n/2, 1)$ -outlaw, but it cannot imply constant rate, constant query LDCs which we know do not exist. In fact we establish a converse to Theorem 1.4, showing that its hypothesis is essentially equivalent to the existence of LDCs in the small query complexity regime.

Theorem 1.5. *If $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a (q, δ, η) -LDC, then there exists a probability distribution μ over 1-smooth degree- q functions on $\{-1, 1\}^n$ such that*

$$\kappa_\mu(\varepsilon) \geq \eta k$$

where $\varepsilon = \eta\delta/(q2^{q/2})$.

Theorem 1.5 can in turn convert the problem of proving lower bounds on the length of LDCs to a problem on Banach space geometry. In particular, for a distribution μ over 1-smooth degree- q functions on $\{0, 1\}^n$, one can upper bound $\kappa_\mu(\varepsilon)$ in terms of type constants of the space of q -linear forms on ℓ_q^{n+1} [Bri16].

Candidate outlaws One scenario in which outlaw distributions can be obtained is using incidence geometry in finite fields. In particular, the following result can be derived from our main theorem (stated a bit informally here, see Section 6.1 for the formal version).

Corollary 1.6. *Let $p > 2$ be a fixed prime. Suppose that for a random set of directions $D \subset \mathbb{F}_p^n$ of size $|D| \leq k$, with probability at least $1/2$, there exists a set $B \subset \mathbb{F}_p^n$ of size $|B| \geq \Omega(p^n)$ which does not contain any lines with direction in D . Then, there exists a p -query LDC sending $\{0, 1\}^{\Omega(k)}$ to $\{0, 1\}^{p^n}$.*

The assumption in Corollary 1.6 that D be random is essential for it to be potentially interesting for LDCs. If we instead ask that every set of directions D satisfies the condition—as we did in the conference version of this paper—then letting D be a subspace shows that k must be smaller than a constant depending only on p and ε by Szemerédi’s Theorem (Theorem 6.2 below) [Fox].

The analogue of Corollary 1.6 in $\mathbb{Z}/N\mathbb{Z}$ where lines correspond to arithmetic progressions and directions correspond to common differences can also be used to construct LDCs. This question was studied in [FLW16], where they show that if D is a random subset of $\mathbb{Z}/N\mathbb{Z}$ of size $\omega(N^{1-1/q})$, then almost surely every dense subset of $\mathbb{Z}/N\mathbb{Z}$ contains a q -term arithmetic progression with common difference in D . Our main result, together with the best-known lower bounds on LDCs [KW04, Woo07] can be used to improve the bounds of [FLW16] on random differences in Szemerédi’s theorem as shown in the following corollary.

Corollary 1.7. *For every $\alpha > 0$, there is $N(\alpha)$ such that for every $N > N(\alpha)$ the following is true: if $D \subset \mathbb{Z}/N\mathbb{Z}$ is a uniformly random subset of size $\tilde{\omega}(N^{1-1/\lceil q/2 \rceil})$ when $q > 2$ and $\omega(\log N)$ when $q = 2$, then with probability $1 - o(1)$, every subset $A \subset \mathbb{Z}/N\mathbb{Z}$ of size at least αN contains a q -term arithmetic progression with common difference in the set D .*

Another setting in which our approach leads to interesting open problems is in relation to pseudorandom hypergraphs. Consider a partition of the complete bipartite graph $K_{n,n}$ into n perfect matchings. It is known that picking $k = O(\log n)$ of these matchings at random will give us a pseudorandom (expander) graph (of degree k). For some particular partitions (e.g., given by an Abelian group) this bound is tight. The questions arising from our approach can be briefly summarized as follows: Can one find an n -vertex hypergraph H (say three uniform to be precise) and a partition of H into matchings so that, to get a pseudorandom hypergraph (defined appropriately) one needs at least k random matchings. This would give a code sending $\Omega(k)$ -bit messages with encoding length $O(n)$ and so, becomes interesting when k is super poly-logarithmic in n . We elaborate on this in Section 6.2

1.2 Techniques

Our proof of Theorem 1.4 proceeds in two steps. The first step consists of turning an outlaw over smooth functions into a seemingly crude type of LDC that is only required to work on average over a uniformly distributed message and a uniformly distributed message index. We call such codes *average-case smooth codes* (see below). The second step consists of showing that such codes are in fact not much weaker than honest LDCs.

From outlaws to average-case smooth codes The key ingredient for the first step is *symmetrization*, a basic technique from high-dimensional probability. We briefly sketch how this is used (we refer to Section 3 for the full proof). Suppose that f_1, \dots, f_k are independent smooth functions distributed according to a (k, ε) -outlaw with expectation \bar{f} . We introduce an independent copy³ f'_i of f_i for each $i \in [k]$ and consider the symmetrically distributed random functions $f_i - f'_i$. Since $\bar{f} = \mathbb{E}[f'_i]$ for each $i \in [k]$, Jensen's inequality and Definition 1.3 imply that

$$\mathbb{E}[\|(f_1 - f'_1) + \dots + (f_k - f'_k)\|_{L_\infty}] \geq \mathbb{E}[\|(f_1 - \mathbb{E}[f'_1]) + \dots + (f_k - \mathbb{E}[f'_k])\|_{L_\infty}] \geq \varepsilon k.$$

Since the random functions $f_i - f'_i$ are independent and symmetric, we get that for independent uniformly random signs $x_1, \dots, x_k \in \{-1, 1\}$, the above left-hand side equals

$$\mathbb{E}[\|x_1(f_1 - f'_1) + \dots + x_k(f_k - f'_k)\|_{L_\infty}].$$

The triangle inequality and the Averaging Principle then give that there exist *fixed* smooth functions f_1^*, \dots, f_k^* such that on average over the random signs, we have

$$\mathbb{E}[\|x_1 f_1^* + \dots + x_k f_k^*\|_{L_\infty}] \geq \varepsilon k/2. \quad (2)$$

To get an average-case smooth code out of this, we view each sequence $x = (x_1, \dots, x_k)$ as a k -bit message and choose an arbitrary n -bit string for which the L_∞ -norm in (2) is achieved to be the its encoding, $C(x)$. This gives a map $C : \{-1, 1\}^k \rightarrow \{0, 1\}^n$ satisfying

$$\mathbb{E}[x_1 f_1^*(C(x)) + \dots + x_k f_k^*(C(x))] \geq \varepsilon k/2.$$

Equivalently, for uniform x and i , we have $\Pr[f_i^*(C(x)) = x_i] \geq \frac{1}{2} + \frac{\varepsilon}{4}$. Finally, we use the smoothness property to transform the f_i^* into decoders with the desired properties. This is done in Section 3. It is in the application of the Averaging Principle where the probabilistic method appears in our construction of LDCs.

Average-case smooth codes are LDCs Katz and Trevisan [KT00] observed that LDC decoders must have the property that they select their queries according to distributions that do not favor any particular coordinate. The intuition for this is that if they did favor a certain coordinate, then corrupting that coordinate would cause the decoder to err with too high a probability. If instead, queries are sampled according to a “smooth” distribution, they will all fall on uncorrupted coordinates with good probability provided the fraction of corrupted coordinates δ and query complexity q aren't too large. The following definition allows us to make this intuition precise.

Definition 1.8 (Smooth code). For positive integers k, n, q and parameters $\eta \in (0, 1/2]$ and $c > 0$, a map $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a (q, c, η) -smooth code if, for every $i \in [k]$, there exists a randomized decoder $\mathcal{A}_i : \{0, 1\}^n \rightarrow \{0, 1\}$ such that

³in this context sometimes referred to as a “ghost copy” as it will later disappear again

- For every $x \in \{0, 1\}^k$,

$$\Pr[x_i = \mathcal{A}_i(C(x))] \geq \frac{1}{2} + \eta. \quad (3)$$

- The decoder \mathcal{A}_i (non-adaptively) queries at most q coordinates of $C(x)$.
- For each $j \in [n]$, the probability that \mathcal{A}_i queries the coordinate $j \in [n]$ is at most c/n .

The formal version of Katz and Trevisan's observation is as follows.

Theorem 1.9 (Katz–Trevisan). *If $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a (q, δ, η) -LDC, then C is also a $(q, q/\delta, \eta)$ -smooth code. Conversely, if $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a (q, c, η) -smooth code, then C is also a $(q, \eta/2c, \eta/2)$ -LDC.*

Our second step in the proof of Theorem 1.4 is a stronger form of the converse part of Theorem 1.9. We show that even smooth codes that are only required to work *on average* can be turned into LDCs, losing only a constant factor in the rate and success probability.

Definition 1.10 (Average-case smooth code). A code as in Definition 1.8 is a (q, c, η) -average-case smooth code if instead of the first item, (3) is required to hold only on average over uniformly distributed $x \in \{0, 1\}^k$ and uniformly distributed $i \in [k]$, which is to say that

$$\Pr[x_i = \mathcal{A}_i(C(x))] \geq \frac{1}{2} + \eta,$$

where the probability is taken over x, i and the randomness used by \mathcal{A}_i .

Lemma 1.11. *Let $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ be a (q, c, η) -average-case smooth code. Then, there exists an $(q, \Omega(\eta/c), \Omega(\eta))$ -LDC sending $\{0, 1\}^l$ to $\{0, 1\}^n$ where $l = \Omega(\eta^2 k / \log(1/\eta))$.*

The idea behind the proof of Lemma 1.11 is as follows. We first switch the message and codeword alphabets to $\{-1, 1\}$ and let $f_i : \{-1, 1\}^k \rightarrow [-1, 1]$ be the expected decoding function $f_i(z) = \mathbb{E}[\mathcal{A}_i(z)]$. The properties of C then easily imply that the set $T \subseteq [-1, 1]^k$ given by $T = \{(f_1(z), \dots, f_k(z)) : z \in \{-1, 1\}^k\}$ has large *Gaussian width*, in particular it holds that for a standard k -dimensional Gaussian vector g , we have $\mathbb{E}[\sup_{t \in T} \langle g, t \rangle] \gtrsim \varepsilon k$.⁴ Next, we employ a powerful result of [MV03] showing that T contains an l -dimensional hypercube-like structure with edge length some absolute constant $c \in (0, 1]$, for $l \gtrsim k$. Roughly speaking, this implies that C is a smooth code on $\{-1, 1\}^l$ whose decoding probability depends on ε and c . Finally, we obtain an LDC via an application of Theorem 1.9. The full proof is given in Section 4.

⁴We write $A \gtrsim B$ and $A = \Omega(B)$ interchangeably to mean that $A \geq cB$ for some absolute constant $c > 0$ independent of all parameters involved.

1.3 Organization

Section 2 contains some preliminaries in Fourier analysis over the Boolean cube. In Section 3, we prove our main theorem (Theorem 1.4) by first showing that outlaw distributions over smooth functions imply existence of average-case smooth codes and using Lemma 1.11 to convert them to LDCs. In Section 4, we prove Lemma 1.11 showing how to convert average-case smooth-codes to LDCs. In Section 5, we show the converse to our main theorem (Theorem 1.5) showing how to get outlaw distributions over smooth functions from LDCs. Finally in Section 6, we give some candidate constructions of outlaw distributions over smooth functions using incidence geometry and hypergraph pseudorandomness.

2 Preliminaries

We recall a few basic definitions and facts from analysis over the n -dimensional Boolean hypercube $\{-1, 1\}^n$. Equipped with the coordinate-wise multiplication operation, the hypercube forms an Abelian group whose group of characters is formed by the functions $\chi_S(x) = \prod_{i \in S} x_i$ for all $S \subseteq [n]$. The characters form a complete orthonormal basis for the space of real-valued functions on $\{-1, 1\}^n$ endowed with the inner product $\langle f, g \rangle = \mathbb{E}_{x \in \{-1, 1\}^n} [f(x)g(x)]$, where we use the notation $\mathbb{E}_{a \in S}$ to denote the expectation with respect to a uniformly distributed element a over a set S . The *Fourier transform* of a function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is the function $\hat{f} : 2^{[n]} \rightarrow \mathbb{R}$ defined by $\hat{f}(S) = \langle f, \chi_S \rangle$. The Fourier inversion formula (which follows from orthonormality of the character functions) asserts that

$$f = \sum_{S \subseteq [n]} \hat{f}(S) \chi_S.$$

Parseval's Identity relates the L_2 -norms of f and its Fourier transform by

$$(\mathbb{E}_{x \in \{-1, 1\}^n} [f(x)^2])^{1/2} = \left(\sum_{S \subseteq [n]} |\hat{f}(S)|^2 \right)^{1/2}.$$

A function f has *degree* q if $\hat{f}(S) = 0$ when $|S| > q$ and the *degree- q truncation* of f , denoted $f^{\leq q}$, is the degree- q function defined by

$$f^{\leq q} = \sum_{|S| \leq q} \hat{f}(S) \chi_S.$$

A function f is a *q -junta* if it depends only on a subset of q of its variables, or equivalently, if there exists a subset $T \subseteq [n]$ of size $|T| \leq q$ such that $\hat{f}(S) = 0$ for every $S \not\subseteq T$. The *i th discrete derivative* $D_i f$ is the function $(D_i f)(x) = (f(x) - f(x^i))/2$, where x^i is the point that differs from x on the i th coordinate. It is easy to show that the i th discrete derivative in of a function f is given by

$$D_i f = \sum_{S \ni i} \hat{f}(S) \chi_S.$$

Hence, it follows that $\|D_i f\|_{\text{Spec}} = \sum_{S \ni i} |\hat{f}(S)|$.

3 From outlaws to LDCs

In this section we prove Theorem 1.4. For convenience, in the remainder of this paper, we switch the message and codeword alphabets of all codes from $\{0, 1\}^n$ to $\{-1, 1\}^n$. We begin by showing that outlaw distributions over degree- q functions give q -query average-case smooth codes. Combined with Lemma 1.11, this implies the second part of Theorem 1.4.

Theorem 3.1. *Let μ be a probability distribution on 1-smooth degree- q functions on $\{-1, 1\}^n$, let $\varepsilon \in (0, 1]$ and let $k = \kappa_\mu(\varepsilon)$. Then, there exists a $(q, 1, \varepsilon/4)$ -average-case smooth code sending $\{-1, 1\}^k$ to $\{-1, 1\}^n$.*

Proof: The proof uses a symmetrization argument. Let $\mathcal{F} = (f_1, \dots, f_k)$ and $\mathcal{F}' = (f'_1, \dots, f'_k)$ be two k -tuples of independent μ -distributed random variables and let $\bar{f} = \mathbb{E}_\mu[f]$. Then, by definition of $\kappa_\mu(\varepsilon)$ and Jensen's inequality,

$$\begin{aligned} \varepsilon &\leq \mathbb{E}_{\mathcal{F}} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - \bar{f}) \right\|_{L_\infty} \right] \\ &= \mathbb{E}_{\mathcal{F}} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - \mathbb{E}_{\mathcal{F}'}[f'_i]) \right\|_{L_\infty} \right] \\ &\leq \mathbb{E}_{\mathcal{F}, \mathcal{F}'} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - f'_i) \right\|_{L_\infty} \right]. \end{aligned}$$

The random variables $f_i - f'_i$ are symmetrically distributed, which is to say that they have the same distribution as their negations $f'_i - f_i$. Since they are independent, it follows that for every $x \in \{-1, 1\}^k$, the random variable $x_1(f_1 - f'_1) + \dots + x_k(f_k - f'_k)$ has the same distribution as $(f_1 - f'_1) + \dots + (f_k - f'_k)$. Therefore,

$$\begin{aligned} \mathbb{E}_{\mathcal{F}, \mathcal{F}'} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - f'_i) \right\|_{L_\infty} \right] &= \mathbb{E}_{x \in \{-1, 1\}^k} \left[\mathbb{E}_{\mathcal{F}, \mathcal{F}'} \left[\left\| \frac{1}{k} \sum_{i=1}^k x_i (f_i - f'_i) \right\|_{L_\infty} \right] \right] \\ &\leq 2 \mathbb{E}_{\mathcal{F}} \left[\mathbb{E}_{x \in \{-1, 1\}^k} \left[\left\| \frac{1}{k} \sum_{i=1}^k x_i f_i \right\|_{L_\infty} \right] \right]. \end{aligned}$$

Applying the Averaging Principle to the outer expectation, we find that there exist 1-smooth degree- q functions $f_1^*, \dots, f_k^* : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that

$$\mathbb{E}_{x \in \{-1, 1\}^k} \left[\left\| \frac{1}{k} \sum_{i=1}^k x_i f_i^* \right\|_{L_\infty} \right] \geq \frac{\varepsilon}{2}. \quad (4)$$

Define the code $C : \{-1, 1\}^k \rightarrow \{-1, 1\}^n$ such that for each $x \in \{-1, 1\}^k$, we have

$$\frac{1}{k} \sum_{i=1}^k x_i f_i^*(C(x)) = \left\| \frac{1}{k} \sum_{i=1}^k x_i f_i^* \right\|_{L_\infty}. \quad (5)$$

For each $i \in [k]$, define the decoder \mathcal{A}_i as follows. Let $\nu_i : 2^{[n]} \rightarrow [0, 1]$ be the probability distribution defined by $\nu_i(S) = |\widehat{f}_i^*(S)| / \|\widehat{f}_i^*\|_{\text{sp}}$. Given a string $z \in \{-1, 1\}^n$, with probability $1 - \|\widehat{f}_i^*\|_{\text{sp}}$, the decoder \mathcal{A}_i returns a uniformly random sign, and with probability $\|\widehat{f}_i^*\|_{\text{sp}}$, it samples a set $S \subseteq [n]$ according to ν_i and returns $\chi_S(z)$. This is a valid probability distribution since for any 1-smooth function f , we have

$$\|f\|_{\text{sp}} = \sum_{S \subseteq [n]} |\widehat{f}(S)| \leq \sum_{S \subseteq [n]} |S| |\widehat{f}(S)| = \sum_{i=1}^n \sum_{S \ni i} |\widehat{f}(S)| \leq n \cdot \frac{1}{n} = 1.$$

Then, \mathcal{A}_i queries at most q coordinates of z and since f_i^* is 1-smooth, the probability that it queries any coordinate $j \in [n]$ is at most $\|D_j f_i^*\|_{\text{sp}} \leq 1/n$. We also have $\mathbb{E}[\mathcal{A}_i(z)] = f_i^*(z)$. Therefore, by (4) and (5), we have

$$\begin{aligned} \mathbb{E}_{x \in \{-1, 1\}^k, i \in [k]} [\Pr[x_i = \mathcal{A}_i(C(x))]] &= \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x \in \{-1, 1\}^k, i \in [k]} [x_i \mathbb{E}[\mathcal{A}_i(C(x))]] \\ &= \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x \in \{-1, 1\}^k, i \in [k]} [x_i f_i^*(C(x))] \\ &= \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x \in \{-1, 1\}^k} \left[\left\| \frac{1}{k} \sum_{i=1}^k x_i f_i^* \right\|_{L_\infty} \right] \\ &\geq \frac{1}{2} + \frac{\varepsilon}{4}. \end{aligned}$$

Hence, C is a $(q, 1, \varepsilon/4)$ -average-case smooth code. \square

The final step before the proof of Theorem 1.4 is to show that for any distribution μ over smooth functions, there exists a distribution $\tilde{\mu}$ over smooth functions of bounded degree that is not much more concentrated than μ .

Lemma 3.2. *Let μ be a probability distribution over 1-smooth functions and let $\varepsilon > 0$. Then, there exists a probability distribution $\tilde{\mu}$ over 1-smooth functions of degree $q = 4/\varepsilon$ such that $\kappa_{\tilde{\mu}}(\varepsilon/2) \geq \kappa_{\mu}(\varepsilon)$.*

Proof: We first establish that smooth functions have low-degree approximations in the supremum norm. If $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is 1-smooth, then

$$q \sum_{|S| > q} |\widehat{f}(S)| \leq \sum_{S \subseteq [n]} |S| |\widehat{f}(S)| = \sum_{i=1}^n \sum_{S \ni i} |\widehat{f}(S)| = \sum_{i=1}^n \|D_i f\|_{\text{sp}} \leq 1.$$

It follows that the degree- q truncation $f^{\leq q}$ satisfies

$$\|f - f^{\leq q}\|_{L_\infty} \leq \sum_{|S| > q} |\widehat{f}(S)| \leq \frac{1}{q} = \frac{\varepsilon}{4}. \quad (6)$$

Define $\tilde{\mu}$ as follows: sample f according to μ and output $f^{\leq q}$. Clearly, $\tilde{\mu}$ is also a distribution over 1-smooth functions. For $k = \kappa_\mu(\varepsilon)$, we have

$$\mathbb{E}_{f_1, \dots, f_k \sim \mu} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - \mathbb{E}[f_i]) \right\|_{L_\infty} \right] \geq \varepsilon.$$

Hence, by the triangle inequality and (6), we have

$$\mathbb{E}_{f_1, \dots, f_k \sim \tilde{\mu}} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_i - \mathbb{E}[f_i]) \right\|_{L_\infty} \right] \geq \frac{\varepsilon}{2},$$

giving the claim. \square

Proof of Theorem 1.4: By applying Lemma 3.2 to μ , we get a distribution $\tilde{\mu}$ over 1-smooth degree $q = O(1/\varepsilon)$ functions with $k' = \kappa_{\tilde{\mu}}(\varepsilon/2) \geq \kappa_\mu(\varepsilon) = k$. By Theorem 3.1, we get a $(q, 1, \Omega(\varepsilon))$ -average-case smooth code $C' : \{-1, 1\}^{k'} \rightarrow \{-1, 1\}^n$. Finally we use Lemma 1.11 to convert C' to a $(q, \Omega(\varepsilon), \Omega(\varepsilon))$ -LDC $C : \{-1, 1\}^\ell \rightarrow \{-1, 1\}^n$ where $\ell = \Omega(\varepsilon^2 k' / \log(1/\varepsilon))$. For the last part of the theorem we can simply apply Theorem 3.1 directly. \square

4 From average-case smooth codes to LDCs

In this section, we prove Lemma 1.11. For this, we need the notion of the Vapnik–Chervonenkis dimension (VC-dimension).

Definition 4.1 (VC-dimension). For $T \subset [-1, 1]^k$ and $w > 0$, $\text{vc}(T, w)$ is defined as the size of the largest subset $\sigma \subset [k]$ such that there exists a shift $s \in [-1, 1]^k$ satisfying the following: for every $x \in \{-1, 1\}^\sigma$, there exists $t \in T$ such that for every $i \in \sigma$, $(t_i - s_i)x_i \geq w/2$.

Observe that if T is convex, then $\text{vc}(T, w)$ is the maximum dimension of a shifted hypercube with edge lengths at least w contained in T .

Definition 4.2 (Gaussian width). Let g be a k -dimensional standard Gaussian vector, with independent standard normal distributed entries. The *Gaussian width* of a set $T \subseteq \mathbb{R}^k$ is defined as

$$E(T) = \mathbb{E}_g[\sup_{t \in T} \langle g, t \rangle].$$

It is easy to see that a large VC-dimension implies a large Gaussian width. The following theorem shows the converse: containing a hypercube-like structure is the only way to have large Gaussian width.

Theorem 4.3 ([MV03]). *Let $T \subset [-1, 1]^k$. Then, the Gaussian width of T is bounded as*

$$E(T) \lesssim \sqrt{k} \int_{\alpha E(T)/k}^1 \sqrt{\text{vc}(T, w) \log(1/w)} dw$$

for some absolute constant $\alpha > 0$.

Finally, we use that fact that, as for LDCs, we can assume that on input $y \in \{0, 1\}^n$, the decoder \mathcal{A}_i of a smooth code first samples a set $S \subseteq [n]$ of at most q coordinates according to a probability distribution that depends on i only and then returns a random sign depending only on i , S and the values of y at S .

Proof of Lemma 1.11: The proof works by showing that the average-case smooth code property implies that the image of the (average) decoding functions should have large Gaussian width. We then use Theorem 4.3 to find a hypercube like structure inside the image, which we use to construct a smooth code. Finally we use Theorem 1.9 to convert the smooth code to an LDC.

Recall the switch of the message and codeword alphabets to $\{-1, 1\}$. For each $i \in [k]$, let $f_i : \{-1, 1\}^n \rightarrow [-1, 1]$ be the expected decoding function $f_i(z) = \mathbb{E}[\mathcal{A}_i(z)]$. Let g be a standard k -dimensional Gaussian vector and $T = \{(f_1(z), \dots, f_k(z)) : z \in \{-1, 1\}^n\}$. By the definition of average-case smooth code we have

$$2\eta k \leq \mathbb{E}_{x \in \{-1, 1\}^k} \left[\sum_{i=1}^k x_i f_i(C(x)) \right] \leq \mathbb{E}_{x \in \{-1, 1\}^k} \left[\sup_{t \in T} \langle x, t \rangle \right] \lesssim \mathbb{E}_g \left[\sup_{t \in T} \langle g, t \rangle \right].$$

(See for instance [Tal14, Lemma 3.2.10] for the last inequality.) By Theorem 4.3, for some constant $\alpha > 0$, we have

$$\eta k \lesssim \sqrt{k} \int_{\alpha\eta}^1 \sqrt{\text{vc}(T, w) \log(1/w)} dt \leq \sqrt{k} \cdot \sqrt{\text{vc}(T, \alpha\eta) \log(1/\alpha\eta)}$$

where we used the fact that $\text{vc}(T, w)$ is decreasing in w . So for $\tau = \alpha\eta$, we have $\text{vc}(T, \tau) \gtrsim \eta^2 k / \log(1/\eta)$. By the definition of VC-dimension, there exists a subset $\sigma \subset [k]$ of size $|\sigma| \geq \text{vc}(T, \tau)$ and a shift $s \in [-1, 1]^k$ such that for every $x \in \{-1, 1\}^\sigma$ there exists $t \in T$ such that $(t_i - s_i)x_i \geq \tau/2$ for every $i \in \sigma$.

Now we will define the code $C' : \{-1, 1\}^\sigma \rightarrow \{-1, 1\}^n$. Given $x \in \{-1, 1\}^\sigma$, there exists $t(x) \in T$ such that $(t(x)_i - s_i)x_i \geq \tau/2$ for every $i \in \sigma$. Define $C'(x) \in \{-1, 1\}^n$ to be one of the preimages of $t(x)$ under f , that is,

$$(f_1(C'(x)), \dots, f_k(C'(x))) = t(x).$$

Let W_p denote a $\{-1, 1\}$ -valued random variable with mean p . The decoding algorithms $\mathcal{A}'_i(y)$ run $\mathcal{A}_i(y)$ internally and give their output as follows:

$$\mathcal{A}'_i(y) = \begin{cases} \text{Output } W_{(1-s_i)/2} & \text{if } \mathcal{A}_i(y) \text{ returns } 1 \\ \text{Output } -W_{(1+s_i)/2} & \text{if } \mathcal{A}_i(y) \text{ returns } -1 \end{cases}$$

Therefore, for every $x \in \{-1, 1\}^\sigma$ and for every $i \in \sigma$,

$$\begin{aligned}
x_i \mathbb{E}[\mathcal{A}'_i(C'(x))] &= x_i \mathbb{E} \left[\frac{(1 + \mathcal{A}_i(C'(x)))}{2} W_{(1-s_i)/2} - \frac{(1 - \mathcal{A}_i(C'(x)))}{2} W_{(1+s_i)/2} \right] \\
&= \frac{x_i}{2} \mathbb{E}[\mathcal{A}_i(C'(x)) - s_i] \\
&= \frac{x_i}{2} (f_i(C'(x)) - s_i) \\
&= \frac{x_i}{2} (t(x)_i - s_i) \\
&\geq \frac{\tau}{4} \gtrsim \eta.
\end{aligned}$$

Since the probability that $\mathcal{A}'_i(C'(x))$ queries any particular location of $C'(x)$ is still at most c/n , it follows that C' is a $(q, c, \Omega(\eta))$ -smooth code. By Theorem 1.9, C' is also a $(q, \Omega(\eta/c), \Omega(\eta))$ -LDC. \square

5 From LDCs to outlaws

In this section we prove Theorem 1.5, the converse of our main result.

Proof of Theorem 1.5: By Theorem 1.9, the map $C : \{-1, 1\}^k \rightarrow \{-1, 1\}^n$ is also a $(q, q/\delta, \eta)$ -smooth code. For each $i \in [k]$, let \mathcal{B}_i be its decoder for the i th index. Let $\nu_i : 2^{[n]} \rightarrow [0, 1]$ be the probability distribution used by \mathcal{B}_i to sample a set $S \subseteq [n]$ of at most q coordinates and let $f_{i,S} : \{-1, 1\}^n \rightarrow [-1, 1]$ be function whose value at $y \in \{-1, 1\}^n$ is the expectation of the random sign returned by $\mathcal{B}_i(y)$ conditioned on the event that it samples S . Since this value depends only on the coordinates in S , the function $f_{i,S}$ is a q -junta.

Fix an $i \in [k]$ and let $f_i : \{-1, 1\}^n \rightarrow [-1, 1]$ be the function given by $f_i = \mathbb{E}_{S \sim \nu_i}[f_{i,S}]$. Then, since a q -junta has degree at most q , so does f_i . We claim that f_i is $\delta/(q2^{q/2})$ -smooth. Since the functions $f_{i,S} : \{-1, 1\}^n \rightarrow \{-1, 1\}$ are q -juntas, it follows from Parseval's identity that they have spectral norm at most $2^{q/2}$. Moreover, for each $j \in [n]$, we have $\Pr_{S \sim \nu_i}[j \in S] \leq q/(\delta n)$. Hence, since $f_{i,S}$ depends only on the coordinates in S , we have

$$\|D_j f_i\|_{\text{sp}} \leq \sum_{S \ni j} \nu_i(S) \|f_{i,S}\|_{\text{sp}} \leq \frac{q2^{q/2}}{\delta n},$$

which gives the claim. By (3), it holds for every $x \in \{-1, 1\}^k$ and every $i \in [k]$ that

$$x_i f_i(C'(x)) \geq 2\eta. \quad (7)$$

Define the distribution μ to correspond to the process of sampling $i \in [k]$ uniformly at random and returning f_i . Let $\bar{g} = (f_1 + \dots + f_k)/k$ be the mean of μ . We show that $\kappa_\mu(\eta) \geq \eta k$. To this end, let $l = \eta k$, let $\sigma : [l] \rightarrow [k]$ be an arbitrary map and define the

functions g_1, \dots, g_l by $g_i = f_{\sigma(i)}$. Let $x \in \{-1, 1\}^k$ be such that for each $i \in [l]$, we have $x_{\sigma(i)} = 1$ and $x_j = -1$ elsewhere. It follows from (7) that $f_{\sigma(i)}(C(x)) \in [2\eta, 1]$ for every $i \in [l]$ and that $f_i(C(x)) \leq 0$ for every other $i \in [k]$. Hence,

$$\begin{aligned} \left\| \frac{1}{l} \sum_{i=1}^l (g_i - \bar{g}) \right\|_{L_\infty} &\geq \left(\frac{1}{l} \sum_{i=1}^l (g_i - \bar{g}) \right)(C(x)) \\ &= \frac{1}{l} \sum_{i=1}^l f_{\sigma(i)}(C(x)) - \frac{1}{k} \sum_{i=1}^k f_i(C(x)) \\ &\geq 2\eta - \frac{l}{k} = \eta. \end{aligned}$$

If σ maps each element in $[l]$ to a uniformly random element in $[k]$, then g_1, \dots, g_l are independent, μ -distributed and satisfy

$$\mathbb{E} \left[\left\| \frac{1}{l} \sum_{i=1}^l (g_i - \bar{g}) \right\|_{L_\infty} \right] \geq \eta,$$

which shows that $\kappa_\mu(\eta) \geq l$. Finally we can scale all the functions in μ to make them 1-smooth, and get a distribution $\tilde{\mu}$ over 1-smooth functions with $\kappa_{\tilde{\mu}}(\eta\delta/(q2^{q/2})) \geq \eta k$. \square

6 Candidate outlaws

In this section we elaborate on the candidate outlaws mentioned in the introduction.

6.1 Incidence geometry

We begin by describing a variant of Corollary 1.6 based on a slightly different assumption and show conditions under which this assumption holds. Let p be an odd prime, let \mathbb{F}_p be a finite field with p elements and let n be a positive integer. For $x, y \in \mathbb{F}_p^n$, the *line* with origin x in direction y , denoted $\ell_{x,y}$, is the sequence $(x + \lambda y)_{\lambda \in \mathbb{F}_p}$. A line is nontrivial if $y \neq 0$.

Corollary 6.1. *For every odd prime p and $\varepsilon \in (0, 1]$, there exist a positive integer $n_1(p, \varepsilon)$ and a $c = c(p, \varepsilon) \in (0, 1/2]$ such that the following holds. Let $n \geq n_1(p, \varepsilon)$ and k be positive integers. Assume that for independent uniformly distributed elements $z_1, \dots, z_k \in \mathbb{F}_p^n$, with probability at least $1/2$, there exists a set $B \subseteq \mathbb{F}_p^n$ of size εp^n such that every nontrivial line through the set $\{z_1, \dots, z_k\}$ contains at most $p - 2$ points of B . Then, there exists a $(p - 1, c, c)$ -LDC sending $\{0, 1\}^l$ to $\{0, 1\}^{p^n}$, where $l = \Omega(c^2 k / \log(1/c))$.*

The proof uses the following version of Szemerédi's Theorem [Tao12, Theorem 1.5.4] and its standard ‘‘Varnavides-type’’ corollary (see for example [TV06, Exercise 10.1.9]).

Theorem 6.2 (Szemerédi's theorem). *For every odd prime p and any $\varepsilon \in (0, 1]$, there exists a positive integer $n_0(p, \varepsilon)$ such that the following holds. Let $n \geq n_0(p, \varepsilon)$ and let $S \subseteq \mathbb{F}_p^n$ be a set of size $|S| \geq \varepsilon p^n$. Then, S contains a nontrivial line.*

Corollary 6.3. *For every odd prime p and any $\varepsilon \in (0, 1]$, there exists a positive integer $n_1(p, \varepsilon)$ and a $c(p, \varepsilon) \in (0, 1]$ such that the following holds. Let $n \geq n_1(p, \varepsilon)$ and let $S \subseteq \mathbb{F}_p^n$ be a set of size $|S| \geq \varepsilon p^n$. Then, S contains at least $c(p, \varepsilon)p^{2n}$ nontrivial lines, that is,*

$$\Pr_{x \in \mathbb{F}_p^n, y \in \mathbb{F}_p^n \setminus \{0\}} [\{(x + \lambda y)_{\lambda=0}^{p-1}\} \subset S] \geq c(p, \varepsilon).$$

Proof. Abusing notation, we identify functions $f : \mathbb{F}_p^n \rightarrow \{-1, 1\}$ with vectors in $\{-1, 1\}^{\mathbb{F}_p^n}$. Let $\phi : \{-1, 1\} \rightarrow \{0, 1\}$ be the map $\phi(\alpha) = (\alpha + 1)/2$. For a function $f : \mathbb{F}_p^n \rightarrow \{-1, 1\}$, let $\phi(f) : \mathbb{F}_p^n \rightarrow \{0, 1\}$ be the function $\phi(f)(x) = \phi(f(x))$ and for $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$, define $\phi^{-1}(f) : \mathbb{F}_p^n \rightarrow \{-1, 1\}$ analogously.

For every $x \in \mathbb{F}_p^n$, let $F_x : \{-1, 1\}^{\mathbb{F}_p^n} \rightarrow \mathbb{R}$ be the degree- $(p-1)$ function

$$F_x(f) = \mathbb{E}_{y \in \mathbb{F}_p^n \setminus \{0\}} \left[\prod_{\lambda \in \mathbb{F}_p^*} \phi(f)(x + \lambda y) \right]. \quad (8)$$

Then, for a set $B \subseteq \mathbb{F}_p^n$, the value $F_x(\phi^{-1}(1_B))$ equals the fraction of all nontrivial lines $\ell_{x,y}$ through x of which B contains the $p-1$ points $\{x + \lambda y : \lambda \in \mathbb{F}_p^*\}$. If B has size at least εp^n , it follows from Corollary 6.3 that $\mathbb{E}_{x \in \mathbb{F}_p^n} [F_x(\phi^{-1}(1_B))] \geq c(p, \varepsilon)$. Moreover, since the monomials in the expectation of (8) involve disjoint sets of variables and can be expanded as

$$\prod_{\lambda \in \mathbb{F}_p^*} \phi(f)(x + \lambda y) = \frac{1}{2^q} \sum_{S \subseteq \mathbb{F}_p^*} \prod_{\lambda \in S} f(x + \lambda y),$$

it follows that each F_x is $2(1 - p^{-n})$ -smooth.

Let μ be the uniform distribution over F_x . We claim that $\kappa_\mu(c(p, \varepsilon)) \geq k$, which implies the result by Theorem 1.4 since μ is supported by degree $(p-1)$ -functions. For every set $A \subseteq \mathbb{F}_p^n$, let $B_A \subseteq \mathbb{F}_p^n$ be a maximal set such that every nontrivial line through A contains at most $p-2$ points of B_A , and let $f_A = \phi^{-1}(1_{B_A})$. Let z be a uniformly distributed random variable over \mathbb{F}_p^n , let z_1, \dots, z_k be independent copies of z and let $A = \{z_1, \dots, z_k\}$. Then, F_{z_1}, \dots, F_{z_k} are independent μ -distributed random functions. Moreover, in the event that both $|B_A| \geq \varepsilon p^n$ and every nontrivial line through A meets B_A in at most $p-2$ points, we have

$$|(F_{z_i} - \mathbb{E}[F_z])(f_A)| = \mathbb{E}[F_z(\phi^{-1}(1_{B_A}))] - F_{z_i}(\phi^{-1}(1_{B_A})) \geq c(p, \varepsilon)$$

for every $i \in [k]$. Since this event happens with probability at least $1/2$, we have

$$\mathbb{E} \left[\left\| \frac{1}{k} \sum_{i=1}^k (F_{z_i} - \mathbb{E}[F_z]) \right\|_{L_\infty} \right] \geq \mathbb{E} \left[\left\| \frac{1}{k} \left(\sum_{i=1}^k (F_{z_i} - \mathbb{E}[F_z]) \right) (f_A) \right\| \right] \geq \frac{c(p, \varepsilon)}{2},$$

which gives the claim. \square

The proof of the formal version of Corollary 1.6 (given below) is similar to that of Corollary 6.1, so we omit it. In the following, $\mathbb{P}\mathbb{F}_p^{n-1}$ is the projective space of dimension $n-1$, which is the space of directions in \mathbb{F}_p^n . The formal version of Corollary 1.6 is then as follows.

Corollary 6.4. *For every odd prime p and $\varepsilon \in (0, 1]$, there exist a positive integer $n_1(p, \varepsilon)$ and a $c = c(p, \varepsilon) \in (0, 1/2]$ such that the following holds. Let $n \geq n_1(p, \varepsilon)$ and k be positive integers. Suppose that for independent uniformly distributed elements $z_1, \dots, z_k \in \mathbb{F}_p^{n-1}$, with probability at least $1/2$, there exists a set $B \subset \mathbb{F}_p^n$ of size $|B| \geq \varepsilon p^n$ which does not contain any lines with direction in $\{z_1, \dots, z_k\}$. Then, there exists a (p, c, c) -LDC sending $\{0, 1\}^l$ to $\{0, 1\}^{p^n}$, where $l = \Omega(c^2 k / \log(1/c))$.*

Feasible parameters for Corollary 6.1 Proving lower bounds on k for which the assumption of Corollary 6.1 holds true thus allows one to infer the existence of $(p-1)$ -query LDCs with rate $\Omega(k/N)$ for $N = p^n$, provided p and ε are constant with respect to n . We establish the following bounds, which imply the (well-known) existence of $(p-1)$ -query LDCs with message length $k = \Omega((\log N)^{p-2})$.

Theorem 6.5. *For every odd prime p there exists an $\varepsilon(p) \in (0, 1]$ such that the following holds. For every set $A \subseteq \mathbb{F}_p^n$ of size $|A| \leq \binom{n+p-3}{p-2} - 1$, there exists a set $B \subseteq \mathbb{F}_p^n$ of size $\varepsilon(p)p^n$ such that every line through A contains at most $p-2$ points of B .*

The proof uses some basic properties of polynomials over finite fields. For an n -variate polynomial $f \in \mathbb{F}_p[x_1, \dots, x_n]$ denote $Z(f) = \{x \in \mathbb{F}_p^n : f(x) = 0\}$. The starting point of the proof is the following standard result (see for example [Tao14]), showing that small sets can be ‘captured’ by zero-sets of nonzero, homogeneous polynomials of low degree.

Lemma 6.6 (Homogeneous Interpolation). *For every $A \subseteq \mathbb{F}_p^n$ of size $|A| \leq \binom{n+d-1}{d} - 1$, there exists a nonzero homogeneous polynomial $f \in \mathbb{F}_p[x_1, \dots, x_n]$ of degree d such that $A \subseteq Z(f)$.*

The next two lemmas show that if f is nonzero, homogeneous and degree d , and if $a \in \mathbb{F}_p^*$ is such that $f^{-1}(a)$ is nonempty, then lines through $Z(f)$ meet $f^{-1}(a)$ in at most d points.

Lemma 6.7. *Let $f \in \mathbb{F}_p[x_1, \dots, x_n]$ be a nonzero homogeneous polynomial of degree d . Let $a \in \mathbb{F}_p^*$ be such that the set $f^{-1}(a)$ is nonempty. Then, every line that meets $f^{-1}(a)$ in $d+1$ points must have direction in $Z(f)$.*

Proof: The univariate polynomial $g(\lambda) = f(x + \lambda y)$ formed by the restriction of f to a line $\ell_{x,y}$ has degree at most d . By the Factor Theorem, such a polynomial must be the constant polynomial $g(\lambda) = a$ to assume the value a for $d+1$ values of λ . Since f is homogeneous, the coefficient of λ^d , which must be zero, equals $f(y)$, giving the result. \square

The following lemma is essentially contained in [BR16].

Lemma 6.8 (Briët–Rao). *Let $f \in \mathbb{F}_p[x_1, \dots, x_n]$ be a nonzero homogeneous polynomial of degree d . Let $a \in \mathbb{F}_p^*$ be such that $f^{-1}(a)$ is nonempty. Then, there exists no line that intersects $Z(f)$, meets $f^{-1}(a)$ in at least d points and has direction in $Z(f)$.*

Proof: For a contradiction, suppose there exists a line $\ell_{x,y}$ through $Z(f)$ that meets $f^{-1}(a)$ in d points and has direction $y \in Z(f)$. Observe that for every $\lambda \in \mathbb{F}_p$, the shifted line $\ell_{x+\lambda y, y}$ also meets $f^{-1}(a)$ in d points. Hence, without loss of generality we may assume that the line starts in $Z(f)$, that is $x \in Z(f)$. Let $g(\lambda) = a_0 + a_1\lambda + \dots + a_d\lambda^d = f(x + \lambda y) \in \mathbb{F}_p[\lambda]$ be the restriction of f to $\ell_{x,y}$. It follows that $a_0 = g(0) = f(x) = 0$ and, since f is homogeneous, that $a_d = f(y) = 0$. Moreover, there exist distinct elements $\lambda_1, \dots, \lambda_d \in \mathbb{F}_p^*$ such that $g(\lambda_i) = f(x + \lambda_i y) = a$ for every $i \in [d]$. Then $g(\lambda) - a$ is a degree $d - 1$ polynomial with d distinct roots. But it cannot be the zero polynomial since it takes value $-a$ when $\lambda = 0$. \square

The final ingredient for the proof of Theorem 6.5 is the DeMillo–Lipton–Schwartz–Zippel Lemma, as it appears in [CT14].

Lemma 6.9 (DeMillo–Lipton–Schwartz–Zippel). *Let $f \in \mathbb{F}_p[x_1, \dots, x_n]$ be a nonzero polynomial of degree d and denote $r = |\mathbb{F}_p|$. Then,*

$$|Z(f)| \leq \left(1 - \frac{1}{r^{d/(r-1)}}\right) r^n.$$

ProofProof of Theorem 6.5: Let $A \subseteq \mathbb{F}_p^n$ be a set of size $|A| \leq \binom{n+p-3}{p-2} - 1$. Let $f \in \mathbb{F}_p[x_1, \dots, x_n]$ be a nonzero degree- $(p-2)$ homogeneous polynomial such that $A \subseteq Z(f)$, as promised to exist by Lemma 6.6. By Lemma 6.9, there exists an $a \in \mathbb{F}_p^*$ such that the set $B = f^{-1}(a)$ has size at least $|B| \geq p^n / p^{(2p-3)/(p-1)}$. By Lemma 6.7, every line that meets B in $p-1$ points must have direction in $Z(f)$, but by Lemma 6.8 no such line can pass through $Z(f)$. Hence, every line through A meets B in at most $p-2$ points. \square

6.2 Hypergraph pseudorandomness

A second candidate for constructing outlaws comes from special types of hypergraphs. A hypergraph $H = (V, E)$ is a pair consisting of a finite vertex set V and an edge set E of subsets of V that allows for parallel (repeated) edges. A hypergraph is t -uniform if all its edges have size t . For subsets $W_1, \dots, W_t \subseteq V$, define the induced edge count by

$$e_H(W_1, \dots, W_t) = \sum_{v_1 \in W_1} \dots \sum_{v_t \in W_t} 1_E(\{v_1, \dots, v_t\}).$$

A perfect matching in a t -uniform hypergraph is a family of vertex-disjoint edges that intersects every vertex. We shall use the following notion of pseudorandomness.

Definition 6.10 (Relative pseudorandomness). Let $H = (V, E)$, $J = (V, E')$ be t -uniform hypergraphs with identical vertex sets. Then J is ε -pseudorandom relative to H if for all $W_1, \dots, W_t \subseteq V$, we have

$$\left| \frac{e_J(W_1, \dots, W_t)}{|E'|} - \frac{e_H(W_1, \dots, W_t)}{|E|} \right| < \varepsilon. \quad (9)$$

The left-hand side of (9) compares the fraction of edges that the sets W_1, \dots, W_t induce in J with the fraction of edges they induce in H . Standard concentration arguments show that if $|E| \geq |V|$, then a random hypergraph J whose edge set E' is formed by independently putting each edge of E in E' with probability $p = p(\varepsilon, t)$, is ε -pseudorandom relative to H with high probability. A deterministic hypergraph J is thus pseudorandom relative to H if it mimics this property of truly random sub-hypergraphs. For graphs, relative ε -pseudorandomness turns into a common notion sometimes referred to as ε -uniformity when H is the complete graph with all loops, in which case (9) says that the number of edges induced by a pair of vertex-subsets W_1, W_2 is roughly equal to the product of their densities $(|W_1|/|V|)(|W_2|/|V|)$. Uniformity in graphs is closely connected to the perhaps better-known notion of spectral expansion [HLW06]. These two notions were recently shown to be equivalent (up-to universal constants) for all vertex-transitive graphs [CZ17].

We shall be interested in hypergraphs whose edge set can be partitioned into a family of “blocks”, such that randomly removing relatively few of the blocks likely leaves a hypergraph that is *not* pseudorandom relative to the original. (Think of a Jenga tower⁵ that’s already in a delicate balance, so that there are only few ways, or perhaps even no way, to remove many blocks without having it collapse.) Our blocks will be formed by perfect matchings. For technical reasons, the formal definition takes the view of building a new hypergraph out of randomly selected matchings, as opposed to obtaining one by randomly removing matchings.

Definition 6.11 (Jenga hypergraph). A t -uniform hypergraph H is (k, ε) -jenga if its edge set can be partitioned into a family \mathcal{M} of perfect matchings such that, with probability at least $1/2$, the disjoint union of k independent uniformly distributed matchings from \mathcal{M} forms a hypergraph is *not* ε -pseudorandom relative to H .

We have the following simple corollary to Theorem 1.4.

Corollary 6.12. *Let n, k, t be positive integers and $\varepsilon \in (0, 1]$. Assume that there exists a t -uniform n -vertex hypergraph that is (k, ε) -jenga. Then, there exists a $(t, 1, \Omega(\varepsilon/t^2))$ -LDC sending $\{0, 1\}^l$ to $\{0, 1\}^{tn}$, where $l = \Omega(\varepsilon^2 k/t^4 \log(t^2/\varepsilon))$.*

Proof: Let $H = (V, E)$ be a hypergraph as assumed in the corollary. Let \mathcal{M} be a partition of E into perfect matchings such that if M_1, \dots, M_k are independent and uniformly distributed over \mathcal{M} , then with probability at least $1/2$, the hypergraph $J = (V, M_1 \uplus \dots \uplus M_k)$ is not ε -pseudorandom relative to H .

Let V_1, \dots, V_t be copies of V . For each $M \in \mathcal{M}$, define $f_M : \mathbb{R}^{V_1 \cup \dots \cup V_t} \rightarrow \mathbb{R}$ by

$$f_M(x[1], \dots, x[t]) = \frac{1}{|M|} \sum_{v_1 \in V_1} \dots \sum_{v_t \in V_t} 1_M(\{v_1, \dots, v_t\}) x[1]_{v_1} \dots x[t]_{v_t}, \quad x[i] \in \mathbb{R}^{V_i}.$$

The function f_M is a degree- t polynomial. Since every one of the tn variables appears in exactly one monomial and $|M| = n/t$, the restriction of f_M to $\{-1, 1\}^{V_1 \cup \dots \cup V_t}$ is t^2 -smooth.

⁵Jenga[®] is a game of dexterity in which players begin with a tower of wooden blocks and take turns trying to remove a block without making the tower collapse.

Moreover, for $J = (V, M)$ and $W_1, \dots, W_t \subseteq V$, we have

$$f_M(1_{W_1}, \dots, 1_{W_t}) = \frac{e_J(W_1, \dots, W_t)}{|M|}.$$

Let M_1, \dots, M_k be independent uniformly distributed matchings from \mathcal{M} and consider the random hypergraph $J = (V, M_1 \uplus \dots \uplus M_k)$. Let $\bar{f} = \mathbb{E}[f_{M_1}]$ be the expectation of the random function f_{M_1} and note that $\mathbb{E}[f_{M_i}] = \bar{f}$ for each $i \in [k]$. Then, since the functions $f_{M_i} - \bar{f}$ are multilinear,

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{1}{k} \sum_{i=1}^k (f_{M_i} - \bar{f}) \right\|_{L_\infty} \right] &\geq \mathbb{E} \left[\max_{W_1, \dots, W_t \subseteq V} \left| \frac{1}{k} \sum_{i=1}^k (f_{M_i} - \bar{f})(1_{W_1}, \dots, 1_{W_t}) \right| \right] \\ &= \mathbb{E} \left[\max_{W_1, \dots, W_t \subseteq V} \left| \frac{e_J(W_1, \dots, W_t)}{k|M|} - \frac{e_H(W_1, \dots, W_t)}{|E|} \right| \right] \\ &\geq \frac{\varepsilon}{2}. \end{aligned}$$

The result now follows from Theorem 1.4. \square

In the context of outlaws and LDCs, the relevant question concerning Jenga hypergraphs is the following. Let $\kappa^J(n, t, \varepsilon)$ denote the maximum integer k such that there exists an n -vertex t -uniform hypergraph that is (k, ε) -jenga.

Question 6.13. *For integer $t \geq 2$ and parameter $\varepsilon \in (0, 1]$, what is the growth rate of $\kappa^J(n, t, \varepsilon)$ as a function of $n \in \mathbb{N}$?*

For $t = 2$ (graphs), the answer to Question 6.13 follows from famous work of Alon and Roichman [AR94] on expansion of random Cayley graphs, which implies that for constant $\varepsilon \in (0, 1]$, we have $\kappa(n, 2, \varepsilon) = \Theta(\log n)$. The lower bound follows for instance by partitioning the edge set of the complete graph with vertex set $V = \mathbb{F}_2^m$ into the collection of matchings of the form $M_y = \{\{x, x+y\} : x \in \mathbb{F}_2^m\}$ for each $y \in \mathbb{F}_2^m \setminus \{0\}$. Any $m-1$ of such matchings give a graph with two disconnected components of equal size, making it $(m-1, \frac{1}{4})$ -jenga. Via Corollary 6.12, this arguably gives the most round-about way to prove the existence of 2-query LDCs matching the parameters of the Hadamard code! Generalizing the above example, [BR16] considered the p -uniform hypergraph on \mathbb{F}_p^m whose edges are the (unordered) nontrivial lines. It was shown that this hypergraph is (m^{p-1}, ε) -jenga for some $\varepsilon = \varepsilon(p)$ depending on p only, by partitioning the edge set according to the directions of the lines, that is, partitioning it with the matchings $M_y = \{\{x+\lambda y : \lambda \in \mathbb{F}_p\} : x \in \mathbb{F}_p^m\}$, $y \in \mathbb{F}_p^m \setminus \{0\}$. To the best of our knowledge, the best upper bounds on $\kappa^J(n, t, \varepsilon)$ for constant $t \geq 3$ and $\varepsilon \in (0, 1]$ follow from upper bounds on LDCs, via Corollary 6.12.

We end with the following natural question concerning Jenga hypergraphs.

Question 6.14. *Is $\kappa^J(n, t, \varepsilon)$ largest for the complete hypergraph?*

Acknowledgements J.B. and S.G. thank the Simons Institute for hosting them during the 2017 Pseudorandomness program, where part of this work was done.

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
- [AR94] Noga Alon and Yuval Roichman. Random Cayley graphs and expanders. *Random Structures & Algorithms*, 5, 1994.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
- [BDG17] Jop Briët, Zeev Dvir, and Sivakanth Gopi. Outlaw distributions and locally decodable codes. In *Proceedings of the 8th Annual Innovations in Theoretical Computer Science conference (ITCS 2017)*, 2017. To appear.
- [BK95] Manuel Blum and Sampath Kannan. Designing programs that check their work. *Journal of the ACM*, 42(1):269–291, 1995.
- [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.
- [BR16] Jop Briët and Shravas Rao. Arithmetic expanders and deviation bounds for random tensors. *arXiv preprint arXiv:1610.03428*, 2016.
- [Bri16] Jop Briët. On embeddings of ℓ_1^k from locally decodable codes. *arXiv preprint: arXiv:1611.06385*, 2016.
- [CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998.
- [CT14] Gil Cohen and Avishay Tal. Two structural results for low degree polynomials and applications. *arXiv preprint arXiv:1404.0654*, 2014.
- [CZ17] David Conlon and Yufei Zhao. Quasirandom Cayley graphs. *Discrete Analysis*, 6, 2017. Available at arXiv:1603.03025 [math.CO].
- [DG15] Zeev Dvir and Sivakanth Gopi. 2-Server PIR with sub-polynomial communication. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 577–584. ACM, 2015.
- [Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the forty-first annual ACM symposium on Theory of computing (STOC 2009)*, pages 39–44, 2009.

- [FLW16] Nikos Frantzikinakis, Emmanuel Lesigne, and Mate Wierdl. Random differences in szemerédi’s theorem and related results. *Journal d’Analyse Mathématique*, 130(1):91–133, 2016.
- [Fox] Jacob Fox. Personal communication.
- [GHSY12] Parikshit Gopalan, Cheng Huang, Huseyin Simitci, and Sergey Yekhanin. On the locality of codeword symbols. *IEEE Trans. Information Theory*, 58(11):6925–6934, 2012.
- [GKO⁺17] Sivakanth Gopi, Swastik Kopparty, Rafael Oliveira, Noga Ron-Zewi, and Shubhangi Saraf. Locally testable and locally correctable codes approaching the gilbert-varshamov bound. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*, pages 2073–2091. SIAM, 2017.
- [GKST02] Oded Goldreich, Howard Karloff, Leonard J Schulman, and Luca Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. In *Computational Complexity, 2002. Proceedings. 17th IEEE Annual Conference on*, pages 143–151. IEEE, 2002.
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43:439–561, 2006.
- [KMRS16] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High-rate locally-correctable and locally-testable codes with sub-polynomial query complexity. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 202–215, 2016.
- [KS07] Tali Kaufman and Madhu Sudan. Sparse random linear codes are locally decodable and testable. In *Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on*, pages 590–600. IEEE, 2007.
- [KSY14] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. *Journal of the ACM (JACM)*, 61(5):28, 2014.
- [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.
- [KW04] 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.
- [MV03] Shachar Mendelson and Roman Vershynin. Entropy and the combinatorial dimension. *Invent. Math.*, 152(1):37–55, 2003.

- [Sha48] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [Tal14] Michel Talagrand. *Upper and lower bounds for stochastic processes*, volume 60 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]*. Springer, Heidelberg, 2014. Modern methods and classical problems.
- [Tao12] Terence Tao. *Higher order Fourier analysis*, volume 142. American Mathematical Society, 2012.
- [Tao14] Terence Tao. Algebraic combinatorial geometry: the polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory. *EMS Surv. Math. Sci.*, 1:1–46, 2014.
- [TV06] Terence Tao and Van Vu. *Additive Combinatorics*. Cambridge University Press, 2006.
- [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.