Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 53 (2013)

High rate locally correctable codes via lifting

Alan Guo *

Abstract

We present a general framework for constructing high rate error correcting codes that are locally correctable (and hence locally decodable if linear) with a sublinear number of queries, based on lifting codes with respect to functions on the coordinates. Our approach generalizes the lifting of affine-invariant codes of Guo, Kopparty, and Sudan and its generalization automorphic lifting, suggested by Ben-Sasson et al, which lifts algebraic geometry codes with respect to a group of automorphisms of the code. Our notion of lifting is a natural alternative to the degree-lifting of Ben-Sasson et al and it carries two advantages. First, it overcomes the rate barrier inherent in degree-lifting. Second, it is extremely flexible, requiring no special properties (e.g. linearity, invariance) of the base code, and requiring very little structure on the set of functions on the coordinates of the code.

As an application, we construct new explicit families of locally correctable codes by lifting algebraic geometry codes. Like the multiplicity codes of Kopparty, Saraf, Yekhanin and the affine-lifted codes of Guo, Kopparty, Sudan, our codes of block-length N can achieve N^{ϵ} query complexity and $1-\alpha$ rate for any given $\epsilon, \alpha > 0$ while correcting a constant fraction of errors, in contrast to the Reed-Muller codes and the degree-lifted AG codes of Ben-Sasson et al which face a rate barrier of $\epsilon^{O(1/\epsilon)}$. However, like the degree-lifted AG codes, our codes are over an alphabet significantly smaller than that obtained by Reed-Muller codes, affine-lifted codes, and multiplicity codes.

*CSAIL, Massachusetts Institute of Technology, 32 Vassar Street, Cambridge, MA, USA. aguo@mit.edu. Research supported in part by NSF grants CCF-0829672, CCF-1065125, and CCF-6922462, and an NSF Graduate Research Fellowship

ISSN 1433-8092

Contents

1 Introduction

We present a general framework for constructing long locally correctable codes from short base codes via the operation of lifting. Our notion of lifting generalizes affine lifting, automorphic lifting, and high-degree sampling defined in previous works, and we use it to obtain new explicit high rate locally correctable codes by lifting certain algebraic geometric codes.

1.1 Error correcting codes and locally correctable codes

We begin with some coding theory preliminaries. A *code* $\mathcal{C}$ of block length N over an alphabet R is a subset of R^N. Elements $f \in \mathcal{C}$ are *codewords*. Typically Σ is used to denote the alphabet, but we use R because it is helpful to think of a codeword f not as a vector in R^N, but as a function $f : D \to R$ (D for domain, R for range), where we identify D with $[N] = \{1, \dots, N\}$. Typically one thinks of $\mathcal{C}$ as the image of some *encoding map* $\mathsf{Enc} : R_0^K \to R^N$ which injectively maps K-symbol *messages* over an alphabet R_0 to N-symbol codewords (here R_0 may be different from R). The *rate* of the code $\mathcal{C}$ is K/N, which measures the efficiency of our encoding. We want to make K/N as large as we can. Another important parameter of a code is the minimum pairwise distance between distinct codewords. The *(Hamming) distance* between two words $f, g \in R^N$ is the number of coordinates in which they differ, i.e.

$$\Delta(f, g) \triangleq \{i \in [N] \mid x_i \neq y_i\}.$$

The distance $\Delta(\mathcal{C})$ of $\mathcal{C}$ is simply $\min\{\Delta(f, g) \mid f, g \in \mathcal{C},\ f \neq g\}$. We want $\Delta(\mathcal{C})$ to be as large as possible. We often look at the normalized distance $\delta(f, g)$, which is simply $\frac{1}{N}\Delta(f, g)$, and similarly $\delta(\mathcal{C}) = \frac{1}{N}\Delta(\mathcal{C})$.

The motivation behind error correcting codes is to make information robust to noise. The original message $m \in R_0^K$ is encoded into some codeword $\mathsf{Enc}(m) \in R^N$. Noise may corrupt some symbols of $\mathsf{Enc}(m)$, resulting in a new word $r \in R^N$, the *received word*. The number of symbols corrupted is exactly $\Delta(\mathsf{Enc}(m), r)$. If the number of errors is small, say less than $\Delta(\mathcal{C})/2$, then $\mathsf{Enc}(m)$ is the unique codeword in $\mathcal{C}$ within Hamming distance $\Delta(\mathcal{C})/2$ of r, and one can uniquely decode r to get m, since Enc is injective.

To decode a received word, it may be necessary to examine the entire word. In some settings, the received word is prohibitively large, and one wishes only to decode one symbol of the message. Codes with which one can do this by querying only a small number of symbols of the input are known as *locally decodable codes*. A related concept is the notion of *locally correctable code*. Such a code allows one to correct a symbol of the *codeword* (rather than a symbol of the message) by querying only a few symbols of the input. The main parameters of interest are the rate and the query complexity, or locality, the number of symbols queried to recover a single symbol. These codes are the focus of this work. We formally define these notions in Section 2.

1.2 Previous work

Until recently, there were no known locally correctable codes with sublinear query complexity and rate greater 1/2. The Reed-Muller code was the first locally correctable code, with the first correction procedure proposed by Reed [5], which happened to be a local correction procedure. The m-variate Reed-Muller over $\mathbb{F}_q$ with degree parameter r consists of all m-variate polynomials of total degree less than r. More precisely, a codeword is the list of evaluations of such a polynomial on all points of $\mathbb{F}_q^m$. The idea behind the local correction procedure is to pick a random line passing through the point whose value we wish to correct, view the restriction of the polynomial to the line as a corrupted Reed-Solomon codeword, and use a Reed-Solomon decoding algorithm to correct the value on the line. For any $\epsilon > 0$, the Reed-Muller codes can achieve query complexity N^ϵ by taking $m = 1/\epsilon$ and $N = q^m$. Unfortunately, the m-variate Reed-Muller code with positive distance (by taking r to be a constant fraction of q) can never exceed $1/m!$ in rate. This certainly never exceeds $1/2$.

The recent work of Kopparty, Saraf, and Yekhanin [4] introduced the first locally correctable codes that can achieve rate greater than $1/2$, and in fact can achieve any rate arbitrarily close to 1. More precisely, for any $\epsilon, \alpha > 0$, the multiplicity code can achieve query complexity N^ϵ and rate $1-\alpha$ while correcting a constant fraction of errors. One may view multiplicity codes as a variant of Reed-Muller codes, where each codeword consists of evaluations of a low-degree polynomial along with its partial derivatives.

An alternative to the multiplicity codes are the lifted Reed-Solomon codes of Guo, Kopparty, and Sudan [3]. These are yet another variant of Reed-Muller codes — more precisely, they are supercodes of Reed-Muller codes with vastly greater dimension but the same distance. The main idea behind lifted codes is the notion of "lifting" — an operation first introduced in [2] to prove negative results in property testing. Essentially, the lifting operation takes a short base code $\mathcal{C} \subseteq \{\mathbb{F}_q^t \to \mathbb{F}_q\}$ and "lifts" it to a longer code $\mathcal{C}' \subseteq \{\mathbb{F}_q^m \to \mathbb{F}_q\}$, for $m > t$, such that codewords of $\mathcal{C}'$ are those $f : \mathbb{F}_q^m \to \mathbb{F}_q$ whose restriction to every t-dimension affine subspace is a codeword of $\mathcal{C}$. Guo et al [3] obtain locally correctable codes with query complexity N^ϵ and rate $1 - \alpha$ by lifting the Reed-Solomon code. Our work generalizes this notion of lifting.

The work of Ben-Sasson et al [1] presents another way to build long locally correctable codes from short base codes via the "degree-lifting" operation. Degree-lifting abstracts the process of obtaining the Reed-Muller codes from the Reed-Solomon code and applies it to algebraic geometry codes. By degree-lifting certain algebraic geometry codes, such as the Hermitian code, Ben-Sasson et al obtain locally correctable codes with Reed-Muller-like properties but significantly smaller alphabet. Unfortunately, degree-lifting faces the same rate barrier that the Reed-Muller codes face, for essentially the same reason. Two key contributions of [1] which we use in our work are the notions of a group being "close" to doubly transitive, and the fractal correction algorithm. In particular, a conceptual contribution of [1] is the observation that the "uniformity" of the automorphism group of an algebraic geometry code yields good local correctability properties. Our work generalizes this observation. Ben-Sasson et al also suggests the idea of "automorphic lifting", a natural generalization of the affine lifting of [3] to apply to algebraic geometry codes. Our work further generalizes this idea. Moreover, our notion of lifting encapsulates the notion of high-degree sampling used in [1] as well. The idea of high-degree sampling is to restrict not to automorphisms, but to "high-degree views". For instance, instead of restricting to lines to decode the Reed-Muller code, one may restrict to curves parametrized by quadratic equations.

1.3 Our results

In this work, we introduce a lifting framework which abstracts the lifting operation used by [3] and the automorphic lifting suggested by [1] as well as the high-degree restrictions used by [1]. Our framework applies to *arbitrary codes* and *arbitrary sets of functions* (as opposed to invariant codes under some group of (generalized) automorphisms). In particular, unlike the degree-lifting operation of [1], our lifting operation does not require an algebraic notion of "degree". Informally, our lifting operation is defined as follows. Let Φ be a set of functions from $D \to D$. The m-variate lift of $\mathcal{C} \subseteq \{D \to R\}$ with respect to Φ is the code whose codewords are those $f : D^m \to R$ such that the univariate function $f(\sigma_1(x), \ldots, \sigma_m(x))$ is a codeword of $\mathcal{C}$ for all $(\sigma_1, \ldots, \sigma_m) \in \Phi^m$. For affine-lifting, the domain is $D = \mathbb{F}_q$ and Φ is the group of affine permutations on $\mathbb{F}_q$, and in [3] the base code is taken to be affine-invariant. More generally, for automorphic lifting, Φ is some group of automorphisms on D under which $\mathcal{C}$ is invariant. Our definition of lifting requires neither $\mathcal{C}$ to be Φ-invariant, nor even Φ to be a group.

A conceptual contribution of our work is to show that if Φ is sufficiently close to uniform in the sense of Ben-Sasson et al [1], then this suffices for the lift to have good distance and be locally correctable. We show that there is nothing essential about the symmetry of the base code under Φ, nor the fact that Φ is a group. Thus, designing good lifted codes "merely" involves choosing a good set Φ with respect to which to lift. On the one hand, including too many functions in Φ kills the rate of the lifted code, since every function adds a constraint on the lifted code. On the other hand, including too few functions in Φ kills the distance of the lifted code, since we want enough functions in Φ to make it "close" to doubly transitive.

As an application, we construct an explicit family of locally correctable codes via lifting. The family arises from lifting the Hermitian code, the algebraic geometry code that [1] degree-lift. We obtain high rate locally correctable codes similar to the lifted Reed-Solomon codes, except over a significantly smaller alphabet.

Though our explicit construction uses algebraic geometry codes as base codes, our exposition is elementary and self-contained. Invoking the language of algebraic function field theory is only necessary to prove the properties of the base codes; the properties themselves can be stated in elementary terms, and we do so. We refer the interested reader who wishes to see the proofs of these facts to the book of Stichtenoth [6] on algebraic function fields and codes.

1.4 Comparison of parameters

We compare the parameters of the constant rate locally correctable codes found in the literature, including the ones constructed in this paper. We start with some easy comparisons. The lifted Reed-Solomon code of Guo, Kopparty, Sudan [3] is strictly better than the Reed-Muller code, as it is a strict supercode with the same distance. In fact, with m variables over $\mathbb{F}_q$, the two codes have the same length, alphabet, and query complexity, but the rate of Reed-Muller is bounded above by $\frac{1}{m!}$ (even as its distance goes to 0) whereas the rate of the lifted Reed-Solomon code approaches 1 as its distance goes to 0. Similarly, the lifted Hermitian code (Theorem 7.1) has the same length, alphabet, and query complexity as that of the degree-lifted Hermitian code of Ben-Sasson et al [1], but the rate of the degree-lifted Hermitian code is bounded above by $\frac{1}{m!}$ whereas the rate of the lifted Hermitian code approaches 1 as its distance goes to 0.

To compare the various families of high rate locally correctable codes, we normalize their parameters. Namely, we fix the block length to N, the rate to $1 - \alpha$, query complexity to N^ϵ, and compare the alphabet size and error correcting rate of each code. The results are summarized in the table below.

Code	Alphabet size	Error correcting rate
Multiplicity [4]	$N^{\Omega((1/\epsilon)^{(1/\epsilon)})}$	$\Omega(\epsilon\alpha)$
Lifted Reed-Solomon [3]	N^ϵ	$\alpha^{O((2/\epsilon)^{(1/\epsilon)}\log(1/\epsilon))}$
Lifted Hermitian (Theorem 7.1)	$N^{\epsilon/3}$	$\alpha^{O((8/\epsilon)^{(2/\epsilon)}\log(1/\epsilon))}$

In order for the lifted Reed-Solomon to match the alphabet size of the lifted Hermitian code (by taking locality $N^{\epsilon/3}$), its error correcting rate must become $\alpha^{O((6/\epsilon)^{(3/\epsilon)}\log(1/\epsilon))}$ which is worse than that of the lifted Hermitian code for sufficiently small ϵ.

In comparison with the multiplicity codes of [4], the lifted Hermitian code achieves a much smaller alphabet but also much poorer (though still positive constant) error correction rate. The smaller alphabet is not necessarily an advantage, since one can simply concatenate the multiplicity codes with a suitably good linear code over an alphabet of constant size and still achieve N^ϵ locality, $1-\alpha$ rate, and constant distance. However, the lifted Hermitian code may outperform the multiplicity code in certain concrete settings of parameters.

Organization. In Section 2 we introduce standard notation and terminology used in the paper. In Section 3 we present the key definitions and notions used in the paper, in particular the definition of lifting. In Section 4 we show that if a set of functions is sufficiently "close to doubly transitive", lifting a code with respect to the set yields a code with good distance. In Section 5 we show in addition that the lifted codes are locally correctable. We emphasize that Sections 3, 4, and 5 apply to arbitrary base codes, not necessarily algebraic or even linear codes. In Section 6, we introduce the base codes used in our constructions. We review the Reed-Solomon code as a warmup, and then present the Hermitian code which we lift in Section 7 to obtain explicit high rate locally decodable codes with small alphabet size. We conclude in Section 8.

2 Preliminaries

2.1 Notation

For integers $a < b$, let $[a, b]$ denote the the set $\{a, a+1, a+2, \ldots, b\}$ and let $[a]$ denote $[1, a]$. Throughout the paper, we let Φ denote a set of functions mapping $D \to D$. We assume that Φ contains the identity $\mathrm{id} : D \to D$ which fixes every element of D. We say Φ *acts* on D.

Let $f : D \to R$ and let $\sigma \in \Phi$ where Φ acts on D. The function $f \circ \sigma : D \to R$ is defined by

$$(f \circ \sigma)(x) = f(\sigma(x))$$

for all $x \in D$. Let $m \geq 1$ and let $\sigma = (\sigma_1, \ldots, \sigma_m) \in \Phi^m$. For a function $f : D^m \to R$, define the function $f|_\sigma : D \to R$ by

$$(f|_\sigma)(x) = f(\sigma_1(x), \ldots, \sigma_m(x))$$

for all $x \in D$. For a set Φ acting on D and a point $u \in D^m$, define the *endomorphisms passing through* u to be

$$\Phi_u \triangleq \{\sigma \in \Phi^m \mid \sigma_1 = \mathrm{id}, \sigma_i(u_1) = u_i \, \forall i \in [2, m]\}.$$

For an event A, let $\mathbb{1}_A$ denote the indicator variable for A, i.e.

$$\mathbb{1}_A = \begin{cases} 1 & \text{if A} \\ 0 & \text{otherwise.} \end{cases}$$

Let $f, g : D \to R$. The *(relative) distance between f and g*, is

$$\delta(f, g) \triangleq \mathbb{E}_{x \in D}[\mathbb{1}_{f(x) \neq g(x)}].$$

For a collection $\mathcal{C} \subseteq \{D \to R\}$ of functions, the distance between $f : D \to R$ and $\mathcal{C}$ is

$$\delta(f, \mathcal{C}) \triangleq \min_{g \in \mathcal{C}} \delta(f, g).$$

For a code $\mathcal{C} \subseteq \{D \to R\}$, the distance of $\mathcal{C}$ is

$$\delta(\mathcal{C}) \triangleq \min_{\substack{f,g \in \mathcal{C} \\ f \neq g}} \delta(f, g)$$

If q is a prime power, let $\mathbb{F}_q$ denote the finite field of order q, which is unique up to isomorphism.

2.2 Terminology

For an algorithm $\mathcal{A}$ and function f, let $\mathcal{A}^f$ denote the algorithm $\mathcal{A}$ given oracle access to f.

Definition 2.1 (Locally correctable code)**.** A code $\mathcal{C} \subseteq \{D \to R\}$ is (q, τ)-*locally correctable* if there exists a randomized algorithm $\mathcal{A}$ satisfying the following properties:

1. $\mathcal{A}^f$ makes at most q queries to f;
2. If there exists $g \in \mathcal{C}$ such that $\delta(f, g) \leq \tau$, then for every $x \in D$ we have $\mathcal{A}^f(x) = g(x)$ with probability at least $2/3$ over the randomness of $\mathcal{A}$.

Definition 2.2 (Locally decodable code)**.** A code $\mathcal{C} \subseteq \{D \to R\}$ is (q, τ)-*locally decodable* if $\mathcal{C}$ is the image of an encoding function $\mathsf{Enc} : R^k \to R^D$ and there exists a randomized algorithm $\mathcal{A}$ satisfying the following properties:

1. $\mathcal{A}^f$ makes at most q queries to f;
2. If there exists $m \in R^k$ such that $\delta(f, \mathsf{Enc}(m)) \leq \tau$, then for every $i \in [k]$ we have $\mathcal{A}^f(i) = m_i$ with probability at least $2/3$ over the randomness of $\mathcal{A}$.

For linear codes, local correctability is stronger than local decodability, since one can arrange the generator matrix of the code such that the message is part of the codeword.

3 Definitions

In this section we give the key definitions in the paper, namely Φ-lifting and the notion of a set Φ being "close" to doubly transitive, which is borrowed from [1].

3.1 Lifting

Definition 3.1. Let Φ act on D and let $\mathcal{C} \subseteq \{D \to R\}$. The *$m$-dimensional Φ-lift* of $\mathcal{C}$, denoted $\mathrm{Lift}^m_\Phi(\mathcal{C})$, is the set

$$\mathrm{Lift}^m_\Phi(\mathcal{C}) \triangleq \{f : D^m \to R \mid f|_\sigma \in \mathcal{C} \text{ for all } \sigma \in \Phi^m\}.$$

We say $\mathcal{C} \subseteq \{D \to R\}$ is Φ-*invariant* if whenever $f \in \mathcal{C}$ and $\sigma \in \Phi$ we also have $f \circ \sigma \in \mathcal{C}$. Notice that Definition 3.1 does not require that $\mathcal{C}$ be Φ-invariant, or even that Φ be a group! Indeed, Φ-invariance only ensures us that

$$\mathrm{Lift}^1_\Phi(\mathcal{C}) = \mathcal{C}$$

and if in addition is a group, then the lift operation composes:

$$\mathrm{Lift}^m_{\Phi^n}(\mathrm{Lift}^n_\Phi(\mathcal{C})) = \mathrm{Lift}^{mn}_\Phi(\mathcal{C})$$

where Φ^n acts on D^m componentwise, i.e. if $\varphi = (\varphi_1, \dots \varphi_m) \in \Phi^m$ and $x = (x_1, \dots, x_m) \in D^m$, then $\varphi(x) = (\varphi_1(x_1), \dots, \varphi_n(x_n))$.

The affine lifting found in [3] is (almost) an example of our notion of lifting. Take $D = R = \mathbb{F}_q$ and Φ to be the group of affine permutations on D, i.e. maps of the form $x \mapsto ax + b$ for $a \in \mathbb{F}_q^*$, $b \in \mathbb{F}_q$. Then $\mathrm{Lift}^m_\Phi(\mathcal{C})$ consists of all $f : \mathbb{F}_q^m \to \mathbb{F}_q$ such that $f|_L \in \mathcal{C}$ for all lines L that are not axis-parallel. The affine-lifted codes in [3] consider every line, including the axis-parallel ones. Though we could have defined Φ-lifting to properly generalize affine-lifting, we chose our definition because it is cleaner to state, makes proofs cleaner, and makes negligible difference in the parameters we care about. We point out that one limitation of our definition is that we can only lift a domain D to a direct product D^m, whereas the affine lifting of [3] allows lifting from $\mathbb{F}_q^m$ to $\mathbb{F}_q^n$ for any $m \le n$.

Though any code can be lifted, our constructions in the paper use linear codes as the base code. A code $\mathcal{C} \subseteq \{D \to R\}$ is *linear* if $R = \mathbb{F}$ is a field and $\mathcal{C}$ is a $\mathbb{F}$-vector space. To argue that the lifted code is large, we argue that it has large dimension by showing it contains many linearly independent codewords. To do so, we need the following fact, which is straightforward to verify.

Proposition 3.2. *If $\mathcal{C}$ is linear over $\mathbb{F}$, then so is $\mathrm{Lift}^m_\Phi(\mathcal{C})$.*

3.2 Double transitivity

Now we define the notions of "closeness" to double transitivity that we will work with. There are two such notions, taken from [1].

Definition 3.3. A set Φ acting on a set D is *doubly transitive* if it is transitive on pairs in Φ, i.e. for every $x_1 \neq x_2 \in D$ and $y_1 \neq y_2 \in D$, there exists $\sigma \in \Phi$ such that $\sigma(x_1) = y_1$ and $\sigma(x_2) = y_2$.

Definition 3.4 ([1])**.** A set Φ acting on a set D is (ϵ, α)-*doubly transitive* if, for every $x_1, x_2 \in D$, for at least $1 - \epsilon$ fraction of points $x \in D$, the random variable $\sigma(x)$ is uniformly distributed on $1 - \alpha$ fraction of D, where σ is chosen uniformly from the set $\{\sigma \in \Phi \mid \sigma(x_1) = x_2\} = \Phi_{(x_1,x_2)}$.

When Φ is a group acting transitively on D, double transitivity is equivalent to $\left(\frac{1}{|D|}, 0\right)$-double transitivity (see [1, Lemmas 6.8, 6.9]). Indeed, given $x_1, x_2 \in D$, for every point $x \neq x_1$, the random variable $\sigma(x)$ is uniformly distributed on D, when σ is drawn from those mapping $\sigma(x_1) = x_2$. However, $\sigma(x_1)$ itself will always equal x_2.

Example 3.5. Let $D = \mathbb{F}_q$ and $\Phi = \{x \mapsto ax + b \mid a \in \mathbb{F}_q^*, b \in \mathbb{F}_q\}$. Then Φ is $(\frac{1}{q}, 0)$-double transitive. This follows from the fact that Φ is doubly transitive on D. Another way to see this is to note that, given $x_1, x_2 \in D$, $\sigma(x_1) = x_2$ implies $\sigma(x) = a(x - x_1) + x_2$ for some $a \in \mathbb{F}_q^*$. Therefore, for every $x \neq x_1$ and every $y \in \mathbb{F}_q$, there exists a unique σ such that $\sigma(x) = y$, namely the one with $a = (y - x_2)(x - x_1)^{-1}$.

The second notion of "closeness" to double transitive involves distributions that are statistically close to uniform. The precise definition is as follows.

Definition 3.6. Let p_1, p_2 be two distributions on D, i.e. $\sum_{x\in D} p_1(x) = \sum_{x\in D} p_2(x) = 1$ and $p_1(x), p_2(x) \geq 0$ for all $x \in D$. The *distance between* p_1 *and* p_2 is

$$\|p_1 - p_2\| \triangleq \max_{A\subseteq D} \left| \sum_{x\in A} p_1(x) - \sum_{x\in A} p_2(x) \right| .$$

Definition 3.7 ([1])**.** A set Φ acting on a set D is (α, ϵ)-*close to c-steps uniform* if, for every $x_1, x_2 \in D$, for at least $1-\epsilon$ fraction of points $x \in D$, if one uniformly randomly chooses $w_1, \dots, w_{c-1} \in D$ and $\sigma_1, \dots, \sigma_c \in \Phi$ such that $\sigma_1(x_1) = x_2$ and $\sigma_i(w_{i-1}) = \sigma_{i-1}(w_{i-1})$ for $2 \leq i \leq c$, then the random variable $\sigma_c(x)$ is α-close to uniformly distributed on D.

The motivation behind Definition 3.7 is the use of *fractal correcting* in [1]. Intuitively, one may think of $f|_\sigma$ as f restricted to some curve in D^m. For simplicity assume $m = 2$. To correct the received word f at a particular point x, the usual approach is to pick a random curve passing through x and correct the shorter word f restricted to the curve. Parametrize the curve by $(x, \sigma(x))$. Then the condition that the curve passes through $x = (x_1, x_2)$ is equivalent to $\sigma(x_1) = x_2$. If the curve samples D uniformly, then with high probability the curve does not contain too many corrupted points. If Φ is not doubly transitive, however, then random curves may not sample D^2 uniformly. The intuition behind fractal correcting is to first pick a random curve σ_1 passing through x (i.e. $\sigma_1(x_1) = x_2$), then pick a random point $(w_1, \sigma_1(w_1))$ sitting on the point, then pick another random curve σ_2 passing through $(w_1, \sigma_1(w_1))$ (i.e. $\sigma_2(w_1) = \sigma_1(w_1)$) and so on. After c steps, the cth curve σ_c will sample the space nearly uniformly. We elaborate on this in Sections 4 and 5.

4 Distance of lifted codes

In this section we show that if $\mathcal{C}$ is a code with constant positive distance, and the set Φ acting on the domain D is nearly doubly transitive, then $\mathrm{Lift}_\Phi^m(\mathcal{C})$ has constant positive distance. Our lower bound on the distance of $\mathrm{Lift}_\Phi^m(\mathcal{C})$ degrades as m grows, but for our purposes m is constant, so the distance of the lift is constant as well. We emphasize that the results in this section apply to *any* code $\mathcal{C}$, even non-linear codes.

We begin by lower bounding the distance of the lift when the set is close to doubly transitive, in the sense of Definition 3.4, i.e. when Φ is (ϵ, α)-double transitive.

The following lemma will be used in proving both Theorems 4.2 and 5.1.

Lemma 4.1. *Let Φ acting on D be (ϵ, α)-double transitive. Let $m \geq 1$ and let $f, g \in \{D^m \to R\}$. Fix $x \in D^m$. Then*

$$\mathbb{E}_{\sigma\in\Phi_x}[\delta(f|_\sigma, g|_\sigma)] \leq \epsilon + \frac{\delta(f,g)}{(1-\alpha)^{m-1}}.$$

Proof. Let $D' \subseteq D$ be the set of $z \in D$ such that $\sigma(z)$ is uniform over $1-\alpha$ fraction of D, when σ

is chosen uniformly from Φ_x, as in Definition 3.4. Note that $|D'| \geq (1-\epsilon)|D|$. We have

$$\begin{aligned}
\mathbb{E}_{\sigma\in\Phi_x}\mathbb{E}_{z\in D}[\mathbb{1}_{f|_\sigma(z)\neq g|_\sigma(z)}] &= \mathbb{E}_{z\in D}\mathbb{E}_{\sigma\in\Phi_x}[\mathbb{1}_{f|_\sigma(z)\neq g|_\sigma(z)}] \\
&= \mathbb{E}_{z\notin D'}\mathbb{E}_{\sigma\in\Phi_x}[\mathbb{1}_{f|_\sigma(z)\neq g|_\sigma(z)}] + \mathbb{E}_{z\in D'}\mathbb{E}_{\sigma\in\Phi_x}\mathbb{1}_{f|_\sigma(z)\neq g|_\sigma(z)}] \\
&\leq \epsilon + \mathbb{E}_{z\in D'}\mathbb{E}_{\sigma\in\Phi_x}[\mathbb{1}_{f|_\sigma(z)\neq g|_\sigma(z)}] \\
&\leq \epsilon + \frac{\delta(f,g)}{(1-\alpha)^{m-1}}
\end{aligned}$$

where the final inequality follows from the fact that the last $m-1$ coordinates of $\sigma(z)$ are uniform over $(1-\alpha)^{m-1}$ fraction of D^{m-1} and in the worst case all the disparate points of f and g all lie in this subset. □

Theorem 4.2. *Let $\mathcal{C} \subseteq \{D \to R\}$ be a code with distance δ, and Φ acting on D is (ϵ,α)-doubly transitive. Then $\delta(\mathrm{Lift}^m_\Phi(\mathcal{C})) \geq (1-\alpha)^{m-1}(\delta-\epsilon)$.*

Proof. Let $f, g \in \mathrm{Lift}^m_\Phi(\mathcal{C})$ be distinct and fix $x \in D^m$ such that $f(x) \neq g(x)$. By Lemma 4.1,

$$\mathbb{E}_{\sigma\in\Phi_x}[\delta(f|_\sigma, g|_\sigma)] \leq \epsilon + \frac{\delta(f,g)}{(1-\alpha)^{m-1}}.$$

Therefore, there exists $\sigma \in \Phi_x$ such that $\delta(f|_\sigma, g|_\sigma) \leq \epsilon + \frac{\delta(f,g)}{(1-\alpha)^{m-1}}$. But $f|_\sigma(x_1) = f(x) \neq g(x) = g|_\sigma(x_1)$, so $f|_\sigma$ and $g|_\sigma$ are distinct codewords of $\mathcal{C}$ and hence $\delta \leq \epsilon + \frac{\delta(f,g)}{(1-\alpha)^{m-1}}$, i.e. $\delta(f,g) \geq (1-\alpha)^{m-1}(\delta-\epsilon)$. □

Next we prove a similar result when Φ is close to doubly transitive in the sense of Definition 3.7, i.e. is to (α,ϵ)-close to c-steps uniform. First, some straightforward but useful facts.

Lemma 4.3. *If X and Y are independent and X is α-close to uniform over S and Y is β-close to uniform over T, then (X,Y) is $\alpha+\beta$-uniform over $S \times T$.*

Corollary 4.4. *If $X_i \in D$ is α-close to uniform for each $i \in [m]$ and are independent, then $(X_1,\ldots,X_m) \in D^m$ is $m\alpha$-close to uniform.*

The following lemma will be used in proving both Theorems 4.6 and 5.3.

Lemma 4.5. *Let Φ acting on D be (α,ϵ)-close to c-steps uniform. Let $m \geq 1$ and let $f, g \in \{D^m \to R\}$. Fix $x \in D^m$. Then*

$$\mathbb{E}_{\sigma_1\in\Phi_x}\mathbb{E}_{w_1\in D}\mathbb{E}_{\sigma_2\in\Phi_{\sigma_1(w_1)}}\cdots\mathbb{E}_{\sigma_c\in\Phi_{\sigma_{c-1}(w_{c-1})}}[\delta(f|_{\sigma_c}, g|_{\sigma_c})] \leq \delta(f,g) + \epsilon + m\alpha.$$

Proof. Let $D' \subseteq D$ be the set of $z \in D$ such that $\sigma_c(z)$ is α-close to uniform, as in Definition 3.7. Note that $|D'| \geq (1-\epsilon)|D|$. Then

$$\begin{aligned}
&\mathbb{E}_{\sigma_1\in\Phi_x}\mathbb{E}_{w_1\in D}\mathbb{E}_{\sigma_2\in\Phi_{\sigma_1(w_1)}}\cdots\mathbb{E}_{\sigma_c\in\Phi_{\sigma_{c-1}(w_{c-1})}}\mathbb{E}_{z\in D}\left[\mathbb{1}_{f|_{\sigma_c}(z)\neq g|_{\sigma_c}(z)}\right] \\
=\ &\mathbb{E}_{z\in D}\mathbb{E}_{\sigma_1\in\Phi_x}\mathbb{E}_{w_1\in D}\mathbb{E}_{\sigma_2\in\Phi_{\sigma_1(w_1)}}\cdots\mathbb{E}_{\sigma_c\in\Phi_{\sigma_{c-1}(w_{c-1})}}\left[\mathbb{1}_{f|_{\sigma_c}(z)\neq g|_{\sigma_c}(z)}\right] \\
\leq\ &\epsilon + \mathbb{E}_{z\in D'}\mathbb{E}_{\sigma_1\in\Phi_x}\mathbb{E}_{w_1\in D}\mathbb{E}_{\sigma_2\in\Phi_{\sigma_1(w_1)}}\cdots\mathbb{E}_{\sigma_c\in\Phi_{\sigma_{c-1}(w_{c-1})}}\left[\mathbb{1}_{f|_{\sigma_c}(z)\neq g|_{\sigma_c}(z)}\right] \\
\leq\ &\epsilon + \delta(f,g) + m\alpha.
\end{aligned}$$

□

Theorem 4.6. *Let $\mathcal{C}$ be a code with distance δ, and Φ is (α, ϵ)-close to c-steps uniform. Then $\delta(\mathrm{Lift}_{\Phi}^{m}(\mathcal{C})) \geq \delta^c - m\alpha - \epsilon$.*

Proof. Let $f, g \in \mathrm{Lift}_{\Phi}^{m}(\mathcal{C})$ be distinct and let $\tau = \delta(f, g)$. Fix $x \in D$ such that $f(x) \neq g(x)$. We claim that, for each $i \in [c]$, there exists $w_{i-1} \in D$ and $\sigma_i \in \Phi_{\sigma_{i-1}(w_{i-1})}$ such that

$$0 < \mathbb{E}_{w_i \in D} \mathbb{E}_{\sigma_{i+1} \in \Phi_{\sigma_i(w_i)}} \cdots \mathbb{E}_{\sigma_c \in \Phi_{\sigma_{c-1}(w_{c-1})}} \mathbb{E}_{z \in D} \left[\mathbb{1}_{f|_{\sigma_c}(z) \neq g|_{\sigma_c}(z)} \right] \leq \frac{\tau + m\alpha + \epsilon}{\delta^{i-1}}.$$

We prove the claim by induction. The base case $i = 1$ follows by taking $\sigma_0 \in \Phi_x$, $w_0 = x_1$, and noting that, by Lemma 4.5, since

$$\mathbb{E}_{\sigma_1 \in \Phi_x} \mathbb{E}_{w_1 \in D} \mathbb{E}_{\sigma_2 \in \Phi_{\sigma_1(w_1)}} \cdots \mathbb{E}_{\sigma_c \in \Phi_{\sigma_{c-1}(w_{c-1})}} \mathbb{E}_{z \in D} \left[\mathbb{1}_{f|_{\sigma_c}(z) \neq g|_{\sigma_c}(z)} \right] \leq \tau + m\alpha + \epsilon,$$

there exists $\sigma_1 \in \Phi_x$ such that

$$\mathbb{E}_{w_1 \in D} \mathbb{E}_{\sigma_2 \in \Phi_{\sigma_1(w_1)}} \cdots \mathbb{E}_{\sigma_c \in \Phi_{\sigma_{c-1}(w_{c-1})}} \mathbb{E}_{z \in D} \left[\mathbb{1}_{f|_{\sigma_c}(z) \neq g|_{\sigma_c}(z)} \right] \leq \tau + m\alpha + \epsilon.$$

Moreover, this expectation is positive because $f(x) \neq g(x)$. Now suppose we have proved the $i - 1$ case. The restrictions $f|_{\sigma_{i-1}}$ and $g|_{\sigma_{i-1}}$ are distinct codewords of $\mathcal{C}$ (since they disagree at w_{i-2}) and hence for at least δ-fraction of $w_{i-1} \in D$ we have $f(\sigma_{i-1}(w_{i-1})) \neq g(\sigma_{i-1}(w_{i-1}))$. Restricting to these w_{i-1}, we get

$$0 < \delta \cdot \mathbb{E}_{\sigma_i \in \Phi_{\sigma_{i-1}(w_{i-1})}} \mathbb{E}_{w_i \in D} \mathbb{E}_{\sigma_{i+1} \in \Phi_{\sigma_i(w_i)}} \cdots \mathbb{E}_{\sigma_c \in \Phi_{\sigma_{c-1}(w_{c-1})}} \mathbb{E}_{z \in D} \left[\mathbb{1}_{f|_{\sigma_c}(z) \neq g|_{\sigma_c}(z)} \right] \leq \frac{\tau + m\alpha + \epsilon}{\delta^{i-2}}$$

and the claim thus follows.

From the $i = c$ case of the claim, it follows that there exists $\sigma_c \in \Phi$ such that

$$0 < \mathbb{E}_{z \in D} \left[\mathbb{1}_{f|_{\sigma_c}(z) \neq g|_{\sigma_c}(z)} \right] \leq \frac{\tau + m\alpha + \epsilon}{\delta^{c-1}}.$$

Thus $f|_{\sigma}$ and $g|_{\sigma}$ are distinct codewords of $\mathcal{C}$, so we have $\delta \leq \frac{\tau + m\alpha + \epsilon}{\delta^{c-1}}$. □

5 Correction algorithms

In this section we describe how to locally correct a lifted code, given a decoding algorithm for the base code. We present two correcting methods. The first is one-shot correcting, which abstracts the local correcting algorithms for Reed-Muller codes and the affine-lifted Reed-Solomon codes of [3], and is also used for correcting degree-lifted AG codes in [1]. The idea is to pick a random curve passing through the point which we would like to correct, view the restriction of the received word to the curve as a received word that should be close to a codeword of the base code, and then use the base code decoder to correct the point. The second method is fractal correcting, which was introduced by Ben-Sasson et al [1]. The idea is to recursively perform one-shot correcting. To correct a point, pick a random curve passing through it. However, now recursively correct each point on the curve. If Φ is close to c-steps uniform, then fractal correcting with recursion depth c should succeed with high probability. The analysis of the fractal correction algorithm is found in [1], but we include a proof here for completeness. We emphasize that, as in Section 4, the results of this section apply to arbitrary codes $\mathcal{C}$.

5.1 One-shot correcting

The one-shot correcting algorithm $\mathcal{A}$ works as follows. To compute $\mathcal{A}^f(x)$:

1. Pick $\sigma \in \Phi_x$ uniformly at random.
2. Use the decoding algorithm for $\mathcal{C}$ to correct $f|_\sigma$ to some function $g \in \mathcal{C}$.
3. Output $g(x_1)$.

Theorem 5.1. *Let $\mathcal{C} \subseteq \{D \to R\}$ be a code with distance δ and suppose Φ is (ϵ, α)-doubly transitive. Let $\mathcal{L} = \mathrm{Lift}_\Phi^m(\mathcal{C})$. Suppose*

$$\delta(f, \mathcal{L}) < (1-\alpha)^{m-1} \cdot \min\{\delta/6 - \epsilon, (\delta - \epsilon)/2\}.$$

Then there exists a unique $\widehat{f} \in \mathcal{L}$ such that $\delta(f, \widehat{f}) \le \delta(f, \mathcal{L})$ and for any $x \in D^m$ we have $\mathcal{A}^f(x) = \widehat{f}(x)$ with probability at least $2/3$ over the randomness of $\mathcal{A}$.

Proof. By Theorem 4.2, $\delta(\mathcal{L}) \ge (1-\alpha)^{m-1}(\delta - \epsilon)$. Since $\delta(f, \widehat{f}) < \delta(\mathcal{L})/2$, $\widehat{f}$ is unique. Fix $x \in D^m$. By Lemma 4.1,

$$\mathbb{E}_{\sigma \in \Phi_x}[\delta(f|_\sigma, \widehat{f}|_\sigma)] \le \epsilon + \frac{\delta(f, \widehat{f})}{(1-\alpha)^{m-1}} \le \epsilon + \frac{\delta(f, \mathcal{L})}{(1-\alpha)^{m-1}}.$$

By Markov's inequality, with probability at least $2/3$, $\delta(f|_\sigma, \widehat{f}|_\sigma) \le 3\left(\epsilon + \frac{\delta(f,\mathcal{L})}{(1-\alpha)^{m-1}}\right) < \delta/2$. Step 2 of the algorithm finds some $g \in \mathcal{C}$ such that $\delta(f|_\sigma, g) < \delta/2$. But both $g, \widehat{f}|_\sigma \in \mathcal{C}$ and $\delta(g, \widehat{f}|_\sigma) < \delta$, so in fact $g = \widehat{f}|_\sigma$. Therefore, $\mathcal{A}^f(x) = g(x_1) = \widehat{f}|_\sigma(x_1) = \widehat{f}(x)$. □

Corollary 5.2. *If $\mathcal{C} \subseteq \{D \to R\}$ has distance δ and Φ acting on D is (ϵ, α)-doubly transitive, then $\mathrm{Lift}_\Phi^m(\mathcal{C})$ is (q, τ)-locally correctable for $q = |D|$ and $\tau = O((1-\alpha)^{m-1}(\delta - \epsilon))$.*

5.2 Fractal correcting

The c-step fractal correction algorithm $\mathcal{A}_c$ works as follows. To compute $\mathcal{A}_c^f(x)$:

1. If $c = 1$, output $\mathcal{A}^f(x)$.
2. Otherwise, $c > 1$. Pick $\sigma \in \Phi_x$ uniformly at random.
3. Compute $f' \triangleq \mathcal{A}_{c-1}^f|_\sigma$. That is, for each $z \in D$ let $f'(z) = \mathcal{A}_{c-1}^f(\sigma(z))$.
4. Use the decoding algorithm for $\mathcal{C}$ to correct f' to some function $g \in \mathcal{C}$.
5. Output $g(x_1)$.

Theorem 5.3. *Let $\mathcal{C} \subseteq \{D \to R\}$ be a code with distance δ and suppose Φ acting on D is (α, ϵ)-close to c-steps uniform. Let $\mathcal{L} = \mathrm{Lift}_\Phi^m(\mathcal{C})$. Suppose*

$$\delta(f, \mathcal{L}) < \min\left\{\frac{1}{3}(\delta/2)^c - \epsilon - m\alpha, (\delta^c - \epsilon - m\alpha)/2\right\}.$$

Then there exists a unique $\widehat{f} \in \mathcal{L}$ such that $\delta(f, \widehat{f}) \le \delta(f, \mathcal{L})$ and for any $x \in D^m$ we have $\mathcal{A}_c^f(x) = \widehat{f}(x)$ with probability at least $2/3$ over the randomness of $\mathcal{A}$.

Proof. By Theorem 4.6, $\delta(\mathcal{L}) \geq \delta^c - \epsilon - m\alpha$. Since $\delta(f, \widehat{f}) < \delta(\mathcal{L})/2$, $\widehat{f}$ is unique. Fix $x \in D^m$. For $i \in [c]$, let p_i denote the average probability that the ith bottom-most level of the recursion fails. Our goal is to show that $p_c \leq 1/3$. We will show in fact that $p_i \leq \frac{1}{3}(\delta/2)^{c-i}$ for all $i \in [c]$. By Lemma 4.5, the average of $\delta(f|_{\sigma_c}, \widehat{f}|_{\sigma_c})$ over all σ_c chosen in the bottom-most level is at most $\delta(f, \mathcal{L}) + \epsilon + m\alpha$, so by Markov's inequality with probability at most $\frac{2}{\delta}(\delta(f, \mathcal{L}) + \epsilon + m\alpha)$ we have $\delta(f|_{\sigma_c}, \widehat{f}|_{\sigma_c}) > \delta/2$, i.e. $p_1 \leq \frac{2}{\delta}(\delta(f, \mathcal{L}) + \epsilon + m\alpha) \leq \frac{1}{3}(\delta/2)^{c-1}$.

Now inductively assume $p_i \leq \frac{1}{3}(\delta/2)^{c-i}$. The average value of $\delta(f|_{\sigma_{c-i+1}}, \widehat{f}|_{\sigma_{c-i+1}})$ is at most p_i. By Markov's inequality, with probability at most $\frac{2}{\delta}p_i$ we have $\delta(f|_{\sigma_{c-i}}, \widehat{f}|_{\sigma_{c-i}}) > \delta/2$, so $p_{i+1} \leq \frac{2}{\delta}p_i \leq \frac{1}{3}(\delta/2)^{c-(i+1)}$. ∎

Corollary 5.4. *If $\mathcal{C} \subseteq \{D \to R\}$ has distance δ for some Φ that is (α, ϵ)-close to c-steps uniform, where $c = O(1)$, then $\mathrm{Lift}_\Phi^m(\mathcal{C})$ is (q, τ)-locally correctable for $q = |D|^c$ and $\tau = O(\delta^c - \epsilon - m\alpha)$.*

6 Base codes

In this section we review existing codes, in particular the Reed-Solomon code and the Hermitian code, the latter which we use in Section 7 to construct new high rate locally correctable codes over small alphabets.

Algebraic geometry codes. The Reed-Solomon and Hermitian codes are instances of *algebraic geometry codes*. Since we can describe our base codes, our lifted codes, and their properties without using any terminology typically used in the context of AG codes (e.g. the language of algebraic function fields), we avoid using such terminology and stick to an elementary exposition. In fact, the only deep results from the theory of algebraic function fields that we use can be stated in elementary terms. The interested reader is referred to [6] for details on the theory of algebraic function fields and codes.

6.1 Reed-Solomon code

Let q be a prime power. The *Reed-Solomon code* $\mathrm{RS}_q[r] \subseteq \mathbb{F}_q[x]/(x^q - x)$ can be defined as

$$\mathrm{RS}[r] \triangleq \mathrm{span}_{\mathbb{F}_q}\{x^i \mid i < r\}.$$

It is a $[q, r, q-r+1]_q$-code. Note that its alphabet size $q = N$ where N is its block size. One can identify $\mathbb{F}_q[x]/(x^q - x)$ with $\{\mathbb{F}_q \to \mathbb{F}_q\}$. Consider the group Φ consisting of all affine permutations on $\mathbb{F}_q$, i.e. $\Phi = \{x \mapsto ax + b \mid a \in \mathbb{F}_q^*, b \in \mathbb{F}_q\}$, which acts on $\mathbb{F}_q$. Clearly $\mathrm{RS}_q[r]$ is Φ-invariant. Moreover, Φ is doubly transitive (Example 3.5) and $|\Phi| = q(q-1)$, so it is just large enough to be doubly transitive. In [3], it was shown that $\mathrm{Lift}_\Phi^m(\mathrm{RS}_q[(1-\delta)q])$ has block length q^m, distance at least $\delta - \frac{1}{q}$ (which also follows from Theorem 4.2), and rate at least $1 - \delta^{\Omega\left(\frac{1}{m^m \log m}\right)}$ when q is a power of 2.

6.2 Hermitian code

Let q be a prime power. The *Hermitian curve* $H \subseteq \mathbb{F}_{q^2}^2$ is the set

$$H \triangleq \{(x, y) \mid N(x) = Tr(y)\}$$

where $N : \mathbb{F}_{q^2} \to \mathbb{F}_q$ is the *norm* $N(x) = x^{1+q}$ and $Tr : \mathbb{F}_{q^2} \to \mathbb{F}_q$ is the *trace* $Tr(x) = x + x^q$. It can be shown that N is multiplicative and is a surjective group homomorphism from $\mathbb{F}_{q^2}^* \to \mathbb{F}_q^*$ (and hence a $(q+1)$-to-1 map on $\mathbb{F}_{q^2}^*$) and that Tr is additive and is a surjective $\mathbb{F}_q$-linear map from $\mathbb{F}_{q^2} \to \mathbb{F}_q$ (and hence a q-to-1 map on $\mathbb{F}_{q^2}$). It follows that $|H| = q^3$, since for every $x \in \mathbb{F}_{q^2}$ there are exactly q values of $y \in \mathbb{F}_{q^2}$ such that $Tr(y) = N(x)$.

The *Hermitian code* $\mathrm{Herm}_q[r] \subseteq \mathbb{F}_{q^2}[x]/(x^{q^2} - x, y^{q^2} - y, N(x) - Tr(y))$ is defined as

$$\mathrm{Herm}_q[r] \triangleq \mathrm{span}_{\mathbb{F}_{q^2}}\{x^i y^j \mid qi + (q+1)j < r, j < q\}.$$

It follows from the Riemann-Roch theorem that $\mathrm{Herm}_q[r]$ is a $[q^3, r-g, q^3-r+1]_{q^2}$-code, where $g = \frac{q(q-1)}{2}$ is the genus of the curve H (one can also deduce this by counting the number of "degrees" d which cannot be obtained by a sum $qi + (q+1)j$). Though the Hermitian code has a worse rate-distance trade-off than the Reed-Solomon code, its alphabet size is significantly smaller (q^2 compared to a block length of q^3).

Consider the group Φ of maps $(x, y) \mapsto (ax + b, a^{q+1}y + ab^q x + c)$ for $a \in \mathbb{F}_{q^2}^*$, $(b, c) \in H$. One can verify that this a group of order $q^3(q^2-1)$ acting on H and moreover $\mathrm{Herm}_q[r]$ is Φ-invariant. For interesting values of r, the group Φ is the largest group under which the Hermitian code is invariant [7]. The group Φ is not doubly transitive, but it is shown in [1] that it is almost doubly transitive, in both the senses of Definitions 3.4 and 3.7. We recall the precise statements.

Proposition 6.1 ([1, Theorem 6.3])**.** *Let Φ be as above. Then Φ is (ϵ, α)-doubly transitive for $\epsilon = \frac{1}{q^2}$ and $\alpha = 1 - \frac{1}{q}$.*

Proposition 6.2 ([1, Theorem 7.3])**.** *Let Φ be as above. Then Φ is (α, ϵ)-close to 2-steps uniform for $\alpha = \epsilon = \frac{1}{q}$.*

Applying Theorem 4.6 and Corollary 5.4 to the above facts, we immediately get the following.

Theorem 6.3. *Let Φ be the group of automorphisms on H of the form $(x, y) \mapsto (ax + b, a^{q+1}y + ab^q x + c)$. Let $r = (1-\delta)q^3$, so that $\mathrm{Herm}_q[r]$ has distance δ. Then $\mathrm{Lift}_\Phi^m(\mathrm{Herm}_q[r])$ has distance at least $\delta^2 - \frac{m}{q}$ and is $(q^6, O(\delta^2 - \frac{m}{q}))$-locally correctable.*

Note that, though the Φ-lift of $\mathrm{Herm}_q[(1-\delta)q^3]$ has distance roughly δ^2 which is less than that of the degree-lift, whose distance is δ (see [1]), its error correcting capability is the same.

7 Explicit Constructions

In this section we prove the following.

Theorem 7.1. *Given $\epsilon, \alpha, N_0 > 0$, for infinitely many $N \geq N_0$ there exists a code of length N, rate $1 - \alpha$, alphabet size $N^{\epsilon/3}$ and is $(N^\epsilon, \alpha^{O((8/\epsilon)^{(2/\epsilon)} \log(1/\epsilon))})$-locally correctable.*

We prove this using lifted Hermitian codes. We defer the proof to the end of the section.

Let $m \geq 1$, let $q = 2^\ell > m$, let $c > 0$ such that $\ell - c > \lceil \log_2 m \rceil$, and let $r = (1 - 2^{-c})q^3$. Let Φ be the group of automorphisms on the Hermitian curve $H \subseteq \mathbb{F}_{q^2}^2$ of the form $(x, y) \mapsto (ax + b, a^{q+1}y + ab^q x + c)$, and let $\mathcal{L} = \mathrm{Lift}_\Phi^m(\mathrm{Herm}_q[r])$. By Theorem 6.3, $\mathcal{L}$ has distance $2^{-2c} - \frac{m}{q}$

and is $O(q^6, O(2^{-2c} - \frac{m}{q}))$-locally correctable. Its length is q^{3m} and alphabet size is q^2. The only missing parameter is the rate, to which we devote the rest of this section.

After lifting, the domain of our code is

$$H^m = \{(x_1, y_1, \ldots, x_m, y_m) \in \mathbb{F}_{q^2}^m \mid N(x_k) = Tr(y_k) \;\; \forall k \in [m]\}.$$

A *monomial on* H^m is a monomial of the form $\prod_{k=1}^m x_k^{i_k} y_k^{j_k}$ with $i_k < q^2$ and $j_k < q$ for all $k \in [m]$. The reason for these conditions is to ensure the monomials define distinct functions on H^m. In fact, one can show the monomials on H^m form a basis of $\{H^m \to \mathbb{F}_{q^2}\}$ as a $\mathbb{F}_{q^2}$-vector space.

Definition 7.2. Let p be a prime. Let $a, b \in \mathbb{N}$ and consider their base p representations $a = \sum_{i\geq 0} a_i p^i$ and $b = \sum_{i\geq 0} b_i p^i$ where each $a_i, b_i \in [0, p-1]$. Then a *is the in the* p*-shadow of* b, denoted $a \leq_p b$, if $a_i \leq b_i$ for all i. Moreover, for $a, b, c \in \mathbb{N}$, we say $(a, b) \leq_p c$ if $a_i + b_i \leq c_i$ for all i.

The following generalized theorem of Lucas will be crucial for our analysis later. For $a + b \leq c$, we let $\binom{a}{b,c}$ denote the standard trinomial coefficient $\frac{a!}{b!c!(a-b)!}$ which is the coefficient of $x^b y^c$ in the expansion of $(x + y + 1)^a$. Note that the standard binomial coefficient is $\binom{a}{b} = \binom{a}{b,0}$

Theorem 7.3 ((Generalized) Lucas' theorem)**.** *Let* $a, b, c \in \mathbb{N}$ *with* p*-ary representations given by* a_i, b_i, c_i*. Then*

$$\binom{a}{b, c} \equiv \prod_{i\geq 0} \binom{a_i}{b_i, c_i} \bmod p.$$

In particular, $\binom{a}{b,c} \bmod p$ *is nonzero only if* $(b, c) \leq_p a$*.*

Our strategy for lower bounding $\dim_{\mathbb{F}_{q^2}} \mathcal{L}$ is to lower bound the number of monomials on H^m in $\mathcal{L}$. For a monomial $f(x_1, y_1, \ldots, x_m, y_m) = \prod_{k=1}^m x_k^{i_k} y_k^{j_k}$ and a map $\sigma \in \Phi^m$ where $\sigma_k(x, y) = (a_k x + b_k, a_k^{q+1} y + a_k b_k^q x + c_k)$, we have

$$\begin{aligned} f(\sigma(x,y)) &= \prod_{k=1}^m (a_k x + b_k)^{i_k} \left(a_k^{q+1} y + b_k^q x + c_k\right)^{j_k} \\ &= \prod_{k=1}^m \left(\sum_{d_k \leq_p i_k} (\cdots) x^{d_k}\right) \left(\sum_{(d'_k, e_k) \leq_p j_k} (\cdots) x^{d'_k} y^{e_k}\right) \\ &= \sum_{\forall k \;\; d_k \leq_p i_k, (d'_k, e_k) \leq_p j_k} (\cdots) x^{\sum_{k=1}^m d_k + d'_k} y^{\sum_{k=1}^m e_k} \end{aligned}$$

where the $(\cdots)$ indicate constants which do not matter. Thus, the monomial f is in $\mathcal{L}$ if the following holds: for all $k \in [m]$, for all $d_k \leq_p i_k$ and all $(d'_k, e_k) \leq_p j_k$, after reducing the monomial $x^{\sum_{k=1}^m d_k + d'_k} y^{\sum_{k=1}^m e_k}$ modulo the ideal $I \triangleq (x^{q^2} - x, y^{q^2} - y, x^{q+1} - y^q - y)$, the resulting sum of monomials $x^i y^j$ all satisfy $qi + (q+1)j < r$. The basis of monomials on H given by $x^i y^j$ with $i < q^2$ and $j < q$ provides a canonical way to reduce monomials modulo I. To reduce $x^i y^j$, we perform the following steps. While $i \geq q^2$ or $j \geq q$, if $i \geq q^2$, reduce $x^i y^j$ to $x^{i-q^2+1} y^j$; if $j \geq q$, reduce $x^i y^j$ to $x^{i+q+1} y^{j-q} - x^i y^{j-q+1}$. At each step, either the degree of x is strictly decreasing or the degree of y is strictly decreasing, and the degree of y never increases, so this process will eventually terminate.

Lemma 7.4. *For $a \in \mathbb{N}$, let a_i denote the ith bit in the binary representation of a, i.e. $a = \sum_{i\geq 0} a_i 2^i$. Let $b = 2 + \lceil \log_2 m \rceil$. Let*

$$\text{Good} = \{(i_1, \ldots, i_m, j_1, \ldots, j_m) \mid \exists s \in [2\ell - c, 2\ell - b - 1] \, \forall t \in [0, b] \, \forall k \, (i_k)_{s+t} = (j_k)_{s+t} = 0\}.$$

If $(i_1, \ldots, i_m, j_1, \ldots, j_m) \in \text{Good}$, then $\prod_{k=1}^m x_k^{i_k} y_k^{j_k} \in \text{Lift}_\Phi^m(\mathcal{C})$.

Proof. For $a \in \mathbb{N}$, the condition $a < r = (1-2^{-c})q^3$ is equivalent to the condition $\exists s' \in [3\ell - c, 3\ell - 1]$ such that $a_{s'} = 0$. For each $k \in [m]$, fix $d_k \leq_p i_k$ and $(d'_k, e_k) \leq_p j_k$. The hypothesis implies that $(d_k)_{s+t} = (d'_k)_{s+t} = (e_k)_{s+t} = 0$ for all $t \in [0, b]$. It suffices to show that after reducing $x^{\sum_{k=1}^m d_k + d'_k} y^{\sum_{k=1}^m e_k}$ modulo I into a sum of monomials $x^i y^j$ with $i < q^2$ and $j < q$, each of them satisfies $(i)_t = (j)_t = 0$ for some $t \in [2\ell - c, 2\ell - 1]$, for this would imply

$$(qi + (q+1)j)_{t+\ell} = (qi)_{t+\ell} + ((q+1)j)_{t+\ell} = (i)_t + (j)_t = 0$$

and since $t + \ell \in [3\ell - c, 3\ell - 1]$ this implies the lemma.

Let $d = \sum_{k=1}^m d_k + d'_k$ and let $e = \sum_{k=1}^m e_k$. Consider three cases.

Case 1. $d < q^2$, $e < q$. In this case, the monomoial $x^d y^e$ does not reduce, so it suffices to show that $(d)_{s+b} = (e)_{s+b} = 0$. The only way one of these is 1 is by carrying from the lower order bits, so we may ignore the higher order bits and assume without loss of generality that $(d_k)_{s'} = (d'_k)_{s'} = (e_k)_{s'} = 0$ for $s' \geq s + b$. Then $d_k, d'_k, e_k < 2^s$, so $\sum_{k=1}^m d_k + d'_k < (m 2^{s+1}) < 2^{s+b}$ and thus $(d)_{s+b} = 0$ and similarly $\sum_{k=1}^m e_k < m 2^s < 2^{s+b}$ so $(e)_{s+b} = 0$.

Case 2. $d \geq q^2$, $e < q$. In this case, the monomial $x^d y^e$ reduces to $x^{d \bmod (q^2-1)} y^e$. By the previous case, $(e)_{s+b} = 0$, so it only remains to show $(d \bmod (q^2 - 1))_{s+b} = 0$. Doubling d cyclically permutes the bits of $d \bmod (q^2-1)$. In particular, $(2d \bmod (q^2-1))_i = (d \bmod (q^2-1))_{i-1 \bmod 3\ell-1}$. Then $(2^{3\ell-1-s-b} d \bmod (q^2-1))_i = (d \bmod (q^2-1))_{i+s+b+1-3\ell}$. Therefore, it suffices to show that $(2^{3\ell-1-s-b} d \bmod (q^2-1))_{3\ell-1} = 0$. Since the bits of order $[s, s+b]$ of d_k, d'_k are zero, the bits of order $[3\ell - 1 - b, 3\ell - 1]$ of $2^{3\ell-1-s-b}$ times d_k, d'_k, e_k are zero, hence $2^{3\ell-1-s-b} d_k \bmod (q^2 - 1) < 2^{3\ell-1-b}$ and similarly for d'_k. Therefore $\sum_{k=1}^m 2^{3\ell-1-s-b}(d_k + d'_k) \bmod (q^2-1) < 2^{3\ell-1}$ so $(2^{3\ell-1-s-b} d \bmod (q^2-1))_{3\ell-1} = 0$, and in fact $(2^{3\ell-1-s-b} d \bmod (q^2-1))_{3\ell-2} = 0$, so we can conclude that $(d \bmod (q^2-1))_{3\ell-1} = (d \bmod (q^2-1))_{3\ell-2} = 0$, which we need in Case 3.

Case 3. $e \geq q$. We induct on the $(q, q+1)$-weighted degree $qd + (q+1)e$. In this case, after reducing the y-degree by one step, the monomial reduces to $x^{d+q+1} y^{e-q} - x^d y^{e-q+1}$. The latter monomial has strictly smaller $(q, q+1)$-weighted degree, so by induction it is in $\mathcal{L}$. Thus it suffices to deal with $x^{d+q+1} y^{e-q}$. Repeating this reduction and ignoring the monomials with strictly smaller $(q, q+1)$-weighted degree, after at most m reductions (since $e_k < q$ and so $e < mq$) we have $x^{d+u(q+1)} y^{e \bmod q}$ for some $u \leq m$, which further reduces to $x^{d+u(q+1) \bmod (q^2-1)} y^{e \bmod q}$. This is almost Case 2, except for the additional $u(q+1)$ in the exponent of x. By Case 2, $(d \bmod (q^2 - 1))_{s+b-1} = (d \bmod (q^2-1))_{s+b} = 0$ and $(e \bmod q)_{s+b} = 0$. Note that since $\lceil \log_2 m \rceil < \ell - c$, $u(q+1) \leq m(q+1) < 2^{2\ell-c} + 2^{\ell-c} < 2^{s+1}$. Write $d \bmod (q^2-1)$ as $d' + 2^{s+b+1} d''$ where $d' < 2^{s+b-1}$. Then $d + u(q+1) \bmod (q^2-1) = d' + u(q+1) + 2^{s+b+1} d'' < 2^{s+b-1} + 2^{s+1} + 2^{s+b+1} d'' < 2^{s+b} + 2^{s+b+1} d''$ so $(d + u(q+1) \bmod (q^2-1))_{s+b} = 0$. □

Lemma 7.5. *Let* Good *be defined as in Lemma 7.4. Let $b = 2 + \lceil \log_2 m \rceil$. Then*

$$|\text{Good}| \geq q^{3m}(1 - (1 - 2^{-mb})^{c/b}).$$

Proof. We show the equivalent assertion that, by picking $i_1, \ldots, i_m < q^2$ and $j_1, \ldots, j_m < q$ uniformly at random, the probability that $(i_1, \ldots, i_m, j_1, \ldots, j_m) \in \text{Good}$ is $1 - (1 - 2^{-mb})^{c/b})$ at least. Note that each $j_k < q$ so we only need to consider the i_k. Partition $[3\ell - c, 3\ell - 1]$ into c/b intervals each of length b. Let E_i be the event that $(i_k)_t = 0$ for all $k \in [m]$ and all t in the ith interval. By Lemma 7.4, if $\bigvee_i E_i$ then $(i_1, \ldots, i_m, j_1, \ldots, j_m) \in \text{Good}$, so the probability of landing in Good is at least

$$\Pr\left[\bigvee_i E_i\right] = 1 - \Pr\left[\bigwedge_i \overline{E_i}\right] = 1 - (1 - 2^{-mb})^{c/b}.$$

□

Putting together Lemmas 7.4 and 7.5 with the discussion above, we immediately obtain the following.

Theorem 7.6. *Let $m \geq 1$, let $c > 0$ and let $\delta = 2^{-c}$. Let q be a power of 2 such that $\delta q > m$, and let $r = (1 - \delta)q^3$. Let Φ be the group of automorphisms on the Hermitian curve $H \subseteq \mathbb{F}_{q^2}^2$ of the form $(x, y) \mapsto (ax + b, a^{q+1}y + ab^q x + c)$ and let $\mathcal{L} = \text{Lift}_\Phi^m(\text{Herm}_q[r])$. Let $b = 2 + \lceil \log_2 m \rceil$. Then the rate of $\mathcal{L}$ is at least $1 - (1 - 2^{-mb})^{c/b} \geq 1 - e^{-c/(b2^{mb})}$.*

Putting everything together, we now prove Theorem 7.1.

Proof of Theorem 7.1. Fix $\epsilon, \alpha, N_0 > 0$. Recall that we want, for infinitely many $N \geq N_0$, a code of length N, rate $1 - \alpha$, alphabet size $N^{\epsilon/3}$, and is $(N^\epsilon, \Omega(1))$-locally correctable.

Set $m = \lceil 2/\epsilon \rceil$. Let $b = 2 + \lceil \log_2 m \rceil$ and set $c \geq b \cdot 2^{mb} \ln \frac{1}{\alpha}$. Let $\delta = 2^{-c}$, set q to be a power of 2 such that $\delta q > m$ and $q^{3m} \geq N_0$. Set $N = q^{3m}$ and set $r = (1 - \delta)q^3$. Let $\mathcal{L} = \text{Lift}_\Phi^m(\text{Herm}_q[r])$ where Φ is the usual automorphism group of the Hermitian curve $H \subseteq \mathbb{F}_{q^2}^2$. By our choice of parameters and Theorem 7.6, $\mathcal{L}$ has block length $q^{3m} = N$, rate at least $1 - e^{-c/b2^{mb}} \geq 1 - \alpha$, alphabet size $q^2 \leq N^{\epsilon/3}$, has query complexity $q^6 \leq N^\epsilon$, and can correct up to $\delta^2 = \alpha^{O((8/\epsilon)^{(2/\epsilon)} \log(1/\epsilon))}$.

□

Explicitness of code. Although a lifted code is not a priori explicit even if the base code is, Lemma 7.4 shows that the lifted Hermitian code (more accurately, a subcode with the same parameter guarantees) is explicit in the following way. Let Good be defined as in Lemma 7.4. The $\mathbb{F}_{q^2}$-span of monomials in Good have the same rate guarantees as the full lift, its block length and alphabet size and locality are the same, and certainly its distance is at least as good, since it is a subcode. Moreover, to encode a message $m \in \mathbb{F}_{q^2}^{\text{Good}}$ into a codeword $\mathsf{Enc}(m) \in \mathbb{F}_{q^2}^{H^m}$, first compute all the monomials in Good, which can be done by iterating over every monomial on H^m and checking if it is in Good, which can be done in polynomial time. Then interpret the symbols of m as coefficients of the monomials in Good and let $\mathsf{Enc}(m)$ be the evaluations of m on every point of H^m.

8 Conclusion

In this work, we presented a general framework for constructing high rate locally correctable codes. Our framework is an abstraction of affine lifting [3], automorphic lifting [1], and high-degree lifting [1]. We showed that the lift of a code with good distance with respect to some Φ that is close to doubly transitive also has good distance, and moreover this holds even when the base code is not invariant under Φ or when Φ is not a group. We showed how one can generalize the construction of the lifted Reed-Solomon code of [3] to lift other algebraic geometry codes, such as the Hermitian code to obtain locally correctable codes that can attain query complexity N^{ϵ} and rate $1-\alpha$ while correcting a constant fraction of errors, for any given $\epsilon, \alpha > 0$.

We believe the lifting framework deserves further study. Lifted codes naturally have good locality properties. A natural direction to explore is the local *testability* of lifted codes. A *local tester* is given oracle access to a word f and must distinguish whether $f \in \mathcal{C}$ or $\delta(f, \mathcal{C}) > \epsilon$ for some given constant $\epsilon > 0$. The work of [3] shows that affine lifting naturally yields affine-invariant locally testable codes. An interesting question is whether lifting algebraic geometry codes yields locally testable codes, and what kind of assumptions on Φ are necessary (for example, that the base code is Φ-invariant or that Φ is a group). In fact, [3] shows that both local correctability and local testability follows generically from affine lifting. In our work, local correctability follows generically from lifting — the instantiation of algebraic geometric base codes is only used to analyze the rate. It would be interesting to see if local testability follows generically from lifting as well.

Acknowledgements

Thanks to Ronald Cramer and Chaoping Xing for pointing out a bug in a proof analyzing the Hermitian tower code, which rendered the relevant sections obsolete. Thanks to Eli Ben-Sasson for helpful discussions and in particular for suggesting Theorem 4.6. Thanks to Madhu Sudan for his helpful comments and encouragement. Thanks to Badih Ghazi and Henry Yuen and for carefully proofreading parts of this paper and for helpful comments on the writing.

References

[1] Eli Ben-Sasson, Ariel Gabizon, Yohay Kaplan, Swastik Kopparty, and Shubhangi Saraf. A new family of locally correctable codes based on degree-lifted algebraic geometry codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:148, 2012.

[2] Eli Ben-Sasson, Ghid Maatouk, Amir Shpilka, and Madhu Sudan. Symmetric LDPC codes are not necessarily locally testable. In *IEEE Conference on Computational Complexity*, pages 55–65. IEEE Computer Society, 2011.

[3] Alan Guo, Swastik Kopparty, and Madhu Sudan. New affine-invariant codes from lifting. In Robert D. Kleinberg, editor, *ITCS*, pages 529–540. ACM, 2013.

[4] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 167–176. ACM, 2011.

[5] Irving S. Reed. A class of multiple-error-correcting codes and the decoding scheme. *IEEE Transactions on Information Theory*, 4:38–49, 1954.

[6] Henning Stichtenoth. *Algebraic function fields and codes.* Universitext. Springer, 1993.

[7] Chaoping Xing. On automorphism groups of the hermitian codes. *IEEE Transactions on Information Theory*, 41(6):1629–1635, 1995.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de