Low Degree Local Correction Over the Boolean Cube

Prashanth Amireddy* Amik Raj Behera† Manaswi Paraashar‡ Srikanth Srinivasan §
Madhu Sudan¶

Abstract

In this work, we show that the class of multivariate degree-d polynomials mapping $\{0,1\}^n$ to any Abelian group G is locally correctable with $\widetilde{\mathcal{O}}_d((\log n)^d)$ queries for up to a fraction of errors approaching half the minimum distance of the underlying code. In particular, this result holds even for polynomials over the reals or the rationals, special cases that were previously not known. Further, we show that they are locally list correctable up to a fraction of errors approaching the minimum distance of the code. These results build on and extend the prior work of Amireddy, Behera, Paraashar, Srinivasan, and Sudan [1] (STOC 2024) who considered the case of linear polynomials ($d = 1$) and gave analogous results.

Low-degree polynomials over the Boolean cube $\{0,1\}^n$ arise naturally in Boolean circuit complexity and learning theory, and our work furthers the study of their coding-theoretic properties. Extending the results of [1] from linear polynomials to higher-degree polynomials involves several new challenges and handling them gives us further insights into properties of low-degree polynomials over the Boolean cube. For local correction, we construct a set of points in the Boolean cube that lie between two exponentially close parallel hyperplanes and is moreover an interpolating set for degree-d polynomials. To show that the class of degree-d polynomials is list decodable up to the minimum distance, we stitch together results on anti-concentration of low-degree polynomials, the Sunflower lemma, and the Footprint bound for counting common zeroes of polynomials. Analyzing the local list corrector of [1] for higher degree polynomials involves understanding random restrictions of non-zero degree-d polynomials on a Hamming slice. In particular, we show that a simple random restriction process for reducing the dimension of the Boolean cube is a suitably good sampler for Hamming slices. Thus our exploration unearths several new techniques that are useful in understanding the combinatorial structure of low-degree polynomials over $\{0,1\}^n$.

1 Introduction

In this paper, we consider the local correction of low-degree polynomial functions over groups evaluated over $\{0,1\}^n$ and give polylogarithmic query local correctors for every constant degree. This extends and generalizes previous work of Amireddy, Behera, Paraashar, Srinivasan and Sudan [1] who considered and solved the analogous problem for the linear (i.e., $d = 1$) case. We define some of the basic terms and review the previous work before describing the challenges in strengthening to higher degrees and the new tools used to overcome them.

Low degree polynomials over groups. The main objects of interest in this paper are polynomial functions mapping $\{0,1\}^n$ to an Abelian group G. Here a function f is a polynomial of degree at most d if it can be expressed as $\sum_{S\subseteq[n]:|S|\leqslant d} c_S \prod_{i\in S} x_i$, where the product is over the integers and the coefficients c_S come from the Abelian group G. We denote the space of polynomial functions of degree at most d by $\mathcal{P}_d(\{0,1\}^n, G)$ (which we compress to $\mathcal{P}_d$ when G and n are known). The standard proof of the Ore-DeMillo-Lipton-Schwartz-Zippel lemma naturally extends to polynomials over groups. It proves that two different degree d polynomials disagree on at least $\delta_d := 2^{-d}$ fraction of the domain (if $d < n$), and thus form natural classes of error-correcting codes.

*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported in part by a Simons Investigator Award and NSF Award CCF 2152413 to Madhu Sudan and a Simons Investigator Award to Salil Vadhan. Email: pamireddy@g.harvard.edu

†Department of Computer Science, University of Copenhagen, Denmark. Supported by Srikanth Srinivasan's start-up grant from the University of Copenhagen. The work begun when the author was a student at the Department of Computer Science, Aarhus University, Denmark and was supported by Srikanth Srinivasan's start-up grant from the Aarhus University. Email: ambe@di.ku.dk

‡Department of Mathematical Sciences, University of Copenhagen, Denmark. Supported by the European Union under the Grant Agreement No 101078107, QInteract. Email: manaswi.isi@gmail.com

§Department of Computer Science, University of Copenhagen, Denmark. Also partially employed by Aarhus University, Denmark. This work was funded by the European Research Council (ERC) under grant agreement no. 101125652 (ALBA). Email: srsr@di.ku.dk

¶School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported in part by a Simons Investigator Award and NSF Award CCF 2152413. Email: madhu@cs.harvard.edu

This paper explores the corresponding correction questions focusing on locality.
A special case that is already of interest is when the group G is the group of real numbers (or rationals) - a setting where relatively few codes are shown to exhibit local correction properties.

Local correction of polynomials. Informally, the local correction problem is that of computing, given oracle access to a function $f : \{0,1\}^n \to G$ and a point $\mathbf{a} \in \{0,1\}^n$, the value $P(\mathbf{a})$ of the nearest degree d polynomial P to the function f at the point $\mathbf{a}$, while making few oracle queries to f. More formally, for functions $f, g : \{0,1\}^n \to G$, let $\delta(f,g)$ denote the fraction of points from the domain where they differ. We say f is ε-close to g if $\delta(f,g) \leqslant \varepsilon$ and ε-far otherwise. For a given G, we say that that $\mathcal{P}_d$ is (δ, q)-locally correctable if for every n there is a probabilistic algorithm that, for every function $f : \{0,1\}^n \to G$ that is $\delta = \delta(n)$-close to some polynomial $P \in \mathcal{P}_d(\{0,1\}^n, G)$ and for every $\mathbf{a} \in \{0,1\}^n$, outputs $P(\mathbf{a})$ with probability at least 3/4 while making at most $q = q(n)$ queries to f.
One of the main quests of this work is to give non-trivial upper bounds on the query complexity q for which $\mathcal{P}_d$ is $(\Omega_d(1), q)$-locally correctable.

List correction of codes. Note that (δ, q)-locally correctability of $\mathcal{P}_d$ requires that δ is less than half the minimum distance of the space, i.e., $\delta < \delta_d/2$. To go beyond one usually resorts to the notion of list-decoding; and in the local setting, to notions like "local list-decoding" and "local list correction". Roughly list-decoding allows the decoder to output a small list of words with the guarantee that all codewords within a given distance are included in the output. Formally we say $\mathcal{P}_d$ is (combinatorially) (δ, L)-list correctable if for every $f : \{0,1\}^n \to G$ there are at most L degree-d polynomials P satisfying $\delta(f,P) \leqslant \delta$.[1] Unlike the unique decoding problem where the maximum δ such that a code is uniquely correctable up to δ errors is well understood, the list-decoding radius for higher values of L is not well-understood. A natural question that we study here (for the first time in this generality) is: *What is the largest δ such that $\mathcal{P}_d$ is $(\delta, \mathcal{O}_d(1))$-list correctable?* We refer to this largest value of δ as the list-decoding radius of $\mathcal{P}_d$.

Local list correction of codes. Local list correction is the notion of list decoding combined with the notion of local correction. Formalizing this definition is a bit more subtle and was first done in [21]. The notion allows the decoder to work in two phases — a preprocessing phase with $q_1 = q_1(n)$ queries to the function f, that outputs up to L algorithms $\phi_1, \ldots, \phi_L$ and a query phase, where given $\mathbf{a} \in \{0,1\}^n$ each algorithm ϕ_i makes $q_2 = q_2(n)$ queries to f and outputs $\phi_L(\mathbf{a})$. We say that $\mathcal{P}_d$ is (δ, q_1, q_2, L)-local list correctable if for every function f and polynomial $P \in \mathcal{P}_d$ that are δ-close, there is a decoder as above such that one of its outputs includes P with high probability (say 3/4). The final goal of this paper is to locally list-correct $\mathcal{P}_d$ using non-trivially small number of queries (in both the preprocessing and query phases) where the fraction of errors approaches the list-decoding radius.

1.1 Motivation and previous work

Local decoding of polynomials over finite fields has played a central role in computational complexity and in particular in breakthrough results like IP=PSPACE and the PCP theorem. While most of these results consider functions over the entire multivariate domain (i.e., $\mathbb{F}^n$), low-degree polynomials over $\{0,1\}^n$ do arise quite naturally in complexity theory, notably in circuit complexity capturing classes like AC^0 [18, 20] and ACC [6], and in learning theory. Many of these results exploit basic distance properties of multivariate polynomials as given by the Ore-DeMillo-Lipton-Schwartz-Zippel lemma. This lemma roughly says that the space of degree-d polynomial functions mapping S^n to a field $\mathbb{F}$ where $S \subseteq \mathbb{F}$ is finite form an error-correcting code of relative distance $d/|S|$ when $d < |S|$, and $|S|^{-d/(|S|-1)}$ when $d \geqslant |S|$.

The special case of $S = \mathbb{F}$ is extensively studied and heavily used, e,g., in PCPs and cryptography. In this setting, the lemma can also be made algorithmic, with the first such instance handling the special case of $\mathbb{F}_2$ dating back to the works of Reed and Muller [19, 16]. More recently, local list correction algorithms were discovered in the works of Goldreich and Levin [12] for linear polynomials, Sudan, Trevisan and Vadhan [21] for higher-degree polynomials over large fields, Gopalan, Klivans and Zuckerman [13] for higher-degree polynomials

[1] The algorithmic version would require the list of nearby polynomials to be algorithmically recoverable from f. We don't consider this notion in this work but move ahead to the harder "local list correction" problem.

over $\mathbb{F}_2$, and Bhowmick and Lovett [5] for polynomials over any small finite field.

The case of general S however has not received much attention and is mostly unexplored. This was first highlighted in a relatively recent work of Kim and Kopparty [14] who gave polynomial time (but not local) unique-decoding algorithms correcting errors up to half the minimum distance. Their work exposed the fact that many other algorithmic and even some coding-theoretic questions were not well understood when $S \neq \mathbb{F}$, and our work aims to fill some gaps in knowledge here.

Another motivation for our work is the design of locally correctable codes over the reals. A series of works [4, 8, 9] has exposed that there are no known $(\Omega(1), \mathcal{O}(1))$-locally correctable codes over the reals of arbitrarily large dimensions. The underlying challenge here leads to novel questions in incidence geometry. Roughly the goal here is to design a finite set of points $T \subseteq \mathbb{R}^n$ such many pairs of points in T are contained in constant-dimensional "degenerate" subspaces, where a q-dimensional subspace is said to be degenerate if it contains $q+1$ points from T. Till recently no sets that possessed this property with $q = o(n)$ were known, and the recent results of [1] may be viewed as showing that the set $T = \{0,1\}^n \subseteq \mathbb{R}^n$ has $\widetilde{O}(\log n)$ dimensional subspaces covering most pairs of points of T. Local correctability of degree d polynomials would translate to showing that moment vectors[2] of $\{0,1\}^n$ (viewed as vectors in $\mathbb{R}^N$ for $N = \mathcal{O}(n^d)$) also exhibit a similar property, thus adding to the body of sets in $\mathbb{R}^N$ that show non-trivial degeneracies.

Turning to previous work on local correction of polynomials over grids, the local correction question when $S = \{0,1\}$ was first explored by Bafna, Srinivasan and Sudan [3], who mainly presented a lower bound of $\widetilde{\Omega}(\log n)$ on the number of queries to recover even when $d = 1$ and from some $o(1)$ fraction of errors and $\mathbb{F} = \mathbb{R}$. On the positive side, for fields of characteristic p, they gave an $\mathcal{O}_{d,p}(1)$ query algorithm to recover from $\Omega_{d,p}(1)$ fraction of errors. This left the case for large and zero characteristic fields open.

The recent work of Amireddy, Behera, Paraashar, Srinivasan, and Sudan [1] investigated the case of general fields, and more generally, polynomials over Abelian groups, for the special case of $d = 1$. For this setting, they consider all three questions posed in the previous section, namely the (unique) local correction limit, the list-decoding radius, and the local list correction problem. In this setting where distinct degree 1 polynomials disagree with each other on at least half the domain, they show that up to $1/4$ fraction of errors can be uniquely locally corrected with $\widetilde{\mathcal{O}}(\log n)$ queries. They further show that the list-decoding radius approaches $1/2$ and that there are $\widetilde{\mathcal{O}}(\log n)$ query algorithms to locally list correct $\mathcal{P}_1$ for any fraction of errors bounded away from $1/2$. Their work exposes a number of technical challenges in going beyond the $d = 1$ case and we address those in this paper.

1.2 Technical challenges in extending beyond the linear case We start by reviewing the main ideas in [1] and outlining the challenges in the higher-degree extension. Their unique local corrector correcting up to half the minimum distance works in three steps: Given an oracle for a function $f(\mathbf{x})$ at a distance less than $1/4$ from a linear polynomial $P(\mathbf{x})$,

- It first provides oracle access to a function $f_1(\mathbf{x})$ at any tiny but constant distance from $P(\mathbf{x})$, using $\mathcal{O}(1)$ queries to $f(\mathbf{x})$
- Then the algorithm provides oracle access to a function $f_2(\mathbf{x})$ at distance $1/\operatorname{poly}(\log n)$ from $P(\mathbf{x})$ while making $\operatorname{poly}(\log \log n)$ queries to $f_1(\mathbf{x})$
- Finally, the algorithm provides oracle access to the linear polynomial $P(\mathbf{x})$ making $\mathcal{O}(\log n)$ queries to f_2. Composing the three steps gives the desired unique local corrector.

The first two steps in their result are general enough to work for all degrees. The third step in the unique local corrector of [1] is the most significant one and does not extend immediately to the higher degree setting. [1] reduce this step to show that any point in $\{0,1\}^n$ can be expressed as a linear combination of $\widetilde{\mathcal{O}}(\log n)$ roughly

[2] d-moment vector of $\mathbf{v} \in \{0,1\}^n$ is a vector in $\{0,1\}^{\binom{n}{d}}$ which is an evaluation vector of $\mathbf{v}$ on all multilinear monomials of degree $\leqslant d$.

balanced vectors[3]. To extend their approach to higher degrees, we would need an analogous result for the degree $\leqslant d$-moment vectors of vectors in $\{0,1\}^n$ but we are unable to find such an extension directly. Instead, as we elaborate further below, we manage to find a new path for this step, which results in a somewhat different proof even for the linear ($d = 1$) case.

Next, we turn to the combinatorial analysis of the list-decoding radius. For simplicity assume that the Abelian group G is a finite field $\mathbb{F}_p$. The analysis in the linear case [1] splits into two cases: the low[4] characteristic case ($p \leqslant 3$) and the high characteristic case ($p > 5$). The former case is handled via a suitable version of the Johnson bound. In the latter setting, a key insight used in [1] for the linear case is that non-sparse linear polynomials tend to be non-zero with very high probability, i.e. *anti-concentration* of non-sparse linear polynomials. In the higher degree setting, both analyses become much more involved. In the low characteristic case, the Johnson bound no longer yields the right answer. For the high characteristic case, the primary obstacle is understanding the anti-concentration statement for non-sparse low-degree polynomials.

Generalizing the result above to all Abelian groups involves multiple steps in [1] - they extend the latter approach above to groups where every element has a sufficiently high order (specifically order at least 5). Then they consider groups where every element has order a power of 2 or 3 separately and analyze these special cases; and finally use some "special intersection properties" of the agreement sets[5] of different polynomials with any given function to apply a counting tool from the work of Dinur, Grigorescu, Kopparty, and Sudan [7] to combine these different steps. While many of the steps extend to the higher degree setting (sometimes needing heavier machinery) the final step involving "special intersection properties" simply does not work in our setting. (Roughly the difference emanates from the fact that the probability that two linearly independent degree-1 functions vanish at a random point is at most $1/4$, which is the square of the probability for a single degree-1 function. This fails for degree 2.) Overcoming this barrier leads to further new challenges in the higher-degree case.

The local list-correction algorithm for degree-1 polynomials of [1] is inspired by the local list-correction algorithm of Reed-Muller codes from [21]. The high-level idea is to start with a function $f(\mathbf{x})$ that is $(1/2 - \varepsilon)$-close to a set List of linear polynomials and produce a small list of oracles such that each polynomial in List is within a small constant distance to an oracle from the list, at which point the unique local corrector becomes applicable. This 'error-reduction' step involves choosing a random subcube C of $\{0,1\}^n$ (as defined in [2, Section 2]) of sufficiently large but constant dimension k and doing a brute-force list decoding on C to find a list List': it is not hard to argue that the restricted version P' of each polynomial $P \in \mathsf{List}$ appears in the list List' with high probability (this needs the combinatorial bound mentioned above). To complete the error-reduction, they need to decode P at a random point $\mathbf{b} \in \{0,1\}^n$. This is done by repeating the above brute-force algorithm with the subcube $\mathsf{C}^{\mathbf{b}}$ 'spanned' by C and $\mathbf{b}$: informally, this is the smallest subcube spanned by C and $\mathbf{b}$ and has dimension $2k$ (see [2, Section 5.2.1] for details.) They now obtain a new list of polynomials List'' and need to isolate the polynomial P'' corresponding to P in this list to get $P(\mathbf{b})$. Here, [1] uses the fact that we know the restriction of P to the subcube C inside $\mathsf{C}^{\mathbf{b}}$. The bad event in this case is that there are two polynomials in $\mathcal{L}''$ that disagree on $\mathbf{b}$ but agree on C. Bounding the probability of this event is the key step in the analysis of [1] and is done by giving a complete understanding of which kinds of distinct polynomials over $\mathsf{C}^{\mathbf{b}}$ can collapse to the same polynomial over C. This kind of understanding seems difficult to obtain for higher degrees, and we need to develop a new analysis for bounding the probability of the bad event in this setting.

We now turn to our results before elaborating on the techniques used to overcome the challenges. We omit the proofs here and just state the main results and lemmas. Full version of the paper can be found at [2].

[3]Hamming weight is very close to $n/2$

[4]The precise constants here are not important, as the analysis works as long as the $p \leqslant \mathcal{O}(1)$.

[5]Agreement set of a polynomial P and a function f is defined as the subset of $\{0,1\}^n$ on which P and f agree.

1.3 Our main results Briefly, our results provide poly-logarithmic query algorithms for unique and list-decoding to the maximal fraction of errors that are allowed. Specifically, the unique decoding algorithm works up to half the distance. We also establish that the list-decoding radius approaches the distance (as the list size tends to infinity) and give matching local algorithms. We give our specific theorems below.

THEOREM 1.1. (LOCAL CORRECTION ALGORITHMS FOR $\mathcal{P}_d$ UP TO THE UNIQUE DECODING RADIUS) *For every Abelian group G and for every constant $\varepsilon > 0$, the space $\mathcal{P}_d$ has a (δ, q)-local correction algorithm where $\delta = \frac{1}{2^{d+1}} - \varepsilon$ and $q = \widetilde{\mathcal{O}}_\varepsilon(\log n)^d$.*

As noted earlier, some dependence on n is needed even when the degree is 1 and G is a field of large characteristic (or characteristic 0), as an $\Omega(\log n/\log\log n)$ lower bound was shown in this setting by earlier work of Bafna, Srinivasan, and Sudan [3] (and shown to be tight up to $\mathrm{poly}(\log\log n)$ factors in [1]). Our upper bound above is thus optimal to within polynomial factors (for constant d). However, we do not know if the query complexity can be improved to, say, $\widetilde{\mathcal{O}}_d(\log n)$ for degree d. Refer to [2, Section 3] for a proof.

We also extend the above algorithm from Theorem 1.1 to the list decoding regime. For this, we first establish a bound on the list-decoding radius. As far as we know, the following result was not known before for G being any group other than the field $\mathbb{F}_2$. We state the theorem below and refer to [2, Section 4] for a proof.

THEOREM 1.2. (COMBINATORIAL LIST DECODING BOUND FOR $\mathcal{P}_d$) *For every Abelian group G and for every constant $\varepsilon > 0$, the space $\mathcal{P}_d$ over any Abelian group G is $(1/2^d - \varepsilon, \exp(\mathcal{O}_d(1/\varepsilon)^{\mathcal{O}(d)})$-list correctable.*

The above is tight in the sense that the number of codewords at distance $1/2^d$ can depend on both n and the size of the group G (and is infinite when G is infinite). We do not know if the dependence on ε is tight. Note that for the setting where $d = 1$, [1] gives a polynomial bound in terms of $1/\varepsilon$. Our bound as stated above is exponential and while we can see a path to improving this to a quasi-polynomial, we don't see a polynomial upper bound using the proofs of this paper even when $d = 1$.

Finally, we state our local list correction result. Refer to [2, Section 5] for proof.

THEOREM 1.3. (LOCAL LIST CORRECTION FOR $\mathcal{P}_d$) *For every Abelian groups G and for every $\varepsilon > 0$, the space $\mathcal{P}_d$ is $(1/2^d - \varepsilon, \mathcal{O}_\varepsilon(1), \widetilde{\mathcal{O}}_\varepsilon(\log n)^d, \exp(\mathcal{O}_d(1/\varepsilon)^{\mathcal{O}(d)}))$-locally list correctable.*

Specifically, there is a randomized algorithm $\mathcal{A}$ that, when given oracle access to a polynomial f and a parameter $\varepsilon > 0$, outputs with probability at least $3/4$ a list of randomized algorithms $\phi_1, \ldots, \phi_L$ ($L \leqslant \exp(\mathcal{O}_d(1/\varepsilon)^{\mathcal{O}(d)})$) such that the following holds. For each $P \in \mathcal{P}_d$ that is $(1/2^d - \varepsilon)$-close to f, there is at least one algorithm ϕ_i that, when given oracle access to f, computes P correctly on every input with probability at least $3/4$.

The algorithm $\mathcal{A}$ makes $\mathcal{O}_\varepsilon(1)$ queries to f, while each ϕ_i makes $\widetilde{\mathcal{O}}_\varepsilon((\log n)^d)$ queries to f.

1.4 Technical tools In the process of proving our main results, we prove several lemmas that we believe are independently interesting.

In the proof of Theorem 1.1, the main step is to construct, for any $\mathbf{a} \in \{0,1\}^n$, a distribution $\mathcal{D}_\mathbf{a}$ over $(\{0,1\}^n)^q$ such that the marginal distribution of each point is close to the uniform distribution over $\{0,1\}^n$, and for any degree-d polynomial P, $P(\mathbf{a})$ can be computed via the evaluations of P on a sample from $\mathcal{D}_\mathbf{a}$. Constructing such a distribution $\mathcal{D}$ reduces to the following problem of constructing a geometric set with some nice algebraic properties. We discuss in [2, Section 3] how such a set leads to the distribution $\mathcal{D}_\mathbf{a}$.

Question 1. *Find two parallel hyperplanes in k dimensions that are ε-close in Euclidean distance such that every non-zero degree-d multilinear polynomial is non-zero on the points of the Boolean cube $\{0,1\}^k$ lying between the two hyperplanes.*

The 'closeness' parameter ε plays a crucial role in the efficiency of the local correction algorithm. It is easy to see (and folklore) that we can get $\varepsilon = 1/k^{\Omega(1)}$. However, we show that we can obtain a construction where ε is *exponentially small* in k. Formally, we have the following lemma. Refer to [2, Section 3.2] for the proof.

LEMMA 1.1. (MAIN LEMMA FOR LOCAL CORRECTION) *Fix a degree parameter $d \geqslant 0$ and a dimension parameter $k \in \mathbb{Z}_{>0}$ such that k is divisible by $10 \cdot (d+1)$. There exists a set $\mathcal{S} \subseteq \{0,1\}^k$ such that for every Abelian group G, the set $\mathcal{S}$ satisfies the following properties:*

- *For every non-zero degree-d polynomial $Q(y_1, \ldots, y_k) \in \mathcal{P}_d(\{0,1\}^k, G)$, there exists a point $\mathbf{z} \in \mathcal{S}$ such that $Q(\mathbf{z}) \neq 0$.*
- *There exists positive integers $w_1, \ldots, w_k$ such that*

$$\mathcal{S} \subseteq \left\{ \mathbf{y} \in \{0,1\}^k \;\middle|\; \left| \sum_{j=1}^{k} w_j y_j - \frac{W}{2} \right| \leqslant \frac{W}{2^{\Omega(k/(d+1))}} \right\},$$

where $W := \sum_{j=1}^{k} w_j$.

Furthermore, $|\mathcal{S}|$ is at most $\mathcal{O}_d(k^d)$.

The proof of the combinatorial list decoding bound is broken down into two cases depending on the order of elements in the group. The first case is when all elements have order larger than a prime $p_0(d)$ (a constant dependent on d) and the second case is when the group is a product of p-groups for $p < p_0$. For the first case, the key step is an understanding of the anti-concentration properties of low-degree non-sparse polynomials. More precisely, we have the following question.

Question 2. *Let $P(x_1, \ldots, x_n)$ be a polynomial of degree d involving at least s monomials. Then how large can $\Pr_{\mathbf{a} \sim \{0,1\}^n}[P(\mathbf{a}) = 0]$ be?*

If we take $P = x_1 x_2 \cdots x_{d-1} \cdot L(x_d, \ldots, x_n)$ where L is a linear polynomial with s monomials, we see that P vanishes with probability approximately $1 - 2^{-(d-1)} - \mathcal{O}(1)/\sqrt{s}$ over (say) the reals. In Lemma 1.2, we build on known anti-concentration results [10, 15] show that this is essentially the worst possible in groups with no elements of small order. Refer to [2, Section 4.1] for the proof.

LEMMA 1.2. (ANTI-CONCENTRATION BOUND FOR NON-SPARSE POLYNOMIALS) *For every positive integers d, s and for every Abelian group G in which all the non-zero elements have order greater than $(s+1)!$, the following holds: For each degree-d polynomial $Q(\mathbf{x})$ over G of sparsity at least s, we have*

$$\Pr_{\mathbf{x} \sim \{0,1\}^n}[Q(\mathbf{x}) \neq 0] \geqslant 1/2^{d-1} - 2^{\mathcal{O}(d^3)}/\sqrt{s}.$$

In the second case, an important step in our proof is a tail inequality for events defined by the vanishing of low-degree polynomials over a field. Let $Q_1(\mathbf{x}), \ldots, Q_t(\mathbf{x})$ be t degree-d polynomials on the same variable set. We know that each is non-zero on a random input with probability at least 2^{-d} and hence that the expected number of polynomials vanishing at a uniformly random point $\mathbf{a} \sim \{0,1\}^n$ is at most $(1 - 2^{-d}) \cdot t$. This leads to the following question.

Question 3. *Suppose we have a collection of t degree-d polynomials. Can we bound the probability that more than $(1 - 2^{-d} + \varepsilon) \cdot t$ many of these polynomials vanish at a random point $\mathbf{a}$ in $\{0,1\}^n$?*

Clearly, we cannot get a strong tail bound unless we impose some 'independence' constraints on the polynomials (for example, we cannot hope for a strong bound if all the polynomials are in the linear span of a small number of polynomials). We show that we can get a Chernoff-style tail bound under the constraint that the 'leading monomials' (under a suitable ordering) of these polynomials are pairwise disjoint. This is done in [2, Section 4.2.1] using the 'Footprint bound' [11] (essentially a tool from commutative algebra) and an idea due to Panconesi and Srinivasan [17].

We build on this result to prove an optimal bound on the list-decoding radius for degree-d polynomials over small finite fields $\mathbb{F}_p$ (and more generally over groups that are products of p-groups, where $p < p_0$). In the setting when we are working with polynomials mapping $\mathbb{F}_p^n$ to $\mathbb{F}_p$, this was done in the works of [13, 5] via an involved mixture of algebraic and analytic techniques. Unfortunately, these do not seem to be applicable here: one significant reason that appears again and again in our work is that we cannot restrict a given function to an arbitrary subspace in our ambient space since the domain $\{0,1\}^n$ does not have this algebraic structure. Instead we use other combinatorial techniques such as the Sunflower lemma in conjunction with the above tail bound to obtain the stated result.

For local list correction, we follow the algorithm of [1] modulo changes in parameters to handle higher degree polynomials. The main innovation involves analyzing the behavior of degree-d polynomials on a Hamming slice (points with a fixed Hamming weight) after a random process of reducing the dimension. In particular, assume that we have a fixed degree-d polynomial $R(x_1, \ldots, x_{2k})$ in $2k$ dimensions such that R is non-zero on the Hamming slice of weight k. We now choose a random subcube C by pairing the $2k$ variables at random into k pairs and identifying the variables in each pair. We would like to upper bound the probability that R is zero on the cube C by a function that goes to 0 with k.[6] We do this by addressing the following two questions.

The first question is on how the density of any fixed set on a slice changes under the aforementioned random process. In this work, we are particularly interested in the middle slice, i.e. points of Hamming weight k in $\{0,1\}^{2k}$.

Question 4. *For any fixed subset S of the middle slice, how does the density of the set $S \cap \mathsf{C}$ (as a subset of the middle slice of C) compare with the density of S?*

In [2, Section 5.1.1] we show that the density is almost preserved under the random process. In other words, this random process is a good sampler for subsets of the middle slice. To prove this lemma, we show that certain kinds of *Johnson graphs* are good spectral expanders.

The second question is on a quantitative estimate of the number of non-zero points of a degree-d polynomial on a Hamming slice. This is a natural question, but has not been addressed before as far as we know.

Question 5. *For a degree-d polynomial R which is non-zero on a Hamming slice, on how many points of the Hamming slice is it non-zero?*

We give a simple lower bound on the number of non-zero points for a degree-d polynomial on the Hamming slice by modifying the proof of the Ore-DeMillo-Lipton-Schwartz-Zippel lemma. We prove the following lemma and refer to [2, Section 5.2.1].

LEMMA 1.3. (A DLSZ LEMMA OVER HAMMING SLICES) *The following holds for any non-negative integers integers n, k, d where $k \leqslant n$ and $d \leqslant \min\{k, n-k\}$. Let $R \in \mathcal{P}_d(\{0,1\}^n, G)$ be a polynomial such that R does not vanish at some point in $\{0,1\}_k^n$. Then*

$$|\{\mathbf{a} \in \{0,1\}_k^n \mid R(\mathbf{a}) \neq 0\}| \geqslant \binom{n-2d}{k-d}.$$

[6] One might hope to prove such a statement under the weaker assumption that R is simply a non-zero multilinear polynomial of degree d. Unfortunately, the simple example $R = x_1 + \cdots + x_{2k}$ over the group $G = \mathbb{F}_2$ shows that such a statement is not possible even in the degree-1 case.

References

[1] Prashanth Amireddy, Amik Raj Behera, Manaswi Paraashar, Srikanth Srinivasan, and Madhu Sudan. Local Correction of Linear Functions over the Boolean Cube . *Electron. Colloquium Comput. Complex.*, TR24-056, 2024.

[2] Prashanth Amireddy, Amik Raj Behera, Manaswi Paraashar, Srikanth Srinivasan, and Madhu Sudan. Low Degree Local Correction Over the Boolean Cube . *Electron. Colloquium Comput. Complex.*, TR24-164, 2024.

[3] Mitali Bafna, Srikanth Srinivasan, and Madhu Sudan. Local decoding and testing of polynomials over grids. *Random Struct. Algorithms*, 57(3):658–694, 2020.

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

[5] Abhishek Bhowmick and Shachar Lovett. The list decoding radius for Reed-Muller codes over small fields. *IEEE Trans. Inf. Theory*, 64(6):4382–4391, 2018.

[6] Abhishek Bhrushundi, Kaave Hosseini, Shachar Lovett, and Sankeerth Rao. Torus polynomials: An algebraic approach to ACC lower bounds. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 13:1–13:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[7] Irit Dinur, Elena Grigorescu, Swastik Kopparty, and Madhu Sudan. Decodability of group homomorphisms beyond the Johnson bound. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 275–284, New York, NY, USA, 2008. Association for Computing Machinery.

[8] Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. Improved rank bounds for design matrices and a new proof of Kelly's theorem. *Forum Math. Sigma*, 2:Paper No. e4, 24, 2014.

[9] Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. Superquadratic lower bound for 3-query locally correctable codes over the reals. *Theory Comput.*, 13:Paper No. 11, 36, 2017.

[10] Paul Erdős. On a lemma of littlewood and offord. *Bulletin of the American Mathematical Society*, 51(12)(4):898–902, 1945.

[11] O. Geil and T. Hoholdt. Footprints or generalized bezout's theorem. *IEEE Transactions on Information Theory*, 46(2):635–641, 2000.

[12] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *ACM Symposium on Theory of Computing (STOC)*, pages 25–32, 1989.

[13] Parikshit Gopalan, Adam R. Klivans, and David Zuckerman. List-decoding Reed-Muller codes over small fields. In *ACM Symposium on Theory of Computing (STOC)*, pages 265–274, 2008.

[14] John Y. Kim and Swastik Kopparty. Decoding Reed-Muller codes over product sets. *Theory Comput.*, 13(1):1–38, 2017.

[15] Raghu Meka, Oanh Nguyen, and Van Vu. Anti-concentration for polynomials of independent random variables. *Theory of Computing*, 12(11):1–17, 2016.

[16] David E. Muller. Application of boolean algebra to switching circuit design and to error detection. *Trans. I R E Prof. Group Electron. Comput.*, 3(3):6–12, 1954.

[17] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff–hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.

[18] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematicheskie Zametki*, 41(4):598–607, 1987. (English translation in *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987).

[19] Irving S. Reed. A class of multiple-error-correcting codes and the decoding scheme. *Trans. IRE Prof. Group Inf. Theory*, 4:38–49, 1954.

[20] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82. ACM, 1987.

[21] Madhu Sudan, Luca Trevisan, and Salil P. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001.