Electronic Colloquium on Computational Complexity, Report No. 37 (2009)

A Fourier-analytic approach to Reed-Muller decoding.

Parikshit Gopalan.
MSR-Silicon Valley
MountainView, CA.
parik@microsoft.com

April 9, 2009

Abstract

We present a Fourier-analytic approach to list-decoding Reed-Muller codes over arbitrary finite fields. We prove that the list-decoding radius for quadratic polynomials equals $1-2/q$ over any field $\mathbb{F}_q$ where $q > 2$. This confirms a conjecture due to Gopalan, Klivans and Zuckerman [GKZ08] for degree 2. Previously, tight bounds for quadratic polynomials were known only for $q = 2, 3$; the best bound known for other fields was the Johnson radius which is roughly $1 - 1/\sqrt{q}$.

We say that a polynomial over $\mathbb{F}_q$ is k-dimensional if it can be expressed as a function of k linear functions. We reduce the Reed-Muller list-decoding problem to list-decoding low-dimensional polynomials and present a new Fourier-based algorithm for the low-dimensional case. The list-decoding radius achieved by this approach for degree 3 and higher depends on questions regarding the weight-distribution of the Reed-Muller code. We propose a conjecture in this regard, which if true, improves on the best known bounds for the list-decoding radius for all d and q. The conjecture holds true for $\mathbb{F}_2$, giving an alternate proof of the main result of [GKZ08].

Departing from previous work on Reed-Muller decoding which relies on some form of self-corrector [GRS00, AS03, STV01, GKZ08], our work applies ideas from Fourier analysis of Boolean functions to low-degree polynomials over finite fields. We believe that the techniques used here could find other applications, we present applications to testing and learning.

ISSN 1433-8092

1 Introduction

Traditional algorithms to decode error-correcting codes require that the received word is within less than half the minimum distance of a codeword, so that the codeword can be uniquely recovered. In the 1950s, Elias [Eli57] and Wozencraft [Woz58] introduced the notion of list-decoding in order to decode beyond this barrier. Rather than returning a single codeword, a list-decoding algorithm outputs all codewords within a specified radius of a received word. It took over thirty years before Goldreich and Levin [GL89] and Sudan [Sud97] gave efficient list-decoding algorithms for Hadamard codes and Reed-Solomon codes, respectively. Since these breakthroughs, there has been much progress in devising list-decoders for various codes [Gur04, Gur06, Sud00]. Indeed, list-decoding algorithms are the only tools that we have for solving the nearest codeword problem beyond half the minimum distance in the adversarial error model.

Algorithms for list-decoding error-correcting codes have proved tremendously useful in computer science (see [Gur04, Chapter 12]), with applications ranging from hardness amplification for weakly hard functions [STV01, Tre03], constructions of hard-core predicates from any one-way function [GL89, AGS03], constructions of extractors and pseudorandom generators [TSZS01, SU05] and the average-case hardness of the permanent [Lip89]. Despite the considerable progress in this area, for several natural and well-studied families of codes including Reed-Solomon and Reed-Muller codes, the *list-decoding* radius, or the largest error radius up to which the list-decoding problem is tractable is as yet unknown. This problem for Reed-Muller codes is the focus of our paper.

Reed-Muller codes were discovered by Muller in 1954. The message space of the code $\mathsf{RM}_q(n,d)$ consists of all degree d polynomials in n variables over $\mathbb{F}_q$, the codewords are the evaluations of these polynomials at all points in $\mathbb{F}_q^n$. Let $\delta_q(d)$ denote the normalized minimum distance of $\mathsf{RM}_q(n,d)$. If $d = a(q-1)+b$ where $0 \leqslant b \leqslant q-1$, then

$$\delta_q(d) = \frac{1}{q^a}\left(1 - \frac{b}{q}\right). \tag{1}$$

The case when $d < q$ is the famous Schwartz-Zippel lemma.

Reed-Muller codes are one of the most well-studied families of error-correcting codes in coding theory [MS77, Ass92]. They are also ubiquitous in computer science, indeed several of the aforementioned applications of list-decoding [Lip89, GL89, STV01, TSZS01, SU05] use Reed-Muller codes. A closely related problem is that of low-degree testing, where we are given a function and asked to test if it is close to a codeword in the Reed-Muller code. This is a problem that has been studied extensively in computer science [BLR93, AS03, AKK+05, JPRZ04, KR04, Sam07], and plays in a key role in the original proof of the celebrated PCP theorem [ALM+98, AS98].

For most applications above, the model of interest is the local-decoding model where we are given an oracle for the received word $R : \mathbb{F}_q^n \to \mathbb{F}_q$ that can be queried at chosen points. The goal is to devise an algorithm whose running time is polynomial in the size of the message rather than the size of the codeword. The message being a degree d polynomial

(d will be constant) in n variables over $\mathbb{F}_q$, our goal is to run in time poly(n). So we are interested in the settings where the list-size is a constant, or at worst poly(n). Our running times are typically polynomial in q.

1.1 Previous Work

For (a family of) codes $\mathcal{C} \subset [q]^n$, let $\ell(\mathcal{C}, \eta)$ denote the maximum list-size at radius η (radius $\eta \in [0,1]$ denotes normalized Hamming distance). $\mathsf{LDR}(\mathcal{C})$ is the largest η for which $\ell(\mathcal{C}, \eta - \varepsilon)$ can be bounded by a function of ε (independent of n) for every $\varepsilon > 0$.

The study of list-decoding algorithms for Reed-Muller codes was initiated by the seminal work of Goldreich and Levin on list-decoding Hadamard codes over $\mathbb{F}_2$ or equivalently $\mathsf{RM}_2(n,1)$ codes [GL89]. They showed that $\mathsf{LDR}(\mathsf{RM}_2(n,1)) = 1/2$. Goldreich, Rubinfeld and Sudan generalized this to Hadamard codes over $\mathbb{F}_q$, showing that $\mathsf{LDR}(\mathsf{RM}_q(n,1)) = 1 - 1/q$ [GRS00]. An important development was the discovery of powerful algorithms for list-decoding univariate polynomials over $\mathbb{F}_q$, due to Sudan [Sud97] and Guruswami and Sudan [GS99]. Sudan, Trevisan and Vadhan used these algorithms to devise a list-decoder that works up to radius $1 - \sqrt{2d/q}$ for [STV01], improving on work by Arora and Sudan [AS03] and Goldreich *et al.* [GRS00] (see also[PW04]).

All of the aforementioned decoding algorithms reach a coding theoretic bound known as the Johnson bound [Joh62, Joh63]. The Johnson bound guarantees that for *any* code of minimum distance δ over $\mathbb{F}_q$, $\mathsf{LDR}(\mathcal{C}) \geqslant \mathsf{J_q}(\delta) = (1 - 1/q)(1 - \sqrt{1 - q\delta/(q-1)})$. Since the Johnson bound is oblivious to the structure of the code apart from its minimum distance, one does not expect it to be tight for every code, yet examples of codes decodeable beyond the Johnson bound are relatively few and recent (see the discussion in[DGKS08, GKZ08]). A tantalizing open problem in this area is whether the Johnson bound is tight for Reed-Solomon codes, this is precisely the radius achieved by the Guruswami-Sudan algorithm [GS99].

Recently, Gopalan, Klivans and Zuckerman (GKZ) considered the problem of list-decoding Reed-Muller codes over $\mathbb{F}_2$ [GKZ08]. They showed that $\mathsf{LDR}(\mathsf{RM}_2(n,d)) = 2^{-d}$ which for $d \geqslant 2$ is much better than the Johnson bound. The GKZ algorithm is a generalization of the Goldreich-Levin algorithm: we assume that we have the correct value of the polynomial given as *advice* on a small random subspace A. This advice allows us to self-correct the values at randomly chosen shifts of A, using a unique decoding algorithm. As pointed out in GKZ, this relies crucially on the coincidence that the ratio of minimum distance to unique decoding radius equals the field size (which is 2), and does not seem to extend to other fields (see Appendix C). They propose the following conjecture:

Conjecture 1. *[GKZ08] For any constants q, d, $\mathsf{LDR}(\mathsf{RM}_q(n,d)) = \delta_q(d)$.*

It is easy to show that $\mathsf{LDR}(\mathsf{RM}_q(n,d)) \leqslant \delta_q(d)$, the crux of the conjecture is the matching lower bound. GKZ show that once we bound $\ell(\mathsf{RM}_q(n,d), \eta)$, (a suitable modification of) the [STV01] algorithm can be used to recover the list of polynomials within radius η. Thus the the algorithmic problem reduces to the combinatorial problem of bounding the list-size. GKZ showed that $\mathsf{LDR}(\mathsf{RM}_q(n,d)) \geqslant \frac{1}{2}\delta_q(d-1)$; by Equation 4 this establishes the

conjecture whenever $d \equiv 0 \bmod q-1$. This bound beats the Johnson bound for d sufficiently large. However when $d = 2$, Conjecture 1 states that agreement exceeding $2/q$ guarantees a small list, the Johnson bound guarantees a small list for agreement $\Omega(1/\sqrt{q})$ whereas the GKZ bound requires agreement exceeding $1/2$. Indeed, we believe that the hard(est) case of Conjecture 1 is when d is small, this precisely is where the gap between $\delta_q(d)$ and known bounds is largest.

2 Our Results

We present a Fourier-analytic approach to Reed-Muller decoding, using a reduction to decoding low-dimensional polynomials. A k-dimensional function is one that can be expressed as a k-junta (a function of at most k variables) under a suitable change of basis for $\mathbb{F}_q^n$.

Definition 1. *The dimension of $F : \mathbb{F}_q^n \to \mathbb{F}_q$ denoted $\dim(F)$ is the smallest k for which there exist linear functions $\alpha_1, \ldots, \alpha_k : \mathbb{F}_q^n \to \mathbb{F}_q$ such that F can be expressed as a function of $\alpha_1, \ldots, \alpha_k$.*

Our approach for bounding the list-size consists of two steps:

1. Bound the number of low-dimensional codewords that are close to any received word. We do this by designing a new Fourier-based algorithm for list-decoding low-dimensional polynomials. This algorithm and its analysis are the principal contributions of this work.

2. Show that the low-weight codewords in the Reed-Muller codes stem from low-dimensional codewords, and without these codewords the minimum distance improves to $\delta_q^h(d) > \delta_q(d)$. Invoking the deletion lemma [GKZ08, GGR09], this allows us to apply the Johnson bound for distance $\delta_q^h(d)$ to bound the number of high-dimensional polynomials in the list.

Our algorithm for low-dimensional polynomials suffices to show that the number of low-dimensional polynomials that lie within radius $\delta_q(d)$ is bounded independent of n. Formally, let $\mathsf{RM}_q^k(n,d)$ be the subcode of $\mathsf{RM}_q(n,d)$ consisting of all polynomials of dimension at most k (where k is constant).

Theorem 2.1. *For all q, k and d it holds that $\mathsf{LDR}(\mathsf{RM}_q^k(n,d)) = \delta_q(d)$.*

In the case of quadratic forms, our notion of dimension coincides with the classical notion of the rank of a quadratic form. It is well known that as the rank of a quadratic form increases, the distribution of its values approaches the uniform distribution over $\mathbb{F}_q$ [LN97]. We use this to prove:

Theorem 2.2. *For all q, it holds that $\mathsf{LDR}(\mathsf{RM}_q(n,2)) = \delta_q(2)$. Further, for any q and $\varepsilon > 0$, we have $\ell(\mathsf{RM}_q(n,2), \delta_q(2) - \varepsilon) = \mathrm{poly}(q, \varepsilon^{-1})$.*

This proves the GKZ conjecture for $d = 2$. In fact their conjecture was only for constant q, whereas our bound is reasonable even for $q = \mathrm{poly}(n)$. Thus there is an efficient algorithm to recover all quadratic polynomials that have agreement $2/q + \varepsilon$. This improves on both the Johnson bound, which requires agreement $(1 + \sqrt{q-1})/q > 1/\sqrt{q}$ and the GKZ bound which requires $(q+1)/2q > 1/2$. Concretely, for $q = 256$, Theorem 2.2 guarantees constant list-size for agreement exceeding $1/128$, whereas Johnson and GKZ require agreement more than $1/16$ and $1/2$ respectively.

For cubic forms and higher, the effectiveness of the method depends on how much the distance improves in Step (2) by *deleting* all low-dimensional polynomials. To formalize this, we define $\delta_q^h(d)$ which is the smallest weight at which codewords of unbounded dimension appear. Let

$$\delta_q^k(d) = \min\{\mathsf{wt}(P) : P \ s.t. \ \deg(P) \leqslant d, \quad \dim(P) = k\}; \quad \delta_q^h(d) = \liminf_{k\to\infty} \delta_q^k(d). \tag{2}$$

While it is *a priori* unclear if $\delta_q^h(d) > \delta_q(d)$, we conjecture that it is in fact substantially larger.

Conjecture 2. *For all d and q it holds that $\delta_q^h(d) \geqslant \delta_q(d-1)$.*

Conjecture 2 fits within the framework of the structure versus randomness dichotomy [Tao07]. For $d < q$, a multivariate degree d polynomial over $\mathbb{F}_q$ vanishes with probability at most d/q, by the Schwartz-Zippel lemma. Though this bound is tight, a random degree d polynomial is very likely to vanish with probability roughly $1/q$, and polynomials that vanish with much higher probability must have special structure. Indeed, Green and Tao [GT07] (see also [KL08]) prove that if a degree d polynomial P vanishes with probability exceeding $1/q$, then it can be expressed as a function of a constant number of polynomials of degree $d-1$ polynomials. Similarly, Conjecture 2 asserts that a degree d polynomial that vanishes with probability exceeding $1 - \delta_q(d-1)$ is a function of constantly many degree 1 polynomials.

Conjecture 2 is easy to verify for quadratic forms over any field. In the case of $\mathbb{F}_2$, it is implied by classical results of Kasami and Tokura [KT70]. This allows us to give an alternate proof of the GKZ result that $\mathsf{LDR}(\mathsf{RM}_2(n,d)) = 2^{-d}$ using the following theorem:

Theorem 2.3. *For all d and q it holds that $\mathsf{LDR}(\mathsf{RM}_q(n,d)) \geqslant \min(\mathsf{J_q}(\delta_q^h(d)), \delta_q(d))$.*

If Conjecture 2 holds, then Theorem 2.3 gives

$$\mathsf{LDR}(\mathsf{RM}(n,d)) \geqslant \min(\mathsf{J_q}(\delta_q(d-1)), \delta_q(d))$$

which improves on the bound of $\max(\frac{1}{2}\delta_q(d-1), \mathsf{J_q}(\delta_q(d)))$ from GKZ for all d and q where their bound is less than $\delta_q(d)$. However, it falls short of proving Conjecture 1 for all d, q. Nevertheless we feel that Conjecture 2 is natural and merits study in its own right; it captures the intuition that restricting to high-dimensional polynomial improves the distance of Reed-Muller codes. In section 5.3 we present some bounds on $\delta_q^h(d)$, and discuss the relation between Conjectures 1 and 2.

2.1 Our Techniques

All previous work on Reed-Muller decoding [GRS00, AS03, STV01, GKZ08] relies on the notion of a self-corrector. Starting the correct values at some point(s) as advice, the algorithm self-corrects the values of the polynomial along some low-dimensional subspace. Our work departs entirely from the self-correction paradigm and draws on ideas from Fourier analysis of Boolean functions; notably (a generalization of) the notion of influence of a variable.

Fourier analytic methods are extensively used in learning, typically for concept classes such as halfspaces [KOS02, KKMS05] or decision trees [KM93] whose Fourier spectra show good *concentration*. Reed-Muller decoding is equivalent to (agnostically) learning low-degree polynomials over $\mathbb{F}_q$. It is not at all clear that Fourier analysis ought to be useful even for $d = 2$, since quadratic forms over $\mathbb{F}_2$ are the canonical examples of *bent* functions whose Fourier spectrum is maximally anti-concentrated [MS77]. However, the deletion lemma allows us to focus on low-degree polynomials which are additionally low-dimensional (dimension at most 6 for quadratic forms). The Fourier spectrum of a k-dimensional polynomial P is supported on a k-dimensional subspace $\mathsf{Spec}(P)$. Our key insight is that *within* $\mathsf{Spec}(P)$*, the Fourier mass is anti-concentrated, which makes it possible to identify* $\mathsf{Spec}(P)$ *via Hadamard decoding, even after the adversary has corrupted the codeword.* We outline the main ideas underlying the proof of this statement below:

1. **Finding $\mathsf{Spec}(P)$:** Fix $q = 2$ for simplicity. The Fourier mass of a k-dimensional polynomial P lies entirely on the subspace $\mathsf{Spec}(P)$ of dimension k. It is easy to recover P if we know $\mathsf{Spec}(P)$. Our goal is to show for any received word F where $\Delta(F, P) \leqslant \delta_q(d)$, the *large* Fourier coefficients of F contain a basis for $\mathsf{Spec}(P)$. Equivalently, the large Fourier coefficients α of F that lie in $\mathsf{Spec}(P)$ should not all fall in a low-dimensional subspace $B \subset \mathsf{Spec}(P)$ satisfying an additional equation $b \cdot \alpha = 0$. One can try and prove this using the Fourier expression for ℓ_2 distance, but this approach fails; owing to counterexamples which are real-valued functions.

2. **The Influence of a Direction:** Given a function F, the Fourier mass that lies in the set $S_b = \{\alpha : b \cdot \alpha \neq 0\}$ captures the influence of direction b, which is defined as $\Pr_{x \in \mathbb{F}_2^n}[F(x) \neq F(x + b)]$. This generalizes the notion of the influence of a variable [KKL88]. Influences in low-degree polynomials P show a dichotomy: they are 0 over a subspace $\mathsf{Inv}(P) = \mathsf{Spec}(P)^{\perp}$, and large for all other b. We use this to show that if $\Delta(F, P) \leqslant \delta_q(d)$, and if b is influential in P, then it has noticeable influence on F. Hence, a noticeable fraction of the Fourier mass of F lies in the set S_b. But it falls short of the claim we really wish to prove, which is that there is noticeable Fourier mass lying in $\mathsf{Spec}(P) \cap S_b$, since F (unlike P) need not be low-dimensional.

3. **Folding the Received word:** The crucial step of our analysis is to go from F to a randomized function $\mathbf{F}$, obtained by folding F over the subspace $\mathsf{Inv}(P) = \mathsf{Spec}(P)^{\perp}$. While we defer the formal definition of folding, the following example is illustrative: if P depends only on $X_1, \ldots, X_k$, then so does $\mathbf{F}$; for each setting of $x_1, \ldots, x_k$, $\mathbf{F}(x_1, \ldots, x_k)$ equals $F(x_1, \ldots, x_n)$ where $x_{k+1}, \ldots, x_n$ are set randomly. From the viewpoint of P,

$\mathbf{F}$ is a received word where the noise added at each point is randomized. The crucial observation is that the noise rate stays the same, so $\Delta(\mathbf{F}, P) \leqslant \delta_q(d)$. Hence every influential direction b of P still has influence on $\mathbf{F}$. But since $\mathbf{F}$ is obtained by folding F over $\mathsf{Inv}(P)$, the Fourier spectrum of $\mathbf{F}$ if just that of F projected on to $\mathsf{Spec}(P)$. Thus we conclude that $\mathbf{F}$ (and hence F) has noticeable Fourier mass lying in $\mathsf{Spec}(P) \cap S_b$. Note that folding is just introduced for the sake of analysis, it plays no role in the algorithm.

4. **Fourier analysis over $\mathbb{F}_q$:** Implementing the above scheme over $\mathbb{F}_q$ is fairly challenging, since it is unclear what *the* Fourier expansion of $F : \mathbb{F}_q^n \to \mathbb{F}_q$ should mean. Our main technical innovation is to associate $q-1$ Fourier polynomials with every such F, this allows us to exactly arithmetize Hamming distance over $\mathbb{F}_q$ and handle randomized functions which is crucial in our setting. We believe that this machinery will find other applications. We use it to prove an equivalence between learning parity with worst-case noise and weaker noise models over $\mathbb{F}_q$, extending a result of [FGKP06] for $\mathbb{F}_2$. We present a Fourier-based analysis of linearity testing over arbitrary finite fields $\mathbb{F}_q$, extending the analysis of Hastad and Wigderson for prime fields [BCH+96, HW03].

Organization: We present Fourier-analytic preliminaries in Section 3, the proofs for this section are deferred to Appendix B. The decoding algorithm for low-dimensional polynomials and its analysis are in Section 4. We present reductions to the low-dimensional case in Section 5, together with some discussion of Conjectures 1 and 2. We present applications to Learning and Testing in Appendix A. We discuss the relation between our work and that of [GKZ08, GGR09] in more detail in Appendix C.

3 Low-Dimensional Functions, Folding and Influences

The proofs for all claims in this Section are in Appendix B.

Fourier analysis

Let $p = \mathsf{char}(q)$ and let $q = p^h$. Let ω be a primitive p^{th} root of unity. Given a random variable Z taking values in $\mathbb{F}_q$, we define the quantities $z^c = \mathbb{E}_Z[\omega^{\mathsf{Tr}(cZ)}]$, which we call the (un-normalized) Fourier coefficients of Z. For two such random variables Y, Z, let $\mathsf{SD}(Y, Z)$ denote their statistical distance. The following relation to the Fourier transform is folklore:

Fact 3.1. *For two random variables Y, Z taking values in $\mathbb{F}_q$, we have*

$$\mathsf{SD}(Y, Z) \leqslant \frac{1}{2} \left(\sum_{c \in \mathbb{F}_q^\star} |y^c - z^c|^2 \right)^{\frac{1}{2}}.$$

Let $\mathsf{Tr}(x) = \sum_{i=0}^{h-1} x^{p^i}$ denote the trace map from $\mathbb{F}_q$ to $\mathbb{F}_p$. The set of all linear functions $\mathbb{F}_q \to \mathbb{F}_p$ is given by $\{\mathsf{Tr}(cx)\}_{c\in\mathbb{F}_q}$. The character group $\hat{\mathbb{F}_q}^n$ of $\mathbb{F}_q^n$ of all homomorphisms $\chi : \mathbb{F}_q^n \to \mathbb{C}$ comprises all functions of the form $\chi_\alpha(x) = \omega^{\mathsf{Tr}(\alpha(x))}$ where $\alpha : \mathbb{F}_q^n \to \mathbb{F}_q$ is a linear function. It is easy to show that the functions χ_α form an orthonormal basis for all functions $f : \mathbb{F}_q^n \to \mathbb{C}$ under the inner-product $\langle f, g\rangle = \mathbb{E}_{x\in\mathbb{F}_q^n} f(x)\overline{g(x)}$. Thus every such f has a Fourier expansion given by

$$f(x) = \sum_{\alpha\in\hat{\mathbb{F}_q}^n} \hat{f}(\alpha)\chi_\alpha(x).$$

We also have $\|f\|_2 = \langle f, f\rangle = \sum_\alpha |\hat{f}(\alpha)|^2$. Given a polynomial $F : \mathbb{F}_q^n \to \mathbb{F}_q$, we associate it with $q-1$ Fourier polynomials mapping $\mathbb{F}_q^n \to \mathbb{C}$, one for every $c \in \mathbb{F}_q^\star$, given by

$$f^c(x) := \omega^{\mathsf{Tr}(cF(x))} = \sum_{\alpha\in\hat{\mathbb{F}_q}^n} \hat{f^c}(\alpha)\chi_\alpha(x).$$

Using $q-1$ polynomials lets us exactly arithmetize agreement and Hamming distance, this is crucial in some of our applications in Section A.

Fact 3.2. *Given functions F, G that map $\mathbb{F}_q^n \to \mathbb{F}_q$,*

$$\mathrm{Ag}(F,G) \;=\; \frac{1}{q}(1 + \sum_{c\in\mathbb{F}_q^\star}\langle f^c, g^c\rangle) \;=\; \frac{1}{q}(1 + \sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha} \hat{f^c}(\alpha)\overline{\hat{g^c}(\alpha)}) \tag{3}$$

$$\Delta(F,G) \;=\; \frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\|f^c - g^c\|_2^2 \;=\; \frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha\in\hat{\mathbb{F}_q}^n} |\hat{f^c}(\alpha) - \hat{g^c}(\alpha)|^2 \tag{4}$$

Randomized Functions

We consider randomized functions $\mathbf{F} : \mathbb{F}_q^n \to \mathbb{F}_q$, where each $\mathbf{F}(x)$ is a random variable taking values in $\mathbb{F}_q$. We define the Fourier polynomials associated with $\mathbf{F}$:

Definition 2. *Given a randomized function $\mathbf{F} : \mathbb{F}_q^n \to \mathbb{F}_q$, for each $c \in \mathbb{F}_q^\star$, we define the polynomial $\mathbf{f}^c : \mathbb{F}_q^n \to \mathbb{C}$ by*

$$\mathbf{f}^c(x) = \mathbb{E}_{\mathbf{F}}[\omega^{\mathsf{Tr}(c\mathbf{F}(x))}] = \sum_{\alpha\in\hat{\mathbb{F}_q}^n} \hat{\mathbf{f}^c}(\alpha)\chi_\alpha(x).$$

Note that $\mathbf{f}^c$ is a (deterministic) function from $\mathbb{F}_q^n \to \mathbb{C}$ and the values $\{\mathbf{f}^c(x)\}_{c\in F_q^\star}$ give us the Fourier transform of $\mathbf{F}(x)$. Given two randomized functions $\mathbf{F}, \mathbf{G} : \mathbb{F}_q^n \to \mathbb{F}_q$, we define

$$d(\mathbf{F}, \mathbf{G}) = \mathbb{E}_{x\in\mathbb{F}_q^n}[\mathsf{SD}(\mathbf{F}(x), \mathbf{G}(x))]$$

generalizing the definitions for deterministic functions.

Fact 3.3. *Given randomized functions* $\mathbf{F}, \mathbf{G}$ *that map* $\mathbb{F}_q^n \to \mathbb{F}_q$,

$$d(\mathbf{F},\mathbf{G}) \leqslant \frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star} \mathbb{E}_x[|\mathbf{f}^c(x)-\mathbf{g}^c(x)|^2]\right)^{\frac{1}{2}} = \frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha\in\hat{\mathbb{F}_q}^n} |\hat{\mathbf{f}^c}(\alpha)-\hat{\mathbf{g}^c}(\alpha)|^2\right)^{\frac{1}{2}}. \quad (5)$$

Low-Dimensional Functions

We first define low-dimensional randomized functions.

Definition 3. *A randomized function* $\mathbf{F}:\mathbb{F}_q^n \to \mathbb{F}_q$ *is* k *dimensional if there exist* k *linear forms* $\alpha_1,\ldots,\alpha_k:\mathbb{F}_q^n\to\mathbb{F}_q$ *such that knowing* $\alpha_1(x),\ldots,\alpha_k(x)$ *fixes the distribution of* $\mathbf{F}(x)$.

Hence $\mathbf{F}$ is a (randomized) function of $\alpha_1,\ldots,\alpha_k$, generalizing Definition 1. Facts 3.4 and 3.5 below are proved in [GKS07, GOS+09] for deterministic functions.

Fact 3.4. *For each* $c\in\mathbb{F}_q^\star$, *let* $\mathsf{Supp}(\mathbf{f}^c)\subseteq\hat{\mathbb{F}_q}^n$ *denote the set of non-zero Fourier coefficients of* $\mathbf{f}^c(x)$. *Let* $\mathsf{Spec}(\mathbf{F}) = \mathsf{Span}(\cup_{c\in\mathbb{F}_q^\star}\mathsf{Supp}(\mathbf{f}^c))$. *Then* $\dim(\mathbf{F}) = \dim(\mathsf{Spec}(\mathbf{F}))$.

Alternatively, low-dimensional functions can be defined via invariant subspaces.

Definition 4. *Given* $h\in\mathbb{F}_q^n$, *if* $\mathbf{F}:\mathbb{F}_q^n\to\mathbb{F}_q$ *satisfies*

$$\mathsf{SD}(\mathbf{F}(x+\lambda h),\mathbf{F}(x)) = 0 \quad \forall\, x\in\mathbb{F}_q^n, \lambda\in\mathbb{F}_q$$

we say that $\mathbf{F}$ *is* h-invariant. *We define* $\mathsf{Inv}(\mathbf{F}) = \{h : \mathbf{F} \text{ is } h\text{-invariant}\}$.

$\mathsf{Inv}(\mathbf{F})$ is clearly a subspace of $\mathbb{F}_q^n$, and is in fact dual to $\mathsf{Spec}(\mathbf{F})$.

Fact 3.5. *We have* $\mathsf{Spec}(\mathbf{F}) = \mathsf{Inv}(\mathbf{F})^\perp$. *Hence* $\dim(\mathbf{F}) = \mathrm{codim}(\mathsf{Inv}(\mathbf{F}))$.

Folding

Folding over subspaces was introduced in [FGKP06] (in the $\mathbb{F}_2$ case). Folding maps high-dimensional functions to lower-dimensional randomized functions.

Definition 5. *Let* H *be a subspace of* $\mathbb{F}_q^n$ *and let* $F:\mathbb{F}_q^n\to\mathbb{F}_q$. *Define the randomized function* $\mathbf{F}(x) = F(x+h)$ *where* $h\in H$ *is chosen randomly. We call* $\mathbf{F}$ *the folding of* F *over* H.

Given an oracle for F, we can simulate an oracle for $\mathbf{F}$: on query x, choose a random point $x+h$ in the coset $x+H$ and return $F(x+H)$. Thus $\mathbf{F}$ is invariant on H. In fact, its Fourier spectrum is obtained by projecting the spectrum of F onto $H^\perp$.

Lemma 3.6. *[FGKP06] Let* $\mathbf{F}$ *be the folding of* F *over* H. *For any* $c\in\mathbb{F}_q^\star$, *we have* $\hat{\mathbf{f}^c}(\alpha) = \hat{f^c}(\alpha)$ *if* $\alpha\in H^\perp$ *and* $\hat{\mathbf{f}^c}(\alpha) = 0$ *otherwise.*

The Influence of a Direction

We define the influence of a direction, which is a generalization of the notion of influence of a variable. Given a vector $b \in \mathbb{F}_q^n \setminus \{0^n\}$, we partition $\mathbb{F}_q^n$ into lines along the direction b, which are the equivalence classes for the relation $x \sim y$ if $x - y = \lambda b$ for some $\lambda \in \mathbb{F}_q$. This partition is just $\mathbb{F}_q^n/\{b\}$, and it is isomorphic to $\mathbb{F}_q^{n-1}$.

Definition 6. (Influence of a direction) *Given $b \in \mathbb{F}_q^n$, and a function $F : \mathbb{F}_q^n \to \mathbb{F}_q$ we define*

$$\mathsf{Inf}_b(F) = \Pr_{x \in \mathbb{F}_q^n, \lambda \in \mathbb{F}_q}[F(x) \neq F(x + \lambda b)].$$

One can relate $\mathsf{Inf}_b(F)$ to the Fourier mass lying outside the subspace of $\hat{\mathbb{F}_q}^n$ given by $b \cdot \alpha = 0$.

Fact 3.7. *Given $b \in \mathbb{F}_q^n$, we have*

$$\mathsf{Inf}_b(F) = \frac{1}{q} \sum_c \sum_{\alpha:\ b \cdot \alpha \neq 0} |\hat{f^c}(\alpha)|^2 \tag{6}$$

We extend the notion of influences to randomized functions (generalizing the above notion). To compute the influence of b for a deterministic function, we pick sample two points on a line (in the direction b) and compute their Hamming distance. For randomized function, we sample two points and compute their statistical distance.

Definition 7. *Given a randomized function $\mathbf{F} : \mathbb{F}_q^n \to \mathbb{F}_q$ and $b \in \mathbb{F}_q^n$, we define $\mathsf{Inf}_b(\mathbf{F})$ as*

$$\mathsf{Inf}_b(\mathbf{F}) = \mathbb{E}_{x \in \mathbb{F}_q^n, \lambda \in \mathbb{F}_q}[\mathsf{SD}(\mathbf{F}(x), \mathbf{F}(x + \lambda b)].$$

One can again bound the influence in terms of the Fourier mass that lies outside the subspace $b \cdot \alpha = 0$.

Lemma 3.8. *Given $b \in \mathbb{F}_q^n$, we have*

$$\mathsf{Inf}_b(\mathbf{F}) \leqslant \frac{1}{\sqrt{2}} \left(\sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha:\ b \cdot \alpha \neq 0} |\hat{\mathbf{f}}^c(\alpha)|^2 \right)^{\frac{1}{2}}.$$

4 List-decoding low-dimensional polynomials

In this section, we prove Theorem 2.1. Assume that we have an efficient procedure Had for finding large Fourier coefficients over $\mathbb{F}_q^n$. Given oracle access to $f : \mathbb{F}_q^n \to \mathbb{C}$ and a parameter μ, $\mathsf{Had}(f, \mu)$ returns all $\alpha \in \hat{\mathbb{F}_q}^n$ so that $|\hat{f}(\alpha)|^2 \geqslant \mu$. The list-size is bounded by $\|f\|_2^2/\mu$. Such algorithms are given by [Man95, GGI+02, AGS03]. Theorem 2.1 is proved by arguing

that the polynomial P will be in the list of polynomials that is returned by the following algorithm.

Algorithm 1. List-Decoding low-dimensional polynomials
Input: d, k, ε, `oracle for` $F : \mathbb{F}_q^n \to \mathbb{F}_q$.
Output: `All` $P : \mathbb{F}_q^n \to \mathbb{F}_q$ `s.t.` $\deg(P) \leqslant d, \dim(P) \leqslant k$ `and` $\Delta(P, F) \leqslant \delta_q(d)(1-\varepsilon)$.

1. `Set` $\mu = \varepsilon^4 \delta_q(d)^2/(8q^{k+1})$.
2. `Run` $\mathsf{Had}(f^c, \mu)$ `for all` $c \in \mathbb{F}_q^\star$.
3. `Let` $\mathcal{L}$ `be the list of all linear functions` α `returned.`
4. `Pick` $\alpha_1, \ldots, \alpha_k$ `from` $\mathcal{L}$.
5. `Return all` $P(\alpha_1, \ldots, \alpha_k)$ `s.t.` $\deg(P) \leqslant d$ `and` $\Delta(P, F) \leqslant \delta_q(d)(1-\varepsilon)$.

4.1 Correctness of the Algorithm

Fix a polynomial P with $\deg(P) \leqslant d$, $\dim(P) \leqslant k$ and $\Delta(F, P) = \eta \leqslant \delta_q(d)(1-\varepsilon)$. Our goal is to prove that the list $\mathcal{L}$ contains a basis for $\mathsf{Spec}(P)$, which implies that P one of the polynomials returned by our algorithm. For the analysis, we work with the randomized function $\mathbf{F}$ obtained by folding F over $\mathsf{Inv}(P)$. Folding over $\mathsf{Inv}(P)$ projects the Fourier spectrum of F on to $\mathsf{Spec}(P)$, which is a *small* subspace with only q^k vectors in it. Our main lemma states that all directions that were influential in P continue to have some influence even in $\mathbf{F}$.

Lemma 4.1. (Main) *For the function* $\mathbf{F}$ *defined above and any* $b \notin \mathsf{Inv}(P)$,

$$\mathsf{Inf}_b(\mathbf{F}) \geqslant \frac{\varepsilon^2}{4}\delta_q(d).$$

Proof. Consider the vector space $V = \mathbb{F}_q^n/\mathsf{Inv}(P) \sim \mathbb{F}_q^k$. We can view P as a function $P : V \to \mathbb{F}_q$. Similarly, we can view $\mathbf{F}$ as a randomized function $\mathbf{F} : V \to \mathbb{F}_q$, obtained by adding random noise of rate η to P. Formally, for each $y \in V$, define the noise rate

$$\eta(y) = \Pr_{\mathbf{F}}[\mathbf{F}(y) \neq P(y)] = \Pr_{x \in y + \mathsf{Inv}(P)}[F(x) \neq P(y)]$$

and note that

$$\mathbb{E}_{y \in V}\, \eta(y) = \Pr_{y \in V, x \in y + \mathsf{Inv}(P)}[F(x) \neq P(y)] = \Pr_{x \in \mathbb{F}_q^n}[F(x) \neq P(x)] = \eta.$$

Our goal is to show that any $b \notin \mathsf{Inv}(P)$ has non-negligible influence on $\mathbf{F}$. Recall that for a randomized function $\mathbf{F} : \mathbb{F}_q^n \to \mathbb{F}_q$ and $b \in \mathbb{F}_q^n$, we defined $\mathsf{Inf}_b(\mathbf{F})$ as

$$\mathsf{Inf}_b(\mathbf{F}) = \mathbb{E}_{x \in \mathbb{F}_q^n, \lambda \in \mathbb{F}_q}[\mathsf{SD}(\mathbf{F}(x), \mathbf{F}(x + \lambda b))].$$

Since $\mathbf{F}$ is invariant on $\mathsf{Inv}(P)$, this is equivalent to

$$\mathsf{Inf}_b(\mathbf{F}) = \mathbb{E}_{y \in V, \lambda \in \mathbb{F}_q}[\mathsf{SD}(\mathbf{F}(y), \mathbf{F}(y + \lambda b))]. \tag{7}$$

Consider $V/\{b\}$, the partition of V into lines along b. We can rewrite Equation 7 as

$$\mathsf{Inf}_b(\mathbf{F}) = \mathop{\mathbb{E}_{L \in V/\{b\}}}_{x,y \in L}[\mathsf{SD}(\mathbf{F}(x), \mathbf{F}(y))]. \tag{8}$$

Let us fix a basis containing the vector b for V: call it $\{a_1, \ldots, a_{k-1}, b\}$. Every vector $y \in V$ can be written in this basis as $y = \sum_{i=1}^{k-1} a_i y_i + b y_k$. The polynomial P can we written as $P(y_1, \ldots, y_k)$ of degree d. Assume that y_k occurs with degree $d_2 \leqslant q-1$ (this might depend on the choice of basis). So we can write

$$P(y_1, \ldots, y_k) = Q(y_1, \ldots, y_{k-1}) y_k^{d_2} + \sum_{e < d_2} Q_e(y_1, \ldots, y_{k-1}) y_k^e.$$

for some Q such that $\deg(Q) = d_1 \leqslant d - d_2$. Fixing values for $(y_1, \ldots, y_{k-1})$ specifies a line in $V/\{b\}$, while fixing y_k specifies a point on that line. Thus we have

$$\mathsf{Inf}_b(\mathbf{F}) = \mathop{\mathbb{E}}_{y_1, \ldots, y_{k-1}, y_k, y_k'}[\mathsf{SD}(\mathbf{F}(y_1, \ldots, y_{k-1}, y_k), \mathbf{F}(y_1, \ldots, y_{k-1}, y_k'))]. \tag{9}$$

We say that a line $\ell = (y_1, \ldots, y_{k-1}) \in V/\{b\}$ is good if $Q(y_1, \ldots, y_{k-1}) \neq 0$. Since $\deg(Q) \leqslant d_1$, $\Pr_\ell[\ell \text{ is good}] \geqslant \delta_q(d_1)$. Conditioning on the event that ℓ is good, $P|_\ell$ is a univariate polynomial of degree d_2. Hence, it takes on any particular value in $\mathbb{F}_q$ no more than d_2 times. In contrast, if ℓ is bad, then $P|_\ell$ is constant.

Define the noise rate $\eta(\ell)$ for a line as $\eta(\ell) = \mathbb{E}_{y \in \ell}[\eta(y)]$. We have $\mathbb{E}_{\ell \in V/\{b\}}[\eta(\ell)] = \eta$. We say that a good line is *quiet* if the noise rate along the line is low:

$$\eta(\ell) \leqslant \left(1 - \frac{d_2}{q}\right)\left(1 - \frac{\varepsilon}{2}\right).$$

We claim that at least $\varepsilon/2$ fraction of good lines are quiet; else we have

$$\mathbb{E}_\ell[\eta(\ell)] \geqslant \delta_q(d_1)\left(1 - \frac{\varepsilon}{2}\right)\left(1 - \frac{d_2}{q}\right)\left(1 - \frac{\varepsilon}{2}\right) \;>\; \delta_q(d_1)\left(1 - \frac{d_2}{q}\right)(1 - \varepsilon) \;\geqslant\; \delta_q(d)(1 - \varepsilon).$$

where the last inequality follows from the following property of $\delta_q(d)$:

$$\delta_q(d) \leqslant \delta_q(d_1)\left(1 - \frac{d_2}{q}\right) \quad \text{for all} \;\; d_1, d_2 \;\; s.t.\; d_1 + d_2 \leqslant d, 0 \leqslant d_2 \leqslant q-1.$$

This is easy to verify from Equation 1. Now fix a quiet line ℓ. We have a polynomial $P|_\ell : \ell \to F_q$ of degree $d_2 \leqslant q-1$ and a randomized received word $\mathbf{F}|_\ell$ such that

$$d(P|_\ell, \mathbf{F}|_\ell) = \mathbb{E}_{x \in \ell}[\mathsf{SD}(P(x), \mathbf{F}(x))] = \mathbb{E}_{x \in \ell}[\eta(x)] \leqslant \delta_q(d_2) - \varepsilon'$$

where $\delta_q(d_2) = 1 - \frac{d_2}{q}$ and $\varepsilon' = \frac{1}{2}\delta_q(d_2)\varepsilon$. The final piece of the argument is to show that for every quiet line, $\mathsf{Inf}_b(\mathbf{F})$ is high, which is essentially a claim about univariate polynomials.

Claim 4.2. *For a quiet line ℓ, we have $\mathbb{E}_{x,y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))] \geqslant \varepsilon'$.*

Let us defer the proof of this claim and finish the proof of Lemma 4.1. We have argued that

$$\Pr_{\ell\in V/\{b\}}[\ell \text{ is quiet}] \geqslant \frac{1}{2}\varepsilon\delta_q(d_1) \tag{10}$$

Conditioned on the event that ℓ is quiet, we have proved that

$$\mathop{\mathbb{E}}_{x,y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))] \geqslant \frac{1}{2}\varepsilon\delta_q(d_2) \tag{11}$$

Plugging this into Equation 8 gives

$$\mathsf{Inf}_b(\mathbf{F}) = \mathop{\mathbb{E}_{\ell\in V/\{b\}}}_{x,y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))] \geqslant \frac{\varepsilon^2}{4}\delta_q(d_1)\delta_q(d_2) \geqslant \frac{\varepsilon^2}{4}\delta_q(d) \tag{12}$$

which completes the proof of Lemma 4.1. □

Proof of Claim 4.2. For the purposes of this claim, we use P and $\mathbf{F}$ to denote $P|_\ell$ and $\mathbf{F}|_\ell$ respectively. Similarly $d(\cdot,\cdot)$ will denote distance between randomized functions on the line ℓ.

For every distribution $\mathcal{D}$ on $\mathbb{F}_q$, we can define the (constant) randomized function $\mathcal{D}^q : \ell \to \mathbb{F}_q$ where $\mathcal{D}^q(x) = \mathcal{D}$ for every $x \in \ell$. We claim that $d(P,\mathcal{D}^q) \geqslant \delta_q(d_2)$ for every such distribution $\mathcal{D}$. In the case where $\mathcal{D} = \mathcal{D}_y$ is concentrated at a single point $y \in \mathbb{F}_q$, this holds since $P(x)$ is a univariate polynomial with $\deg(P) = d_2$ and so $\Pr_x[P(x) = y] \leqslant d_2/q$. More generally, we have

$$\begin{aligned} d(P,\mathcal{D}^q) = \mathbb{E}_x[\mathsf{SD}(P(x),\mathcal{D})] \;&=\; \sum_{x\in F_q}\frac{1}{q}(1-\mathcal{D}(P(x)) \;=\; \sum_{y\in F_q}\Pr[P(x)=y](1-\mathcal{D}(y)) \\ &= \sum_{y\in F_q}\Pr[P(x)=y] - \sum_{y\in\mathbb{F}_q}\Pr[P(x)=y]\mathcal{D}(y) \;\geqslant\; 1-\frac{d_2}{q} \end{aligned}$$

where the last inequality uses $\Pr_x[P(x) = y] \leqslant d_2/q$ as $\deg(P) \leqslant d_2$. By the triangle inequality

$$d(\mathbf{F},\mathcal{D}^q) \geqslant d(P,\mathcal{D}^q) - d(\mathbf{F},P) \geqslant \delta_q(d_2) - (\delta_q(d_2)-\varepsilon') = \varepsilon'.$$

We compute $\mathbb{E}_{x,y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))]$ by first sampling $x \in \ell$ and then computing the distance between $\mathbf{F}$ and the distribution $\mathcal{D}^q$ where $\mathcal{D} = \mathbf{F}(x)$.

$$\mathbb{E}_{x,y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))] = \mathbb{E}_{x\in\ell}[\mathbb{E}_{y\in\ell}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(y))]] \;=\; \mathbb{E}_{x\in\ell}[d(\mathbf{F}(x)^q,\mathbf{F})] \;\geqslant \varepsilon'.$$

This finishes the proof of Claim 4.2. □

With the Main lemma in hand, Theorem 2.1 follows easily by the following claim:

Lemma 4.3. *The list $\mathcal{L}$ returned contains a basis for* $\mathsf{Spec}(P)$.

Proof. Assume that the Fourier coefficients in $\mathcal{L} \cap \mathsf{Spec}(P)$ do not span all of $\mathsf{Spec}(P)$, rather they span a subspace B of it that satisfies the additional constraint $b \cdot \alpha = 0$ for $b \in\notin \mathsf{Inv}(P)$. We have

$$\frac{1}{\sqrt{2}} \left(\sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha : b \cdot \alpha \neq 0} |\hat{\mathbf{f}}^c(\alpha)|^2 \right)^{\frac{1}{2}} \geqslant \mathsf{Inf}_b(\mathbf{F}) \geqslant \frac{1}{4} \varepsilon^2 \delta_q(d) \tag{13}$$

where the first inequality is from Lemma 3.8 and the second from Lemma 4.1. Applying Lemma 3.6 to the function $\mathbf{F}$ which is F folded over $\mathsf{Inv}(P)$, we get $\hat{\mathbf{f}}^c(\alpha) = \hat{f}^c(\alpha)$ for $\alpha \in \mathsf{Spec}(P)$ and $\hat{\mathbf{f}}^c(\alpha) = 0$ otherwise. Combining these equations, we get

$$\sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha \in \mathsf{Spec}(P) \setminus B} |\hat{f}^c(\alpha)|^2 \geqslant \frac{1}{8} \varepsilon^4 \delta_q(d)^2$$

Since we sum over $(q^k - q^{k-1})(q-1) < q^{k+1}$ Fourier coefficients on the LHS, at least one of them is as large as the average. Thus, there exist $c \in \mathbb{F}_q^\star$ and $\alpha \in \mathsf{Spec}(P) \setminus B$ so that

$$|\hat{f}^c(\alpha)|^2 > \frac{1}{8} \frac{\varepsilon^4 \delta_q(d)^2}{q^{k+1}}.$$

This coefficient α must belong to the list $\mathcal{L}$, which contradicts the assumption that $\mathcal{L} \cap \mathsf{Spec}(P)$ is contained within B. □

A simple calculation which we omit gives the following bound on the list-size for $\mathsf{RM}_q^k(n,d)$ (we have not attempted to optimize this bound). There exists a constant $c > 0$ such that

$$\ell(\mathsf{RM}_q^k(n,d), \delta_q(d)(1-\varepsilon)) \leqslant \frac{c^k q^{k^d + k^2 + 2k}}{\varepsilon^{4k} \delta_q(d)^{2k}}. \tag{14}$$

The running time of Algorithm 1 is polynomial in n^d, q and the list-size.

5 Reductions to the low-dimensional case.

We use the [GGR09] version of the deletion lemma from [GKZ08].

Lemma 5.1. *[GKZ08, GGR09]* (Deletion Lemma) *Let $\mathcal{C} \subset \mathbb{F}_q^n$ be a linear code over $\mathbb{F}_q$. Let $\mathcal{C}' \subseteq \mathcal{C}$ be a (possibly non-linear) subset of codewords so that $c' \in \mathcal{C}'$ iff $-c' \in \mathcal{C}'$, and every codeword $c \in \mathcal{C} \setminus \mathcal{C}'$ has $\mathsf{wt}(c) \geqslant \delta^h$. Let $\eta = \mathsf{J}_\mathsf{q}(\delta^h) - \gamma$ for $\gamma > 0$. Then $\ell(\mathcal{C}, \eta) \leqslant \gamma^{-2} \ell(\mathcal{C}', \eta)$.*

5.1 Quadratic Forms

For quadratic forms $Q : \mathbb{F}_q^n \to \mathbb{F}_q$, $\dim(Q)$ coincides with the well-studied notion of the rank of a quadratic form. Theorems 6.26, 6.27 and 6.32 from Chapter 6 of [LN97] give the following bound:

Lemma 5.2. *Let $Q : \mathbb{F}_q^n \to \mathbb{F}_q$ be a quadratic form such that $\dim(P) = k$. Then*

$$\mathsf{wt}(Q) \geqslant 1 - \frac{1}{q} - \frac{1}{q^{k/2}}.$$

We use this to complete the proof of Theorem 2.2.

Proof of of Theorem 2.2. By Lemma 5.2, if $\dim(Q) \geqslant 6$, then we have

$$\mathsf{wt}(Q) \geqslant 1 - \frac{1}{q} - \frac{1}{q^3}; \quad \mathsf{J_q}\left(1 - \frac{1}{q} - \frac{1}{q^3}\right) > 1 - \frac{2}{q}.$$

Hence we can apply Lemma 5.1 with $\mathcal{C}' = \mathsf{RM}_q^6(n, 2)$ to conclude that there exists C so that

$$\ell(\mathsf{RM}_q(n,2), \delta_q(2) - \varepsilon) \leqslant \frac{1}{\varepsilon^2}\ell(\mathsf{RM}_q^6(n,2), \delta_q(2) - \varepsilon) \leqslant C\frac{q^{84}}{\varepsilon^{26}}.$$

□

5.2 The $\mathbb{F}_2$ case revisited

Using our techniques, we can give an alternate proof of the GKZ result that $\mathsf{LDR}(\mathsf{RM}_2(n,d) = 2^{-d}$. A classical result of Kasami and Tokura allows us to bound the rank of any codeword of $\mathsf{RM}_2(n,d)$ which has dimension less than $2\delta_2(d)$.

Lemma 5.3. *[KT70] Let $d \geqslant 2$. Let $P : \mathbb{F}_2^n \to \mathbb{F}_2$ with $\deg(P) \leqslant d$ and $\mathsf{wt}(P) < 2\delta_2(d)$. Then P is of one of the following two types:*

1. $P(\alpha_1, \ldots, \alpha_{d+t}) = \alpha_1 \cdots \alpha_{d-t}(\alpha_{d-t+1} \cdots \alpha_d + \alpha_{d+1} \cdots \alpha_{d+t}) \quad 3 \leqslant t < d.$
2. $P(\alpha_1, \ldots, \alpha_{d+2t-2}) = \alpha_1 \cdots \alpha_{d-2}(\alpha_{d-1}\alpha_d + \alpha_{d+1}\alpha_{d+2} + \cdots + \alpha_{d+2t-3}\alpha_{d+2t-2}).$

where the α_is are independent linear forms.

Strictly speaking, the α_is are affine rather linear, but we can safely ignore this issue.

Corollary 5.4. *Let $P : \mathbb{F}_2^n \to \mathbb{F}_2$ be a degree d polynomial with $\dim(P) = k \geqslant 2d$. Then $\mathsf{wt}(P) \geqslant 2\delta_2(d) - 2^{-(k+d)/2}$.*

Proof. Assume that $\mathsf{wt}(P) < 2\delta_2(d)$, else the claim is trivial. Now applying Lemma 5.3, P must be of type (2), since polynomials of type (1) have dimension less than $2d$. A simple calculation shows that for polynomials of type (2), if $\dim(P) = k$, then $\mathsf{wt}(P) \geqslant 2\delta_2(d) - 2^{-(k+d)/2}$. □

We can now reprove the main result from [GKZ08]. Our dependence is polynomial in ε^{-r}, though the exact bound is inferior to GKZ, who also showed a lower bound of $\varepsilon^{-\Omega(r)}$.

Theorem 5.5. *[GKZ08] For all $d \geqslant 1$, it holds that* $\mathsf{LDR}(\mathsf{RM}_2(n,d)) = 2^{-d}$.

Proof. Pick $k = 3d$. Take $\mathcal{C}' = \mathsf{RM}_2^k(n,d)$. By Equation 14, we have

$$\ell(\mathcal{C}', \delta_2(d) - \varepsilon) \leqslant C\varepsilon^{-12d}$$

for some constant $C = C(d)$ that depends on d. By Corollary 5.4, if $\dim(P) \geqslant 3d$,

$$\mathsf{wt}(P) \geqslant 2 \cdot 2^{-d} - 2^{-2d}; \quad \mathsf{J}_2(2 \cdot 2^{-d} - 2^{-2d}) > 2^{-d}.$$

Hence applying Lemma 5.1, we get

$$\ell(\mathsf{RM}_2(n,d), \delta_2(d) - \varepsilon) \leqslant C\varepsilon^{-(12d+2)}$$

which completes the proof. □

5.3 The Case of arbitrary d and q.

In Equation 2 defining $\delta_q^k(d)$, we minimize over the infinite set of all degree d polynomials P with $\dim(P) = k$, the number of variables n could be arbitrary. But since $\dim(P) = k$, we may assume that P is on exactly k variables. Thus we are in effect minimizing over the finite set of $P : \mathbb{F}_q^k \to \mathbb{F}_q$ s.t. $\deg(P) = d$ and $\dim(P) = k$, so $\delta_q^k(d)$ is well-defined.

Combining the deletion lemma with Theorem 2.1 lets us complete the proof of Theorem 2.3.

Proof of Theorem 2.3. Let $\eta = \min(\delta_q(d), \mathsf{J_q}(\delta_q^h(d))) - \varepsilon$. Our goal is to show that for any $\varepsilon > 0$, $\ell(\mathsf{RM}_q(n,d), \eta)$ which is the list-size at radius η can be bounded independent of n.

Since $\eta \leqslant \delta_q(d) - \varepsilon$, by Theorem 2.1

$$\ell(\mathsf{RM}_q^k(n,d), \eta) \leqslant \ell(d,k,q,\varepsilon) = \frac{c^k q^{k^d + k^2 + 2k}}{\varepsilon^{4k} \delta_q(d)^{2k}}.$$

Choose k large enough so that

1. $\delta_q^{k'}(d) \geqslant \delta_q^k(d)$ for all $k' \geqslant k$.
2. $\mathsf{J_q}(\delta_q^k(d)) \geqslant \mathsf{J_q}(\delta_q^h(d)) - \varepsilon/2$.

The second condition implies that

$$\eta \leqslant \mathsf{J_q}(\delta_q^h(d)) - \varepsilon \leqslant \mathsf{J_q}(\delta_q^k(d)) - \varepsilon/2.$$

Every codeword outside of $\mathsf{RM}_q^k(n,d)$ has $\dim(P) \geqslant k$, and hence $\mathsf{wt}(P) \geqslant \delta_q^k(d)$. Thus we can invoke Lemma 5.1 with $\mathcal{C}' = \mathsf{RM}_q^k(n,d)$ to conclude that

$$\ell(\mathsf{RM}_q(n,d),\eta) \leqslant \frac{4}{\varepsilon^2}\ell(d,k,q,\varepsilon) = \ell'(d,k,q,\varepsilon).$$

This shows that the list-size at radius $\min(\delta_q(d), \mathsf{J_q}(\delta_q^h(d))) - \varepsilon$ is bounded independent of n for every $\varepsilon > 0$, which proves the claim. □

If Conjecture 2 holds true, then Theorem 2.3 implies that

$$\mathsf{LDR}(\mathsf{RM}_q(n,d)) \geqslant \min(\mathsf{J_q}(\delta_q(d-1)), \delta_q(d)).$$

This would improve on both the GKZ and the Johnson bound by the following claim:

Claim 5.6. *For all d and q, it holds that*

$$\min(\mathsf{J_q}(\delta_q(d-1)), \delta_q(d)) \geqslant \max\left(\mathsf{J_q}(\delta_q(d)), \frac{1}{2}\delta_q(d-1)\right).$$

The inequality is strict except when $d = 1$ and $d \equiv 0 \bmod q-1$, and in both those cases the RHS equals $\delta_q(d)$.

Proof. Note that for all $\eta \in [0, 1-1/q]$, we have

$$\eta/2 \leqslant \mathsf{J_q}(\eta) \leqslant \eta$$

with $\mathsf{J_q}(\eta) = \eta$ iff $\eta = 1 - \frac{1}{q}$ and $\mathsf{J_q}(\eta) = \eta/2$ iff $\eta = 0$. Also $J(\eta)$ increase monotonically with η.

Further, if $d = a(q-1) + b$ for $1 \leqslant b \leqslant q-1$,

$$\begin{aligned} \delta_q(d-1) &= \delta_q(d)\left(1 + \frac{1}{q-b}\right) \\ \Rightarrow \quad \frac{q}{q-1}\delta_q(d) \;&\leqslant\; \delta_q(d-1) \;\leqslant\; 2\delta_q(d). \end{aligned}$$

The former is tight when $d \equiv 1 \bmod (q-1)$, the latter when $d \equiv q-1 \bmod (q-1)$.

We now prove the above claim. Firstly, note that from the above inequalities, we have

$$\mathsf{J_q}(\delta_q(d-1)) > \frac{1}{2}\delta_q(d-1) \;\text{ and }\; \mathsf{J_q}(\delta_q(d-1)) > \mathsf{J_q}(\delta_q).$$

Secondly, we also have

$$\delta_q(d) \geqslant \frac{1}{2}\delta_q(d-1) \;\text{ and }\; \delta_q(d) \geqslant \mathsf{J_q}(\delta_q(d)).$$

The first inequality is strict, except when $d \equiv q-1 \bmod (q-1)$. In this case, the GKZ bound is already tight. Similarly, the second inequality is strict except when $\delta_q(d) = 1 - \frac{1}{q}$, which holds when $d = 1$ or Hadamard codes, in which case the Johnson bound is tight. □

Note that conjecture 2 holds if $d = 2$ and q is arbitrary, or if $q = 2$ and d is arbitrary. The quadratic case follows from Lemma 5.2 while the $\mathbb{F}_2$ case follows from Corollary 5.4. Indeed, in both cases $\delta_q(d-1) = \delta_q^h(d) = \lim_{k\to\infty} \delta_q^k(d)$.

Assuming the truth of Conjecture 2, one could ask how close it gets us to proving Conjecture 1. This depends on how $\mathsf{J_q}(\delta_q(d-1))$ compares to $\delta_q(d)$. When $d = 2$, we have $\mathsf{J_q}(\delta_q(1)) > \delta_q(2)$. But for $d = 3$ and larger, assuming q is a large constant, we have

$$\mathsf{J_q}(\delta_q(d-1)) \approx 1 - \sqrt{\frac{d-1}{q}} < 1 - \frac{d}{q} = \delta_q(d).$$

Hence the minimum of the two quantities is $\mathsf{J_q}(\delta_q(d-1))$. As d gets larger, $\delta_q(d-1)$ decreases towards 0, hence $\mathsf{J_q}(\delta_q(d-1)) \approx \frac{1}{2}\delta_q(d-1)$. So our approach approaches the GKZ bound.

5.4 An upper bound on $\delta_q^h(d)$.

We now present an upper bound on $\delta_q^h(d)$.

Lemma 5.7. *For all $d \geqslant 3$ and q, we have*

$$\delta_q^h(d) \leqslant \delta_q(d-2)\left(1 - \frac{1}{q}\right). \tag{15}$$

Proof. Let k be odd. Define $Q : \mathbb{F}_q^k \to \mathbb{F}_q$ to be a dimension k quadratic form which is completely unbiased. For instance, we can take

$$Q(Y_1, \ldots, Y_k) = \sum_{i \leqslant \lfloor k/2 \rfloor} Y_{2i}Y_{2i+1} + Y_k.$$

Take $R(Z_1, \ldots, Z_\ell)$ to be a minimum weight codeword on degree $d-2$, this requires $\ell = \lceil \frac{d-2}{q} \rceil$. Now set $P = QR$, so that $P : \mathbb{F}_q^{k+\ell} \to \mathbb{F}_q$, $\deg(P) = d$, $\dim(P) = k + \ell$ and

$$\mathsf{wt}(P) = \delta_q(d-2)\left(1 - \frac{1}{q}\right).$$

Thus by taking k sufficiently large, we can construct polynomials of the desired weight whose dimension grows unbounded. □

Let us compare this upper bound with the lower bound of $\delta_q(d-1)$ claimed in our conjecture. Let $d - 1 = a(q-1) + b$, where $1 \leqslant b \leqslant q-1$.

$$\begin{aligned} \delta_q(d-1) &= \frac{1}{q^a}\left(1 - \frac{b}{q}\right), \\ \delta_q(d-2)\left(1 - \frac{1}{q}\right) &= \frac{1}{q^a}\left(1 - \frac{b-1}{q}\right)\left(1 - \frac{1}{q}\right) = \delta_q(d-1)\left(1 + \frac{1}{q-b}\right)\left(1 - \frac{1}{q}\right). \end{aligned}$$

Thus the ratio between the bounds grows from 1 when $b = 1$ to $2(1 - \frac{1}{q})$ when $b = q - 1$.

Conclusions

We feel that Theorem 2.1 is an important step towards identifying the right list-decoding radius for Reed-Muller codes. The tight examples of configurations with large list-size at radius $\delta_q(d)$ stem from low-dimensional polynomials [GKZ08]. Theorem 2.1 shows that low-dimensional polynomials are not an obstacle to the GKZ conjecture, which might be considered as evidence in its favor. The weakness of our argument is in applying the Johnson bound for the high-dimensional case. Indeed, we believe that the quantity $\delta_q^h(d)$ itself might have a significant role to play in identifying the right list-decoding radius and (dis)proving the GKZ conjecture. We propose determining its precise value and resolving Conjecture 2 as natural open problems.

Acknowledgments

I thank Prasad Raghavendra and Venkatesan Guruswami for numerous enjoyable discussions about this problem, without which this paper would not exist. I thank Sergey Yekhanin, David Zuckerman, Jaikumar Radhakrishnan, Madhu Sudan, Amir Shpilka, Eli Ben-Sasson, Irit Dinur, Ran Raz and Alex Samorodnitsky for generously sharing their time and enthusiasm with me. Thanks to Venkatesan Guruswami, Ryan O'Donnell, Rocco Servedio and Tali Kaufman for useful pointers to the literature.

References

[AGS03] A. Akavia, S. Goldwasser, and S. Safra. Proving hard-core predicates using list decoding. In *Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS'03)*, 2003.

[AKK+05] N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron. Testing Reed-Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005.

[ALM+98] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.

[AS98] S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[AS03] S. Arora and M. Sudan. Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003.

[Ass92] E. F. Assmus. On the Reed-Muller codes. *DMATH: Discrete Mathematics*, 107, 1992.

[BCH+96] M. Bellare, D. Coppersmith, J. Hstad, M. Kiwi, and M. Sudan. Linearity testing in characteristic two. *IEEE Transactions on Information Theory*, 42(6):1781–1795, 1996.

[BKW03] A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.

[BLR93] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

[DGKS08] I. Dinur, E. Grigorescu, S. Kopparty, and M. Sudan. Decodability of group homomorphisms beyond the Johnson bound. In *Proc.* 40^{th} *ACM Symposium on Theory of Computing (STOC'08)*, pages 275–284, 2008.

[Eli57] P. Elias. List decoding for noisy channels. Technical Report 335, Research Laboratory of Electronics, MIT, 1957.

[FGKP06] V. Feldman, P. Gopalan, S. Khot, and A. K. Ponnuswami. New results for learning noisy parities and halfspaces. In *Proc.* 47^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'06)*, 2006.

[GGI+02] A. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss. Near-optimal sparse Fourier representations via sampling. In *Proc.* 34^{th} *ACM Symposium on the Theory of Computing (STOC'02)*, pages 389–398, 2002.

[GGR09] P. Gopalan, V. Guruswami, and P. Raghavendra. List-decoding tensor products and interleaved codes. In *Proc.* 41^{st} *ACM Symposium on the Theory of Computing (STOC'09)*, 2009.

[GKS07] P. Gopalan, S. Khot, and R. Saket. Hardness of reconstructing multivariate polynomials over finite fields. In *Proc.* 48^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'07)*, pages 349–359, 2007.

[GKZ08] P. Gopalan, A. Klivans, and D. Zuckerman. List-decoding Reed-Muller codes over small fields. In *Proc.* 40^{th} *ACM Symposium on the Theory of Computing (STOC'08)*, 2008.

[GL89] O. Goldreich and L. Levin. A hard-core predicate for all one-way functions. In *Proc.* 21^{st} *ACM Symposium on the Theory of Computing (STOC'89)*, pages 25–32, 1989.

[GOS+09] P. Gopalan, R. O'Donnell, A. Shpilka, R. Servedio, and K. Wimmer. Testing Fourier dimensionality and sparsity. In *Manuscript*, 2009.

[GRS00] O. Goldreich, R. Rubinfeld, and M. Sudan. Learning polynomials with queries: The highly noisy case. *SIAM J. Discrete Math.*, 13(4):535–570, 2000.

[GS99] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and Algebraic-Geometric codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999.

[GT07] B. Green and T. Tao. The distribution of polynomials over finite fields, with applications to the gowers norms. In *Submitted*, 2007.

[Gur04] V. Guruswami. *List Decoding of Error-Correcting Codes*, volume 3282 of *Lecture Notes in Computer Science*. Springer, 2004.

[Gur06] V. Guruswami. *Algorithmic Results in List Decoding*, volume 2 of *Foundations and Trends in Theoretical Computer Science*. Now Publishers, 2006.

[HW03] J. Hstad and A Wigderson. Simple analysis of graph tests for linearity and pcp. *Random Struct. Algorithms*, 22(2):139–160, 2003.

[Joh62] S. M. Johnson. A new upper bound for error-correcting codes. *IEEE Transactions on Information Theory*, 8:203–207, 1962.

[Joh63] S. M. Johnson. Improved asymptotic bounds for error-correcting codes. *IEEE Transactions on Information Theory*, 9:198–205, 1963.

[JPRZ04] C. Jutla, A. Patthak, A. Rudra, and D. Zuckerman. Testing low-degree polynomials over prime fields. In *Proc.* 45^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'04)*, 2004.

[KKL88] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proc.* 29^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'88)*, pages 68–80, 1988.

[KKMS05] A. T. Kalai, A. R. Klivans, Y. Mansour, and R. A. Servedio. Agnostically learning halfspaces. In *Proc.* 46^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'05)*, pages 11–20, 2005.

[KL08] T. Kaufman and S. Lovett. Worst-case to average-case reductions for polynomials. In *Proc.* 49^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'08)*, 2008.

[KM93] E. Kushilevitz and Y. Mansour. Learning decision trees using the Fourier spectrum. *SIAM Journal of Computing*, 22(6):1331–1348, 1993.

[KOS02] A. Klivans, R. O'Donnell, and R. Servedio. Learning intersections and thresholds of halfspaces. In *Proc.* 43^{rd} *IEEE Symp. on Foundations of Computer Science (FOCS'02)*, pages 177–186, 2002.

[KR04] T. Kaufman and D. Ron. Testing polynomials over general fields. In *Proc.* 45^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'04)*, 2004.

[KT70] T. Kasami and N. Tokura. On the weight structure of Reed-Muller codes. *IEEE Transactions on Information Theory*, 16(6):752–759, 1970.

[Lip89] R.J. Lipton. New directions in testing. In *Proc. DIMACS workshop on Distributed Computing and Cryptography*, 1989.

[LN97] R. Lidl and H. Neiderreiter. *Finite Fields*. Cambridge University Press, 1997.

[LN98] A. Lapidoth and P. Narayan. Reliable communication under channel uncertainty. *IEEE Transactions on Information Theory*, 44(6):2148–2177, 1998.

[Man95] Y. Mansour. Randomized interpolation and approximation of sparse polynomials. *SIAM Journal of Computing*, 24(2):357–368, 1995.

[MS77] F. MacWilliams and N. Sloane. *The Theory of Error-Correcting Codes*. North-Holland, 1977.

[Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc.* 41^{st} *ACM Symposium on the Theory of Computing (STOC'09)*, 2009.

[PW04] R. Pellikaan and X. Wu. List decoding of q-ary Reed-Muller codes. *IEEE Transactions on Information Theory*, 50(4):679–682, 2004.

[Sam07] A. Samorodnitsky. Low-degree tests at large distances. In *Proc.* 39^{th} *ACM Symposium on the Theory of Computing (STOC'07)*, pages 506–515, 2007.

[STV01] M. Sudan, L. Trevisan, and S. P. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001.

[SU05] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudorandom generator. *J. ACM*, 52(2), 2005.

[Sud97] M. Sudan. Decoding of Reed-Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180–193, 1997.

[Sud00] M. Sudan. List decoding: Algorithms and applications. *SIGACT News*, 31(1):16–27, 2000.

[Tao07] T. Tao. Structure and randomness in combinatorics. In *Proc.* 48^{th} *IEEE Symp. on Foundations of Computer Science (FOCS'07)*, 2007.

[Tre03] L. Trevisan. List-decoding using the XOR lemma. In *Proc.* 44^{th} *IEEE Symposium on Foundations of Computer Science (FOCS'03)*, page 126, 2003.

[TSZS01] A. Ta-Shma, D. Zuckerman, and S. Safra. Extractors from Reed-Muller codes. In *Proc.* 42^{nd} *IEEE Symp. on Foundations of Computer Science (FOCS'01)*, pages 638–647, 2001.

[Woz58] J. Wozencraft. List decoding. Technical Report 48:90-95, Quarterly Progress Report, Research Laboratory of Electronics, MIT, 1958.

A Applications to Learning and Testing

The machinery of Fourier analysis over $\mathbb{F}_q$ developed in previous sections allows to extend results which were previously only known to hold over $\mathbb{F}_p$ or sometimes $\mathbb{F}_2$ to arbitrary fields, we present examples from learning and testing.

Noisy Parity over all fields

We consider the problem of learning noisy parity over various fields under the uniform distribution. The Noisy Parity problem is a central problem in learning theory [BKW03, FGKP06], with connections to coding and cryptography. There are cryptosystems whose security is based on the assumption that learning parity with random noise is hard over large fields (see [Pei09] and references therein). Feldman *et al.* show that many of the central open problems in uniform distribution learning reduce to the noisy party problem over $\mathbb{F}_2$ [FGKP06].

Feldman *et al.* [FGKP06] gave a reduction from learning parity with adversarial noise to learning parity with random noise over $\mathbb{F}_2$. No such result was known for any other field. Unlike the $\mathbb{F}_2$ case, there are many possible models for random noise over $\mathbb{F}_q$ of varying sophistication [LN98]. We present a reduction to the Discrete Memoryless Channel DMC model. This is a well-studied noise model lying in between the adversarial model and the additive random noise model. The idea, as in the [FGKP06] reduction is to fold $\mathbf{F}$ over a random subspace H. We show that with reasonable probability, the resulting randomized function is a parity function with random noise. To prove this, we need to simultaneously work with all $q-1$ Fourier polynomials, as opposed to a single polynomial in [FGKP06].

Linearity Testing

The linearity testing problem is perhaps the most basic problem in all of property testing. Given a function $F : \mathbb{F}_q^n \to \mathbb{F}_q$, we are asked to test if it is close to a linear function. This test was proposed and first analyzed in the seminal work of [BLR93]. A tight Fourier based analysis was presented in the case of $\mathbb{F}_2$ by Bellare *et al.* [BCH+96] and for prime fields $\mathbb{F}_p$ by Hastad and Wigderson [HW03]. We provide the first Fourier based analysis of the BLR test over arbitrary finite fields. Our bound matches that obtained by [BCH+96, HW03]. Further, since we work with $q-1$ polynomials, we do not need to assume that the function over $\mathbb{F}_q$ which we are testing is *folded* [HW03].

A.1 Learning Parity with Noise over Arbitrary Fields

For simplicity, we consider the problem as learning affine functions, it is easy to see that this is equivalent to learning linear functions. We use η for the (non-trivial) agreement rate rather than the noise rate.

Adversarial Channel: We are given examples $\langle x, \mathbf{F}(x)\rangle$ from some randomized function $\mathbf{F} : \mathbb{F}_q^n \to \mathbb{F}_q$ where $x \in \mathbb{F}_q^n$ is drawn uniformly at random and asked to find a linear function $L : \mathbb{F}_q^n \to \mathbb{F}_q$ so that $\mathrm{Ag}(\mathbf{F}, L) \geqslant \frac{1}{q} + \eta$, if one exists.

Discrete Memoryless Channel (DMC): In this model, we are required to learn some linear function $L : \mathbb{F}_q^n \to \mathbb{F}_q$, from samples of the form $\langle x, L(x)\rangle$, The noise is modeled by a $q \times q$ stochastic matrix W, where $w_{ij} = \Pr[\mathbf{F}(x) = j \mid L(x) = i]$. Thus the noise added may depend on $L(x)$ but not on x itself, unlike the adversarial model. But the DMC model is stronger than the additive noise model where the noise added is a random variable that is independent of the label.

However, the matrix W is not known to the algorithm, the guarantee that we are given is that

$$\Pr_{x \in \mathbb{F}_q^n}[\mathbf{F}(x) = L(x)] \geqslant \frac{1}{q} + \eta.$$

The adversarial channel model seems harder, being a generalization of the DMC model. In the adversarial setting there could a list of up to $\frac{1}{\eta^2}$ whereas in the DMC model, the affine shifts of $\alpha(x)$ are the only functions with the desired agreement. So in this model, we require an algorithm to just return $\alpha(x)$, it is easy to then figure out which shifts give good agreement.

Fix $\alpha \in \mathbb{F}_q^n \setminus 0^n$. Let $\mathbf{G} = \mathbf{G}(\alpha)$ be the randomized function obtained by folding $\mathbf{F}$ over $\alpha^\perp$. By the definition of $\mathbf{G}$, for every $c \in \mathbb{F}_q^\star$ we have

$$\hat{\mathbf{g}}^c(\beta) = \begin{cases} \hat{\mathbf{f}}^c(\beta) & \text{if } \beta = d\alpha \text{ for } d \in \mathbb{F}_q, \\ 0 & \text{otherwise.} \end{cases}$$

We claim that $\mathbf{G}$ preserves the agreement between F and every affine shift of $\alpha(x)$.

Lemma A.1. *Let $L(x) = \alpha(x) + e$ for $e \in \mathbb{F}_q$ be an affine shift of $\alpha(x)$. Then* $\mathrm{Ag}(\mathbf{G}, L) = \mathrm{Ag}(\mathbf{F}, L)$.

Proof. Partition $\mathbb{F}_q^n$ into cosets C_b where $L(x) = b$. Then

$$\mathrm{Ag}(\mathbf{F}, L) = \Pr_{x \in \mathbb{F}_q^n}[\mathbf{F}(x) = L(x)] = \Pr_{b \in \mathbb{F}_q} \Pr_{x \in C_b}[\mathbf{F}(x) = b].$$

However, for any $x' \in C_b$ we have $\mathbf{G}(x') = \mathbf{F}(x)$ where $x \in C_b$ is chosen at random. Thus

$$\mathrm{Ag}(\mathbf{F}, L) = \Pr_{b \in \mathbb{F}_q} \Pr_{x \in C_b}[\mathbf{G}(x) = b] = \Pr_{x \in \mathbb{F}_q^n}[\mathbf{G}(x) = L(x)] = \mathrm{Ag}(\mathbf{G}, L).$$

□

Of course, to sample from $\mathbf{G}$, we need to fold over $\alpha^\perp$, and the aim of the algorithm is to find α (equivalently $\alpha^\perp$). We circumvent this by showing that folding over a random subspace of suitable dimension gives a function that is close to $\mathbf{G}$ with reasonable probability. We begin with the following lemma which is an $\mathbb{F}_q$ analogue of Lemma 3 in [FGKP06].

Lemma A.2. *Fix any $\alpha \in \mathbb{F}_q^n \setminus \{0^n\}$. Pick $h_1, \ldots, h_k \in \mathbb{F}_q^n$ randomly and let $H = \mathsf{Span}(h_1, \ldots, h_k)$. Let $\mathbf{H}$ be the function obtained by folding f over H. Then*

$$\Pr_H[d(\mathbf{G}, \mathbf{H}) \leqslant q^{-(k-1)/2}] \geqslant \frac{1}{2q^k}.$$

Proof. We will show that with probability $\frac{1}{2q^k}$, the following two events hold:

1. $\alpha \in H^\perp$.
2. $\sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} \hat{\mathbf{h}}^c(\beta)^2 \leqslant 2q^{-(k-1)}$ for every $c \in \mathbb{F}_q^\star$.

We have $\alpha \in H^\perp$ if $\alpha(h_i) = 0$ for every $i \in [k]$, this happens with probability q^{-k}. Conditioning on this event, for any $\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)$, we have $\Pr_H[\beta \in H^\perp] = q^{-k}$ as the events $\alpha \in H^\perp$ and $\beta \in H^\perp$ are pairwise independent. Fix any $c \in \mathbb{F}_q^\star$. Note that

$$\hat{\mathbf{h}}^c(\beta) = \begin{cases} \hat{\mathbf{f}}^c(\beta) & \text{for } \beta \in H^\perp \\ 0 & \text{otherwise.} \end{cases} \qquad \Rightarrow \sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} \hat{\mathbf{h}}^c(\beta)^2 = \sum_{\beta \in H^\perp \setminus \mathsf{Span}(\alpha)} |\hat{\mathbf{f}}^c(\beta)|^2$$

Hence we have

$$\begin{aligned} \mathbb{E}_H \left[\sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} \hat{\mathbf{h}}^c(\beta)^2 \right] &= \mathbb{E}_H \left[\sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} |\hat{\mathbf{f}}^c(\beta)|^2 \mathbf{I}(\beta \in H^\perp) \right] \\ &= \sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} |\hat{\mathbf{f}}^c(\beta)|^2 q^{-k} \leqslant q^{-k}. \end{aligned}$$

So by Markov's inequality,

$$\Pr_H \left[\sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} \hat{\mathbf{h}}^c(\beta)^2 \geqslant \frac{2}{q^{k-1}} \right] \leqslant \frac{1}{2q} \text{by Markov's inequality.}$$

Taking the union bound over all $c \in \mathbb{F}_q^\star$, this holds for every c with probability $\frac{1}{2}$.

Thus both conditions (1) and (2) hold with probability $\frac{1}{2q^k}$. Assuming this happens, by Equation 5, we have

$$\begin{aligned} d(\mathbf{G}, \mathbf{H}) &\leqslant \frac{1}{2} \left(\sum_{c \in \mathbb{F}_q^\star} \sum_{\beta \in \mathbb{F}_q^n} |\hat{\mathbf{g}}^c(\beta) - \hat{\mathbf{h}}^c(\beta)|^2 \right)^{\frac{1}{2}} \\ &\leqslant \frac{1}{2} \left(\sum_{\beta \in \mathbb{F}_q^n \setminus \mathsf{Span}(\alpha)} \hat{\mathbf{h}}^c(\beta)^2 \right)^{\frac{1}{2}} \leqslant q^{-(k-1)/2}. \end{aligned}$$

□

We are now ready to prove our main theorem:

Theorem A.3. *Assume there is an algorithm $\mathcal{A}$ that solves the noisy parity problem over $\mathbb{F}_q$ in the* DMC *model in time $T(\eta, n)$ using $S(\eta, n) \leqslant T(\eta, n)$ samples. Then there is an algorithm $\mathcal{B}$ that solves the noisy parity problem over $\mathbb{F}_q$ in the adversarial noise model in time* $\mathrm{poly}(q, T(\eta, n))$.

Proof. Fix $\alpha \in \mathbb{F}_n^q \setminus \{0^n\}$ so that $\mathrm{Ag}(\mathbf{F}, L) > \frac{1}{q} + \eta$ for $L(x) = \alpha(x) + c$ for some $c \in \mathbb{F}_q$. Assume that the algorithm $\mathcal{A}$ uses $S = S(\eta, n)$ examples, time $T = T(\eta, n)$, and returns $\alpha(x)$ with probability $\frac{3}{4}$. Pick k so that $q^{-(k-1)/2} < \frac{1}{4S}$. We pick a random subspace H and let $\mathbf{H}$ be the function obtained by folding f over H. Assume that $d(\mathbf{G}, \mathbf{H}) \leqslant q^{-(k-1)/2}$. which happens with probability at least $\frac{1}{2q^k}$, by Lemma A.2.

Let $\mathbf{H}^S$ denote the distribution $\{\langle x_1, \mathbf{H}(x_1)\rangle, \ldots, \langle x_S, \mathbf{H}(x_S)\rangle\}$, where the $x_i \in_R \mathbb{F}_q^n$s are independent, define $\mathbf{G}^S$ similarly. By the natural coupling between $\langle x, \mathbf{G}(x)\rangle$ and $\langle x, \mathbf{H}(x)\rangle$, we have

$$\mathsf{SD}(\langle x, \mathbf{G}(x)\rangle, \langle x, \mathbf{H}(x)\rangle) \leqslant \mathbb{E}_x[\mathsf{SD}(\mathbf{G}(x), \mathbf{H}(x))] = d(\mathbf{G}, \mathbf{H}) \leqslant q^{-(k-1)/2},$$

$$\text{Hence } \mathsf{SD}(\mathbf{G}^S, \mathbf{H}^S) \leqslant Sq^{-(k-1)/2} \leqslant \frac{1}{4}.$$

Secondly, it is easy to simulate random examples from $\mathbf{H}$: following [FGKP06] draw a random example $\langle x, \mathbf{F}(x)\rangle$ and return $\langle x + h, \mathbf{F}(x)\rangle$. We sample from $\mathbf{H}^S$ and run algorithm $\mathcal{A}$ on the samples. Since $\mathcal{A}$ returns $\alpha(x)$ with probability $3/4$ when run on $\mathbf{G}^S$, it will now return $\alpha(x)$ with probability at least $3/4 - 1/10 > 1/2$. Thus the probability of finding $\alpha(x)$ is at least $\frac{1}{2q^k}$. We repeat this experiment $O(q^k) = O((qS)^2)$ times to improve the probability of success to a constant. □

A.2 Linearity Testing for all fields

We analyze the following natural generalization of the BLR [BLR93] test:

1. Pick $x, y \in \mathbb{F}_q^n$, $\lambda \in F_q^\star$ at random.
2. Test if $F(x) + \lambda F(y) = F(x + \lambda y)$.

A Fourier-based analysis over $\mathbb{F}_2$ was given by [BCH+96], it was extended to prime fields by Hastad and Wigderson [HW03]. Our analysis matches their parameters over arbitrary fields. It is clear that the test is complete, the non-trivial part is the soundness.

Theorem A.4. *If $F : \mathbb{F}_q^n \to \mathbb{F}_q$ passes the linearity test with probability $\frac{1}{q} + \eta$, there is a linear function $\alpha : \mathbb{F}_q^n \to \mathbb{F}_q$ so that $\mathrm{Ag}(F, \alpha) \geqslant \frac{1}{q} + \eta$.*

Proof. Firstly, we claim that for any linear function $\alpha : \mathbb{F}_q^n \to \mathbb{F}_q$,

$$\mathrm{Ag}(F, \alpha) = \frac{1}{q}(1 + \sum_{c \in \mathbb{F}_q^\star} \hat{f}^c(c\alpha)) \tag{16}$$

which follows from Equation 3 and the observation that the c^{th} Fourier polynomial corresponding to α is $\chi_{c\alpha}(x)$. We can arithmetize the acceptance probability as

$$\begin{aligned}
\Pr_{x,y,\lambda}[\text{Test accepts}] &= \frac{1}{q}\,\mathbb{E}_{x,y,\lambda}[1 + \sum_{c \in \mathbb{F}_q^\star} \omega^{\mathsf{Tr}(c(F(x)+\lambda F(y)-F(x+\lambda y)))}] \\
&= \frac{1}{q}\,\mathbb{E}_{x,y,\lambda}[1 + \sum_{c \in \mathbb{F}_q^\star} \omega^{\mathsf{Tr}(c(F(x)))}\omega^{\mathsf{Tr}(c\lambda F(y))}\omega^{-\mathsf{Tr}(cF(x+\lambda y))}] \\
&= \frac{1}{q}\,\mathbb{E}_{x,y,\lambda}[1 + \sum_{c \in \mathbb{F}_q^\star} f^c(x) f^{\lambda c}(y)\overline{f^c(x+\lambda y)}] \\
&= \frac{1}{q}\,\mathbb{E}_{x,y,\lambda}[1 + \sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha,\beta,\gamma} \hat{f}^c(\alpha)\hat{f}^{\lambda c}(\beta)\overline{\hat{f}^c(\gamma)}\chi_\alpha(x)\chi_\beta(y)\overline{\chi_\gamma(x+\lambda y)}] \\
&= \frac{1}{q}\,\mathbb{E}_{x,y,\lambda}[1 + \sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha,\beta,\gamma} \hat{f}^c(\alpha)\hat{f}^{\lambda c}(\beta)\overline{\hat{f}^c(\gamma)}\chi_\alpha(x)\chi_\beta(y)\overline{\chi_\gamma(x)\chi_{\lambda\gamma}(y)}] \\
&= \frac{1}{q}(1 + \sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha} |\hat{f}^c(\alpha)|^2\,\mathbb{E}_\lambda[\hat{f}^{\lambda c}(\lambda\alpha)])
\end{aligned}$$

Now assume that the test accepts with probability (exactly) $\frac{1}{q} + \eta$. So we get

$$\sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha} |\hat{f}^c(\alpha)|^2\,\mathbb{E}_\lambda[\hat{f}^{\lambda c}(\lambda\alpha)] = q\eta \quad \Rightarrow \quad \sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha} \frac{1}{q-1}|\hat{f}^c(\alpha)|^2 \sum_{\lambda} \hat{f}^{\lambda c}(\lambda\alpha) = q\eta$$

Define a distribution $\mathcal{D}$ on pairs (c, α) where we sample $c \in \mathbb{F}_q^\star$ at random, and then pick α with probability $|\hat{f}^c(\alpha)^2|$. Then we get

$$\mathbb{E}_{(c,\alpha)\leftarrow\mathcal{D}}\left[\sum_{\lambda} \hat{f}^{\lambda c}(\lambda\alpha)\right] = q\eta$$

So there exists some $c \in \mathbb{F}_q^\star, \alpha \in \mathbb{F}_q^n$ so that $\sum_\lambda \hat{f}^{\lambda c}(\lambda\alpha) = q\eta$. Writing $c' = \lambda c$, and $\alpha' = c^{-1}\alpha$, we get $\sum_{c' \in \mathbb{F}_q^\star} \hat{f}^{c'}(c'\alpha') = q\eta$. But by Equation 16, this implies that

$$\mathrm{Ag}(F, \alpha') = \frac{1}{q}(1 + q\eta) = \frac{1}{q} + \eta$$

□

B Proofs from Section 3

Proof of Fact 3.1. Let $f, g : \mathbb{F}_q \to \mathbb{R}$ denote the p.d.f.s of Y, Z respectively. We use the inner-product

$$\langle f, g \rangle = \mathbb{E}_{x \in \mathbb{F}_q} f(x)\overline{g(x)}$$

for which the function $\omega^{\mathsf{Tr}(cx)}$ are an orthonormal basis. Then we have

$$f(x) = \sum_{c \in \mathbb{F}_q} \hat{f}(c)\omega^{\mathsf{Tr}(cx)}$$

Under this inner product, we have the Fourier coefficients $\hat{f}(c) = \langle f, \omega^{Tr(cx)} \rangle$ and hence

$$\mathbb{E}_x[|f(x) - g(x)|] \leqslant (\mathbb{E}_x[|f(x) - g(x)|^2])^{\frac{1}{2}} = \left(\sum_{c \in \mathbb{F}_q^\star} |\hat{f}(c) - \hat{g}(c)|^2 \right)^{\frac{1}{2}} \tag{17}$$

where in the last line, we use the fact that $\hat{f}(\phi) = \hat{g}(\phi) = \frac{1}{q}$ since f, g are p.d.f.s over $\mathbb{F}_q$.

Observe that

$$\mathbb{E}_x[|f(x) - g(x)|] = \frac{1}{q} \sum_{x \in \mathbb{F}_q} |f(x) - g(x)| = \frac{2}{q}\mathsf{SD}(Y, Z) \tag{18}$$

and that for every $c \in \mathbb{F}_q^\star$. Finally, we rewrite $\hat{f}(c)$ and $\hat{g}(c)$ in terms of y^c and z^c.

$$\hat{f}(c) = \frac{1}{q} \sum_{x \in \mathbb{F}_q} f(x)\omega^{-\mathsf{Tr}(cx)} \;=\; \frac{1}{q} \sum_{x \in \mathbb{F}_q} f(x)\omega^{\mathsf{Tr}(-cx)} \;=\; \frac{1}{q}\,\mathbb{E}[\omega^{\mathsf{Tr}(-cY)}] = \frac{1}{q} y^{-c} \tag{19}$$

where we use $-\mathsf{Tr}(c) = \mathsf{Tr}(-c)$ which holds because Tr is $\mathbb{F}_p$-linear. Plugging equations 18 and 19 into Equation 17 we get

$$\frac{2}{q}\mathsf{SD}(Y, Z) \leqslant \frac{1}{q} \left(\sum_{c \in \mathbb{F}_q^\star} |y^c - z^c|^2 \right)^{\frac{1}{2}} \;\Rightarrow\; \mathsf{SD}(Y, Z) \leqslant \frac{1}{2} \left(\sum_{c \in \mathbb{F}_q^\star} |y^c - z^c|^2 \right)^{\frac{1}{2}} \tag{20}$$

□

Proof of Fact 3.2.

$$\begin{aligned} \mathrm{Ag}(F, G) = \mathbb{E}_x[\frac{1}{q} \sum_{c \in \mathbb{F}_q} \omega^{\mathsf{Tr}(c(F(x) - G(x)))}] \;&=\; \frac{1}{q}(1 + \sum_{c \in \mathbb{F}_q^\star} \mathbb{E}_x[f^c(x)\overline{g^c(x)}]) \\ = \;\frac{1}{q}(1 + \sum_{c \in \mathbb{F}_q^\star} \langle f^c, g^c \rangle) \;&= \frac{1}{q}(1 + \sum_{c \in \mathbb{F}_q^\star} \sum_{\alpha} \hat{f}^c_\alpha \overline{\hat{g}^c_\alpha}). \end{aligned}$$

By symmetry, we also have

$$\mathrm{Ag}(F,G) = \frac{1}{q}(1 + \sum_{c\in\mathbb{F}_q^\star}\langle g^c, f^c\rangle) = \frac{1}{q}(1 + \sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha}\overline{\hat{f}^c_\alpha}\hat{g}^c_\alpha).$$

Similarly, we can write Hamming distance between F and G as

$$\begin{aligned}\Delta(F,G) = 1 - \mathrm{Ag}(F,G) &= \frac{1}{2q}\left(2(q-1) - \sum_{c\in\mathbb{F}_q^\star}\langle f^c, g^c\rangle + \langle g^c, f^c\rangle\right)\\ &= \frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\|f^c - g^c\|_2^2 \qquad \text{Since } \|f^c\|_2 = \|g^c\|_2 = 1.\\ &= \frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha\in\hat{\mathbb{F}_q}^n}|\hat{f}^c_\alpha - \hat{g}^c_\alpha|^2\end{aligned}$$

□

Proof of Fact 3.3. We have

$$\begin{aligned}d(\mathbf{F},\mathbf{G}) &= \mathbb{E}_{x\in\mathbb{F}_q^n}[\mathsf{SD}(\mathbf{F}(x),\mathbf{G}(x))]\\ &\leqslant \mathbb{E}_{x\in\mathbb{F}_q^n}\left[\frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}|\mathbf{f}^c(x) - \mathbf{g}^c(x)|^2\right)^{\frac{1}{2}}\right] && \text{by Fact 3.1}\\ &\leqslant \frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}\mathbb{E}_{x\in\mathbb{F}_q^n}[|\mathbf{f}^c(x) - \mathbf{g}^c(x)|^2]\right)^{\frac{1}{2}} && \text{by Cauchy-Schwartz}\\ &= \frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha\in\mathbb{F}_q^n}[|\hat{\mathbf{f}}^c(\alpha) - \hat{\mathbf{g}}^c(\alpha)|^2]\right)^{\frac{1}{2}}.\end{aligned}$$

□

Proof of Fact 3.4. Consider $\mathbf{f}^c(x) = \mathbb{E}_{\mathbf{F}}[\omega^{\mathsf{Tr}(c\mathbf{F}(x))}]$. Since $\mathbf{F}(x)$ is a function of $\alpha_1(x),\ldots,\alpha_k(x)$, so is $\mathbf{f}^c(x)$. Thus, the Fourier spectrum of $\mathbf{f}^c$ is supported on $\mathsf{Span}(\alpha_1,\ldots,\alpha_k)$ for every c, so $\mathsf{Spec}(\mathbf{F}) \subseteq \mathsf{Span}(\alpha_1,\ldots,\alpha_k)$. Hence $\dim(\mathsf{Spec}(\mathbf{F})) \leqslant \dim(\mathbf{F})$.

In the other direction, fix any basis $(\alpha_1,\ldots,\alpha_k)$ for $\mathsf{Spec}(\mathbf{F})$. Then knowing $\alpha_1(x),\ldots,\alpha(x)$ fixes $\mathbf{f}^c(x)$ for all $c\in\mathbb{F}_q^\star$. But knowing the Fourier coefficients of the random variables $\mathbf{F}(x)$ allows us to determine the distribution of $\mathbf{F}(x)$. Thus $\dim(\mathbf{F}) \leqslant \dim(\mathsf{Spec}(\mathbf{F}))$. □

Proof of Fact 3.5. For every $h\in\mathsf{Inv}(\mathbf{F})$, we have for any $\lambda\in F_q$,

$$\mathbf{f}^c(x) = \mathbf{f}^c(x+\lambda h) = \sum_{\alpha}\hat{\mathbf{f}}^c(\alpha)\chi_\alpha(x)\chi_\alpha(\lambda h).$$

By the uniqueness of the Fourier expansion of every function $f : \mathbb{F}_q^n \to \mathbb{R}$, it follows that for every $\alpha \in \mathsf{Spec}(\mathbf{F})$ and $\lambda \in \mathbb{F}_q$ we have $\chi_\alpha(\lambda h) = 1$. But note that

$$\chi_\alpha(\lambda h) = \omega^{\mathsf{Tr}(\alpha(\lambda h))} = \omega^{\mathsf{Tr}(\lambda \cdot \alpha(h))}.$$

Thus $\mathsf{Tr}(\lambda \cdot \alpha(h)) = 0$ for all $\lambda \in \mathbb{F}_q$, which implies $\alpha(h) = \alpha \cdot h = 0$. Thus the Fourier spectrum is supported entirely on $\mathsf{Inv}(\mathbf{F})^\perp$, implying that $\mathsf{Spec}(\mathbf{F}) \subseteq \mathsf{Inv}(\mathbf{F})^\perp$.

In the other direction, take a basis $(\alpha_1, \ldots, \alpha_k)$ for $\mathsf{Spec}(\mathbf{F})$. For any $h \in \mathsf{Spec}(\mathbf{F})^\perp$ and $\lambda \in \mathbb{F}_q$, we have $\alpha_i(x + \lambda h) = \alpha_i(x)$. But since $\mathbf{F}(x)$ is a function of $(\alpha_1, \ldots, \alpha_k)$, we have $\mathbf{F}(x) = \mathbf{F}(x + \lambda h)$, showing that $\mathsf{Spec}(\mathbf{F})^\perp \subseteq \mathsf{Inv}(\mathbf{F})$ hence $\mathsf{Inv}(\mathbf{F})^\perp \subseteq \mathsf{Spec}(\mathbf{F})$. □

Proof of Lemma 3.6. We have

$$\begin{aligned}
\mathbf{f}^c(x) = \mathbb{E}_{\mathbf{F}}[\omega^{\mathsf{Tr}(c\mathbf{F}(x))}] &= \mathbb{E}_{h\in H}[\omega^{\mathsf{Tr}(cF(x+h))}] = \mathbb{E}_{h\in H}[f^c(x+h)] \\
&= \mathbb{E}_{h\in H}[\sum_{\alpha\in\hat{F_q}^n} \hat{f}^c(\alpha)\chi_\alpha(x+h)] \\
&= \sum_{\alpha\in\hat{F_q}^n} \hat{f}^c(\alpha)\chi_\alpha(x)\,\mathbb{E}_{h\in H}[\chi_\alpha(h)].
\end{aligned}$$

To analyze this last term, note that if $\alpha \in H^\perp$, then $\alpha(h) = 0$ for every $h \in H$, so $\mathbb{E}_{h\in H}[\chi_\alpha(h)] = 1$. On the other hand, when $\alpha \notin H^\perp$ the variable $\alpha(h)$ is uniformly distributed over $\mathbb{F}_q$, hence $\mathbb{E}_{h\in H}[\chi_\alpha(h)] = 0$. Thus we have

$$\mathbf{f}^c(x) = \sum_{\alpha\in H^\perp} \hat{f}^c(\alpha)\chi_\alpha(x).$$

□

Proof of Fact 3.7. If we define the function $G_\lambda(x) = F(x + \lambda b)$ then we have

$$g^c_\lambda(x) = \sum_{\alpha\in\hat{F_q}^n} \hat{f}^c(\alpha)\chi_\alpha(x+\lambda b) = \sum_{\alpha\in\hat{F_q}^n} \hat{f}^c(\alpha)\chi_\alpha(x)\omega^{\mathsf{Tr}(\lambda\alpha(b))}$$

We have

$$\begin{aligned}
\mathsf{Inf}_b(F) &= \mathbb{E}_{\lambda\in\mathbb{F}_q}[\Delta(F, G_\lambda)] \\
&= \mathbb{E}_{\lambda\in F_q}[\frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha\in\hat{\mathbb{F}_q}^n} |\hat{f}^c(\alpha) - \hat{g^c_\lambda}(\alpha)|^2] \\
&= \mathbb{E}_{\lambda\in F_q}[\frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0} |\hat{f}^c(\alpha)(1 - \omega^{\mathsf{Tr}(\lambda\alpha(b))})|^2] \\
&= \frac{1}{2q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0} |\hat{f}^c(\alpha)|^2 \cdot \mathbb{E}_{\lambda\in\mathbb{F}_q}[|1 - \omega^{\mathsf{Tr}(\lambda\alpha(b))}|^2] \\
&= \frac{1}{q}\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:b\cdot\alpha\neq 0} |\hat{f}^c(\alpha)|^2 \qquad (21)
\end{aligned}$$

□

Proof of Lemma 3.8. We have

$$\mathsf{Inf}_b(\mathbf{F}) = \mathbb{E}_{x\in\mathbb{F}_q^n,\lambda\in\mathbb{F}_q}[\mathsf{SD}(\mathbf{F}(x),\mathbf{F}(x+\lambda b))] = \mathbb{E}_{\lambda_\in\mathbb{F}_q}\, d(\mathbf{F}(x),\mathbf{F}(x+\lambda b))$$

We set $\mathbf{G}(x) = \mathbf{F}(x+\lambda b)$ and compute its Fourier polynomials. We have

$$\mathbf{g}^c(x) = \sum_{\alpha\in\mathbb{F}_q^n} \hat{\mathbf{f}}^c(\alpha)\chi_\alpha(x)\omega^{\mathsf{Tr}(\lambda\alpha(b))}.$$

where the last line uses the linearity of Tr. By Equation 5, we get

$$\begin{aligned}
\mathsf{Inf}_b(\mathbf{F}) &= \mathbb{E}_{\lambda\in\mathbb{F}_q}\left[\frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0}|\hat{\mathbf{f}}^c(\alpha)(1-\omega^{\mathsf{Tr}(\lambda\alpha(b))})|^2\right)^{\frac{1}{2}}\right] \\
&\leqslant \frac{1}{2}\left(\mathbb{E}_{\lambda\in\mathbb{F}_q}\left[\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0}|\hat{\mathbf{f}}^c(\alpha)(1-\omega^{\mathsf{Tr}(\lambda\alpha(b))})|^2\right]\right)^{\frac{1}{2}} \quad \text{Since } \mathbb{E}[X]\leqslant\mathbb{E}[X^2]^{\frac{1}{2}} \\
&= \frac{1}{2}\left(\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0}|\hat{\mathbf{f}}^c(\alpha)|^2\,\mathbb{E}_{\lambda\in\mathbb{F}_q}[|1-\omega^{\mathsf{Tr}(\lambda\alpha(b))}|^2]\right)^{\frac{1}{2}} \\
&= \frac{1}{\sqrt{2}}\left(\sum_{c\in\mathbb{F}_q^\star}\sum_{\alpha:\alpha(b)\neq 0}|\hat{\mathbf{f}}^c(\alpha)|^2\right)^{\frac{1}{2}} \qquad \text{Since } \mathbb{E}_\lambda[|1-\omega^{\mathsf{Tr}(\lambda)}|^2] = 2
\end{aligned}$$

□

C Relation to the work of [GKZ08] and [GGR09]

It is interesting to contrast our approach to that of [GKZ08]. While their bound also involves a *dimension reduction* step, the term refers to restricting the received word to a random low-dimensional subspace, which is very different from what we do. The GKZ algorithm is based on a self-corrector that works correctly given the right advice. The self-correction argument already shows that $\mathsf{LDR}(\mathsf{RM}_2(n,d)) = 2^{-d}$. More precisely, it proves that $\ell(\mathsf{RM}_2(n,d), 2^{-d}-\varepsilon)$ is quasi-polynomial in ε^{-1}. The deletion lemma is used only to improve the bounds to polynomial in ε^{-1}. Indeed the self-corrector is crucial to their combinatorial bound as we describe below.

Assume we are trying to decode $\mathsf{RM}_2(n,d)$ from error rates approaching 2^{-d}. Fix a codeword P from the list. If we know the polynomial P correctly on a subspace A, then we can try to self-correct the value at a random shift $b+A$ using unique decoding, since the error rate on the combined subspace containing both A and $b+A$ drops by a factor of $\frac{1}{2}$. The right advice string can be found by enumerating over all possibilities which would give a quasi-polynomial bound in ε^{-1}, while the deletion lemma gives a polynomial bound in ε^{-1}.

If we try to use this argument over $\mathbb{F}_q$, the error rate only drops by a factor of $\frac{q-1}{q}$. If we begin with error-rates approaching δ, this is insufficient to bring the error-rate within the unique-decoding radius. This appears to be a serious bottleneck in using a self-corrector over larger fields. This is especially true when d is much smaller than q, since now the error rate is very close to 1. Hence a self-corrector that requires more than one point is very unlikely to ever get noise-free examples.

Our approach is inspired by the list-decoding algorithms of [GGR09] for list-decoding linear transformations (and more generally, interleaved codes). As in their work, we use the deletion lemma to reduce the decoding problem to the low-dimensional case. In their setting, codewords are matrices and dimension refers to the rank of these matrices. However, bounding the number of low-dimensional codewords is much easier in their setting, and is done via simple combinatorial arguments. In contrast, in our setting, handling the low-dimensional case seems much harder and this is where the machinery of Fourier analysis is utilized.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/