Optimal Testing of Reed-Muller Codes with an Online Adversary

Esty Kelman
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA
Boston University, Cambridge, MA, USA

Uri Meir
Blavatnik School of Computer Science, Tel Aviv University, Israel

Kai Zhe Zheng
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Motivated by applications to property testing in the online-erasure model of Kalemaj, Raskhodnikova, and Varma (ITCS 2022 and Theory of Computing 2023), we define and analyze *semi-sample-based testers* for Reed-Muller codes. The task in Reed-Muller testing is to determine whether an input function $f : \mathbb{F}^n \to \mathbb{F}$ belongs to the Reed-Muller code or is far from it, using as few point queries to f as possible. Reed-Muller testing is a well-studied task with its roots in both the Property Testing and Probabilistically Checkable Proofs literature. The online-erasure model introduces a twist: after each query made, an adversary may erase up to t points of the input function, potentially thwarting any test in which the queries follow a predictable pattern.

Semi-sample-based testers are a hybrid between sample-based testers – which can only make uniformly random queries to the input function – and standard testers, which can choose their queries freely. They are designed with the online-erasure model in mind and operate by first choosing some subset S of the domain and then making their queries uniformly at random inside of S. We describe semi-sample-based testers for the Reed-Muller code and give an optimal analysis of their soundness.

Consequently, we show that semi-sample-based testers are indeed effective in the presence of online erasures, and thereby achieve optimal query complexity for testing the Reed-Muller code in the online-erasure model. This result improves upon prior work of Minzer and Zheng (SODA 2024). As an added bonus, we show that semi-sample-based testers also exist for the lifted affine-invariant codes of Guo, Kopparty, and Sudan (ITCS 2013), thereby providing the first known testers for these codes in the online-erasure model.

2012 ACM Subject Classification Mathematics of computing → Coding theory; Theory of computation → Error-correcting codes; Theory of computation → Streaming, sublinear and near linear time algorithms

Keywords and phrases Property testing, Low degree testing

Digital Object Identifier 10.4230/LIPIcs.CCC.2026.31

Related Version *Full Version*: https://arxiv.org/abs/2605.22813

Funding *Esty Kelman*: Supported in part by NSF TRIPODS program (award DMS-2022448), NSF Grant No. 2022446, and NSF CCF award 2227876.
Kai Zhe Zheng: Supported by the NSF GRFP DGE-2141064.

Acknowledgements Part of this research was carried out during the authors' visit to the Simons Institute for the Theory of Computing at UCB. The authors wish to thank Pooya Hatami for bringing [14] and the references therein to our attention; Corollary 1.4 there essentially proves Theorem 3 by generalizing the lower from [5, Theorem 1.2] to all fields, showing that our new algorithm is optimal.

1 Introduction

Local testing of the Reed-Muller code is a fundamental problem in property testing and complexity theory. In this problem, we are given oracle access to a multivariate function over a finite field, $f : \mathbb{F}^n \to \mathbb{F}$, and must decide whether f is a polynomial of total degree at most d, or is instead far from every polynomial of degree at most d. The notion of distance used here is the fractional Hamming distance. The goal is to design a randomized algorithm, called a *tester*, that accepts f with probability 1 in the former case and rejects f with probability at least $2/3$ in the latter case, while making as few queries to f as possible. Ideally, the query complexity should be independent of n.

The problem of Reed-Muller testing has been extensively explored over the past few decades. Its study began with the degree-1 case, which was considered in the seminal work of Blum, Luby, and Rubinfeld [11]. Their work introduced the well-known BLR linearity test and, more broadly, initiated the field of property testing. Subsequently, a large body of work gave and optimized testers for the general degree d case, $d > 1$ [1, 21, 10, 17, 32, 20, 27]. Additionally, a related variant of the Reed-Muller testing problem was instrumental in the early construction of Probabilistically Checkable Proofs (PCPs) [2, 31], and has continued to inspire subsequent research, including [30, 6, 26, 18, 7].

More recently, a new model of property testing introduced by Kalemaj, Raskhodnikova, and Varma, called the *online-erasure model* [19], has garnered significant attention. Several works have revisited Reed-Muller testing and related questions in this setting [19, 28, 5, 33, 3, 23, 24]. In this model, an adversary may *erase* up to t inputs after each query $f(x)$ is made, where typically t is assumed to be small relative to the domain size $|\mathbb{F}|^n$. Testers that work in the online-erasure model are called *online-erasure-resilient*.

We continue this line of work and give improved online-erasure-resilient testers for the Reed-Muller code. Our testers have query complexity that improves upon the previous state of the art in [28]. In addition, we give a matching lower bound showing that, for fixed field size and degree parameter, our query complexity is optimal as a function of the erasure rate parameter t.

Toward this result, we find it convenient to define a new class of testers, which we refer to as *semi-sample-based testers*. At a high level, semi-sample-based testers make their queries uniformly at random within some moderately large set of points S. In the case of Reed-Muller testing, the set $S \subset \mathbb{F}^n$ is chosen to have size independent of n, but still sufficiently large as a function of t, d, and $q := |\mathbb{F}|$. By making S large enough relative to the number of queries made, we ensure that our semi-sample-based testers are also online-erasure-resilient. Beyond Reed-Muller codes, we note that our approach extends naturally to give semi-sample-based testers for the lifted affine-invariant codes of [15]. To the best of our knowledge, these are the first testers for these codes that are resilient to online erasures.

We believe that the notion of semi-sample-based testing is conceptually useful and may find further applications in property testing, particularly in adversarial models such as the online-erasure model. To motivate this notion and explain why it arises naturally in this setting, we examine several natural candidate testers for Reed-Muller codes and show why they fail under online erasures. These candidates represent two extremes, which semi-sample-based testers naturally interpolate between. For simplicity, we restrict the discussion below to the case $d = 1$, as these issues already arise in this setting.

The Standard Tester. A first approach one might try is to run the standard tester in the online-erasure model. This tester is the well-known BLR linearity tester, which chooses $x, y \in \mathbb{F}^n$ uniformly at random, queries $f(x)$, $f(y)$, and $f(x + y)$, and then checks whether

$f(x) + f(y) = f(x + y)$. While this tester has optimal query complexity in the standard model, it fails completely in the online-erasure model. Indeed, an adversary can erase the value $f(x + y)$ after the first two queries are made, making it impossible to carry out the test. The issue here is that the queries are *too structured*.

Sample-based testing. At the other extreme, one can design a tester whose queries are as unstructured as possible. For example, one may consider a tester that queries f at uniformly random points $x \in \mathbb{F}^n$. Such testers were introduced in [12] and are called *sample-based testers*. In this case, it is unlikely that an erased point is ever queried (since $t \ll |\mathbb{F}|^n$), but this robustness comes at a cost: the tester requires at least n queries to succeed. Indeed, as shown in [13], with high probability, any set of fewer than n uniformly random points is linearly independent, and therefore there always exists a linear function that agrees with f on those points. Thus, while sample-based testers are robust to online erasures, they are not query efficient.

Semi-sample-based testing. Semi-sample-based testers interpolate between these two extremes. They operate by first choosing a subset of the domain and then taking uniformly random samples from within that subset. Our tester, in particular, chooses a k-dimensional linear subspace U and queries f at uniformly random points from U. Note that the case $k = n$ corresponds to sample-based testing, while the case $k = 1$ closely resembles the BLR tester.

By choosing the dimension of U appropriately, we obtain the best of both worlds. We achieve erasure resilience by taking k sufficiently large relative to t, while retaining query efficiency by keeping k independent of n.

1.1 Main Results

Our main result is the construction and analysis of a new tester for the Reed-Muller code using the semi-sampled-based approach. Along the way, we note that our semi-sample-based testers can also be made to work for other affine-invariant families, such as lifted affine-invariant codes [15]. Throughout this section, we say that a function $f : \mathbb{F}^n \to \mathbb{F}$ is ε-far from a given family of functions if, for every function in that family, f differs from it on at least an ε-fraction of all possible inputs.

Reed-Muller Codes. Fix a degree parameter $d \in \mathbb{N}$ and a finite field $\mathbb{F}$ of size q, where q is a power of a prime p. The starting point for our tester is the simple observation that the restriction of any polynomial of degree at most d to a linear subspace remains a polynomial of degree at most d.

Formally, let $U \subset \mathbb{F}^n$ be a linear subspace of dimension k with basis $v_1, \dots, v_k$. If $f : \mathbb{F}^n \to \mathbb{F}$ has degree at most d, then the induced k-variate function $f|_U$, defined by

$$f|_U(z_1, \dots, z_k) := f\left(\sum_{i=1}^{k} z_i v_i\right),$$

is a polynomial of degree at most d in the variables $z_1, \dots, z_k$. This observation underlies most previous Reed-Muller testers, including those of [1, 10, 17, 20].

In the standard model, the tester proceeds by choosing a subspace U of dimension $k = k_{q,d} + 1$, where

$$k_{q,d} := \left\lceil \frac{d+1}{q - q/p} \right\rceil,$$

and checking whether $f|_U$ has degree at most d by querying $f(x)$ for every point $x \in U$. Our semi-sample-based variant, which is instead parameterized by $k \geq k_{q,d}+1$ and denoted by SSB_k, queries f on a random subset of s_k points from U, rather than on all of U. The precise choice of s_k is stated in Theorem 23, but crucially it satisfies

$$s_k = O\left(q^{k_{q,d}+1}\left(\frac{e(d+k)}{k}\right)^k\right).$$

The tester $\mathbf{SSB}_{\boldsymbol{k}}$. The degree-$d$ tester with parameter k proceeds as follows:
1. Choose a linear subspace $U \subset \mathbb{F}^n$ of dimension k uniformly at random.
2. Choose s_k points from U uniformly at random, and let S denote the resulting set.
3. Accept if there exists a degree d function that agrees with f on every point in S. Otherwise, reject.

We note SSB_k makes at least as many queries as the $q^{k_{q,d}+1}$ made by the standard tester, and its query complexity increases with k. However, SSB_k has two properties which are important for us: (1) all queries made are uniformly random inside U; and (2) the *fraction* of points queried in the chosen subspace U equals to s_k/q^k which goes to 0 as k increases. These properties turn out to be crucial for our desired application to testing in the online erasure model.

The focus of this paper is thus to analyze SSB_k above, and our main result is the following.

▶ **Theorem 1** (Informal Version of Theorem 23)**.** *The tester* SSB_k *for degree d Reed–Muller codes over $\mathbb{F}^n$ satisfies the following:*
- Completeness: *If $\deg(f) \leq d$, then the tester accepts with probability* 1.
- Soundness: *If f is ε-far from every degree d function, then the tester rejects with probability at least* $\min\{1/128,\ s_k\varepsilon/8\}$.

To obtain a tester for Reed-Muller codes in the online erasure model, we set k sufficiently large with respect to the erasure parameter t, momentarily setting $\varepsilon = \Omega(1)$. Concretely, we choose $k = \Theta(d) + \log_q t$ and observe that for k at least this large, the probability of querying an erased point is negligible. This observation leads to the following result.

▶ **Theorem 2** (Informal Version of Theorem 28)**.** *For any degree parameter d, field $\mathbb{F}$ of size q, dimension n, and erasure parameter $t \leq O(q^{n-\Theta(d)})$, there exists a tester for the degree d Reed-Muller code over $\mathbb{F}^n$ in the t-online-erasure model with query complexity* $\left(O_q\left(\frac{\log t}{d}\right)\right)^d$ *that satisfies the following:*
- Completeness: *If $\deg(f) \leq d$, then the tester accepts with probability* 1.
- Soundness: *If f is $\Omega(1)$-far from every degree d function, then the tester rejects with probability at least* $\Omega(1)$.

Prior to our work, the best-known online-erasure-resilient Reed-Muller testers used $O_q(\log t)^{3d}$ queries over prime fields and $O_q(\log t)^{3d+3q}$ queries over non-prime fields. Our results improve upon these bounds in all cases. We complement this upper bound with an almost matching lower bound, showing that our query complexity is optimal (up to constant factors, for constant q, d) as a function of the erasure parameter for fixed field size and degree.

▶ **Theorem 3** (Informal Version of the Erasure-Resilient Testing Lower Bound for Reed-Muller Codes)**.** *For any degree parameter d, field $\mathbb{F}$ of size q, dimension n, and erasure parameter t, any tester for the degree d Reed-Muller code over $\mathbb{F}^n$ must make at least* $\left(\frac{\log_q t}{d}\right)^d$ *queries.*

Previously, this lower bound was shown over $\mathbb{F}_2$ in [5]. Leveraging recent results from [4, 14], we note that the same argument extends to general finite fields. A more detailed discussion, including a formal statement of Theorem 3, is provided in the full version.

Lifted Affine Invariant Codes. In addition to testing the Reed-Muller code, we note that our semi-sample-based testers can also be defined quite naturally for a family of functions called lifted affine-invariant codes [15]. These codes are a generalization of Reed-Muller codes and were previously studied in the context of property testing in [15, 16, 20]. They are defined as follows. Let $\mathcal{C}_j \subseteq \{\mathbb{F}^j \to \mathbb{F}\}$ be an *affine-invariant* family of functions, meaning $f \in \mathcal{C}_j$ if and only if $f \circ T \in \mathcal{C}_j$ for any affine transformation $T : \mathbb{F}^j \to \mathbb{F}^j$. Then, the n-dimensional lift of $\mathcal{C}_j$ is $\mathcal{C}^{j \nearrow n}$ and is the set of functions $f : \mathbb{F}^n \to \mathbb{F}$ such that $f|_U \in \mathcal{C}_j$ for every j-dimensional affine subspace $U \subset \mathbb{F}^n$.

It is straightforward to design a semi-sample-based tester for $\mathcal{C}^{j \nearrow n}$. Indeed, one may set any $k \geq j$ and adjust the number of random queries in step 2 suitably (denoted by Q_k below). Then, in step 3 the tester checks if there is a function from $\mathcal{C}^{j \nearrow k}$ agreeing with the queries. Our main result for lifted affine-invariant codes is fully formalized and proved in the full version, and is stated informally below.

▶ **Theorem 4** (Informal Version of the Semi-sample-based Tester for Lifted Affine-invariant Codes)**.** *Let $\mathcal{C}_n = \mathcal{C}^{j \nearrow n}$ be the n-dimensional lift of the affine invariant code $\mathcal{C}_j \subseteq \{\mathbb{F}^j \to \mathbb{F}\}$ and let δ_0 be the smallest minimum distance of the codes $\mathcal{C}_j, \cdots, \mathcal{C}_n$. Choose dimension $k \geq j$ and set the number of queries to be*

$$Q_k := \left\lceil \frac{100 \ln |\mathcal{C}_k|}{\delta_0} \right\rceil + 1.$$

The semi-sample-based tester for $\mathcal{C}^{j \nearrow n}$ satisfies the following:

- *Completeness: If $f \in \mathcal{C}^{j \nearrow n}$, then the tester accepts with probability* 1.
- *Soundness: If f is ε-far from $\mathcal{C}^{j \nearrow n}$, then the tester rejects with probability at least* $\min\{Q_k\varepsilon/8, 1/128\}$.

Once again, by choosing k suitably, we also get the following online-erasure-resilient tester for lifted affine-invariant codes. We state it informally below, saving the precise upper threshold for the erasure parameter to be formally defined in the full version.

▶ **Theorem 5** (Informal Version of Lifted Affine-invariant Codes Online-resilient Tester)**.** *Let $\mathcal{C}^{j \nearrow n}$ be a lifted affine-invariant code and let the erasure parameter t and dimension parameter k be chosen suitably. Then $\mathcal{C}^{j \nearrow n}$ has an $O(Q_k)$ query tester in the t-online erasure model that satisfies the following*

- *Completeness: If $f \in \mathcal{C}_n$, then the tester accepts with probability* 1.
- *Soundness: If f is $\Omega(1)$-far from $\mathcal{C}^{j \nearrow n}$, the tester rejects with probability at least* $2/3$.

Optimal Soundness Testing and Comparison with [28]. In [28], a similar tester is analyzed in the online-erasure model and a version of Theorem 2 is obtained with a higher query complexity[1] of $O_q(\log(t))^{3d}$. We discuss the differences between our result and the main result of [28]. The tester in [28] can be viewed as SSB_k with a modified third step. Instead of checking if there is a degree d-function agreeing with f on all queried points, the tester of [28] finds a structured dual codeword, over domain U, that is supported on the queried

[1] The result of [28] states $3d + 3$ in the exponent, but their analysis easily shows an exponent of $3d$.

points, and checks if $f|_U$ is indeed dual to this codeword. That is, [28] finds a function $h : U \to \mathbb{F}_q$ such that every $f : U \to \mathbb{F}_q$ with $\deg(f) \leq d$ has inner product 0 with h (i.e. $\sum_{x \in U} f(x)h(x) = 0$) and checks whether this is the case. Note that computing this inner product only requires knowing $f(x)$ for x in the support of h, hence the reason for requiring h to be supported on the queried points. From here, [28] rely on the generic soundness result of [22], which guarantees, under certain conditions satisfied here, that a Q-query tester rejects constant-far functions with probability roughly Q^{-2}. Thus, while their base tester makes $O_q(\log(t))^d$ queries, the soundness they obtain for this test requires $O_q(\log(t))^{2d}$ further repetitions to reject with constant probability, and hence leads to a version of Theorem 2 with $O_q(\log(t))^{3d}$.

By adapting the third step of SSB_k to accept if and only if there is a degree d function explaining the queried points, we are able to obtain (in Theorem 1) a better soundness analysis for SSB_k which rejects constant far functions with constant probability (and hence leads to Theorem 2 with the stated query complexity). In fact, our soundness analysis for the SSB_k is *optimal* up to constant factors. As described in [20], a tester making Q queries is called an *optimal tester* if it achieves rejection probability $\min\{\Omega(1), \Omega(Q\varepsilon)\}$ in the soundness case. Such testers are optimal in the sense that this is the best probability with which one can hope to reject an ε-far function with Q queries. Indeed, thinking of an ε-far function as a valid function with ε-fraction of its entries perturbed, the probability of even querying a perturbed point, and hence having any hope of rejecting, is at most $\min\{\Theta(1), Q\varepsilon\}$.

Showing that a fixed tester has such optimal soundness is typically a difficult task. For example, the first optimal Reed-Muller testers were only shown relatively recently: in [10] when $|\mathbb{F}| = 2$, and in [17, 20, 27] when $|\mathbb{F}| > 2$. The works mentioned only show that the specific tester they consider is optimal, and do not easily extend to other testers such as our semi-sample-based ones or the general testers in [22].

Finally, an additional factor of $(\Theta(d))^d$ in the query complexity comes from a precise choice of $k = \log_q t + \Theta(d)$ (rather than $k = \Theta(d \log_q t)$). This results in the main term of $(\Theta(\frac{\log t}{d}))^d$ in the query complexity, matching with our lower bound in Theorem 3.

1.2 Technical Overview

We describe our techniques towards our main theorem for semi-sample-based Reed-Muller testers Theorem 1. The approach for lifted affine-invariant codes is similar. Our analysis follows the inductive approach pioneered by [10]. However, for our purposes, we need to make a few key enhancements, which use techniques from the PCP literature. In particular, we show a hyperplane agreement lemma (in the low error regime), that is motivated by the Plane versus Point analysis of [31], as well as the high-error hyperplane agreement lemmas that were key to recent breakthroughs on 2-to-2 games [25] and optimal alphabet-soundness tradeoff PCPs [29]. Similar ideas were also recently used in the testing of direct-sums [33].

Proof outline. Let s be the number of queries made, and consider f of distance exactly ε from having degree d. Let $f = g + h$ be such that g is the degree d function that is closest to f and h is a function with fractional support size ε. As in [10], we split the soundness analysis into a small distance case, say where $\varepsilon < \frac{1}{10s}$, and a large distance case, where $\varepsilon \geq \frac{1}{10s}$. In the small distance case, one can show directly that with probability roughly $s \cdot \varepsilon$, the tester will only query one input from the support of h and in this case, one can easily argue that the tester will indeed reject.

More difficult to analyze is the large distance case, where one relies on induction on the ambient dimension n and assumes good soundness is guaranteed for all smaller ambient dimensions. Here we want to show that the test rejects with some constant probability, say $1/100$. The key piece of the inductive step is a lemma of the following form:

▶ **Lemma 6.** *If there are $M = q^{\Theta(1)} \cdot s$ hyperplanes $W_1, \dots, W_M$ on which the restriction of f to any W_i is $\frac{1}{100s}$-close to degree d, then f is $\frac{1}{10s}$-close to degree d.*

Observe that we can perform the semi-sample-based tester of Theorem 1 by first choosing a hyperplane W and then performing the tester on the restriction of f to W, denoted by $f|_W$, which is a function over a space of dimension $n-1$. Thus, we can write the rejection probability as $\Pr[f \text{ is rejected}] = \mathbb{E}_W \Pr[f|_W \text{ is rejected}]$. If there are more than M hyperplanes on which $f|_W$ is $\frac{1}{100s}$-close to degree d, then we get that $f|_W$ is $\frac{1}{10s}$-close to degree d and can apply the small distance analysis. On the other hand, if there are fewer than M such hyperplanes, then with probability roughly $1 - M/q^{n-1}$ over the choice of W, the function $f|_W$ is still relatively far from degree d, and by the inductive hypothesis we reject with some constant probability. In order to ensure that we still get constant soundness overall though, we must ensure that M/q^{n-1} is a small constant, and in particular, this means we must start our induction at dimension roughly $n = \log_q(M)$. Fortunately, Lemma 6 gives a strong enough bound on M above for the Reed-Muller code and take n roughly $\Theta(k)$ to get Theorem 1.

We note that similar statements to Lemma 6 are shown in [10, 17], but neither of these is sufficient for our purposes. The reason is because [10] only works over field size $q = 2$ while [17] requires obtains $M = c(q) \cdot s$ for some $c(q)$ with tower type dependence on q (rather than polynomial in q as in Lemma 6), due to the use of the density Hales-Jewett theorem. Additionally, these arguments do not apply beyond Reed-Muller codes, whereas we ultimately want our techniques to generalize to other affine-invariant codes and function families.

Our proof of Lemma 6 instead employs techniques from agreement testing and PCPs. In [31], for example, a similar statement is shown except with planes in place of hyperplanes and non-trivial agreement with a degree d function on the planes, instead of being close to degree d on the planes. Similarly [25, 29] show a similar statement to Lemma 6, but with non-trivial agreement on the hyperplanes. Thus, Lemma 6 can be viewed as a low-error version of the lemmas in [31, 25, 29], and at a high level, the proof proceeds in the same way. We consider the *consistency graph* over the M hyperplanes, $W_1, \dots, W_M$, and say that two vertices $i, j \in [M]$ are adjacent if the local degree d functions on W_i, W_j, call them f_i, f_j, agree on the intersection $W_i \cap W_j$. By showing that this graph is "nearly transitive" we find a large clique, which corresponds to many of the f_i's agreeing in pairs. From here we arbitrarily extrapolate the f_i's from the large clique to some global function F and show that 1) F is close to f; and 2) F is close to degree d. By the triangle inequality, this establishes Lemma 6.

While Lemma 6 on its own right is not hard to establish with existing techniques, we believe its statement and the way we use to obtain testing results are generic enough and could find more applications, making it one novel aspect of our work. In particular, using the same outline we extend our results to families of affine-invariant properties. The main bulk of our strategy follows through, save for step 2).

In order to accomplish 2), we rely on known local testing results to show that F passes a local test with high probability and is thus close to the desired family of functions. Indeed, if all the tester's queries fall within one of the W_i's above in the clique, then the test surely passes, as here F is consistent with the legitimate function f_i. If M is sufficiently large, then this probability is high. That is, step 2) relies on a known (typically structured) test, and leads the way to obtaining unstructured tests for the same property.

Currently, we require some lower bound on the soundness of known testers for the analysis of our tester to work (specifically in step 2). For future work, one could hope to perform step 2) algebraically for some fixed properties, as was done for Reed Muller codes in [10, 17]. For non-linear affine-invariant properties, such as some of those considered in [9, 8], such an approach could yield significantly improved soundness.

2 Preliminaries

Notation. In general, we let $\mathbb{F}$ denote a finite field with characteristic p and cardinality $|\mathbb{F}| = q = p^\ell$. At times we will use a subscript such as $\mathbb{F}_q$ or $\mathbb{F}_2$ to specify a specific field size, but we omit the subscript when it is clear from context. For a function $f : \mathbb{F}^n \to \mathbb{F}$, it is well known that there is an n-variate polynomial with degree in each variable at most $|\mathbb{F}| - 1$ whose evaluation over $\mathbb{F}^n$ is equal to f. We let $\deg(f)$ denote the total degree of this polynomial. Through the paper, for the notion of distance, we will use the fractional Hamming distance. For two functions $f, g : \mathbb{F}^n \to \mathbb{F}$, the fractional Hamming distance is

$$\mathsf{dist}(f, g) := \frac{|\{x \in \mathbb{F}^n : f(x) \neq g(x)\}|}{q^n}.$$

We extend the definition to a family of functions $\mathcal{F} \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$ by

$$\mathsf{dist}(f, \mathcal{F}) := \min_{g \in \mathcal{F}} \mathsf{dist}(f, g) \ .$$

The function families that we consider typically form an error-correcting code, meaning any two members of the family are far apart in fractional hamming distance. For a function family $\mathcal{F}$, we define its distance as

$$\delta(\mathcal{F}) = \min_{f,g \in \mathcal{F}, f \neq g} \mathsf{dist}(f, g).$$

We let $\mathcal{G}_k$ be the set of linear subspaces of dimension k inside of $\mathbb{F}^n$, where the ambient space $\mathbb{F}^n$ will always be clear from context. We denote the number of distinct linear subspaces of dimension k inside a linear space of dimension n by

$$\begin{bmatrix} n \\ k \end{bmatrix}_q := |\mathcal{G}_k| = \prod_{i=0}^{k-1} \frac{q^n - q^i}{q^k - q^i}.$$

We will also work with affine subspaces of dimension k, which are sets of the form $x + U = \{x + y \mid y \in U\}$, where $U \in \mathcal{G}_k$. We use $\dim(U)$ to denote the dimension of U when it is either an affine or linear subspace. It will always be clear from context which kind of subspace U is.

Affine Invariance. A transformation $T : \mathbb{F}^n \to \mathbb{F}^n$ is called an affine transformation if it can be expressed as $T(x) = Mx + b$, for some matrix $M \in \mathbb{F}^{n \times n}$ and affine shift $b \in \mathbb{F}^n$.

A family of functions $\mathcal{C}_n \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$ is called affine invariant if for any affine transformation, even non-invertible, $T : \mathbb{F}^n \to \mathbb{F}^n$ the following holds: $f \in \mathcal{C}_n \implies f \circ T \in \mathcal{C}_n$, where $f \circ T$ is the function given by $f \circ T(x) = f(T(x))$. Note that we require the above to hold even for T non-invertible.

For a k-dimensional subspace U (either linear or affine), we let $f|_U : \mathbb{F}^k \to \mathbb{F}$ denote the restriction of f to U. For our purposes, we will not care about the parametrization of U, so we can define $f|_U = f \circ T$ by taking any affine transformation $T : \mathbb{F}^k \to \mathbb{F}^n$ whose image is U. For every dimension k, let $\mathcal{C}_k = \{f|_U \mid f \in \mathcal{C}_n, U \in \mathcal{G}_k\}$.

One can check that if $\mathcal{C}_n$ is affine invariant, then so is $\mathcal{C}_k$. A common local test for $\mathcal{C}_n$, which we call the k-space test, goes as follows:

- Choose $U \in \mathcal{G}_k$ uniformly at random.
- Accept if $f|_U \in \mathcal{C}_k$. Otherwise, reject.

Note that since $\mathcal{C}_k$ is affine invariant, it does not matter what parametrization we choose for $f|_U$ when performing the test. We remark that typically $\mathcal{C}_k$ is defined using affine-subspaces (or flats), instead of only linear subspaces as we do above, and the local test considered is usually based on restrictions to flats, rather than to linear subspaces. This test is typically called the $(k-1)$-flat test, but note that we can simulate a $(k-1)$-flat test using a k dimensional linear subspace test. Indeed, instead of checking $f|_{U'}$ for a random $(k-1)$-flat U', we can instead sample a random linear subspace U of dimension k containing U'. Over random U', U is uniformly random in $\mathcal{G}_k$, and hence performing the k-space test above with $f|_U$ is just as good as the $(k-1)$-flat test.

The Reed-Muller Code. Given a finite field vector space $\mathbb{F}_q^n$, we refer to the set of degree at most d polynomials as the degree d Reed-Muller code and denote this by $\mathsf{RM}[n, q, d]$.

We will need the following well-known facts about the rate, distance and local testability of Reed-Muller codes.

▶ **Fact 1.** *The code $\mathsf{RM}[n, q, d]$ has distance $q^{-\frac{d}{q-1}}$ and rate at most $\binom{n+d}{d}/q^n$. That is, for any two distinct functions $f, g \in \mathsf{RM}[n, q, d]$,*

$$\Pr_{x\in\mathbb{F}_q^n}[f(x) \neq g(x)] \geq q^{-\frac{d}{q-1}}, \quad \textit{and} \quad \frac{\log_q |\mathsf{RM}[n, q, d]|}{q^n} \leq \frac{\binom{n+d}{d}}{q^n}\ .$$

For the rate we are using the fact that $\mathsf{RM}[n, q, d]$ is spanned by all total degree at most d monomials over n variables and there are at most $\binom{n+d}{d}$ such monomials.

Finally, given field size q and degree d, we define the testing dimension of $\mathsf{RM}[n, q, d]$ to be

$$t_{q,d} = \left\lceil \frac{d+1}{q - q/p} \right\rceil + 1.$$

We will use $t_{q,d}$ to refer to the above quantity throughout the paper. The following result from [20] shows the following soundness for the local test which restricts to a random $t_{q,d} - 1$ affine subspace.

▶ **Theorem 7** ([20, Theorem 1.3])**.** *There exists a constant c_{KM} such that if f is ε-far from $\mathsf{RM}[n, q, d]$, then $\Pr_{U'\subset\mathbb{F}_q^n}[\deg(f|_{U'}) > d] \geq q^{-c_{\mathsf{KM}}} \cdot \min\{1, q^{t_{q,d}}\varepsilon\}$, where U' is a random* affine *subspace of dimension $t_{q,d} - 1$.*

As discussed above, we can also simulate the test above by restricting to a linear subspace of one higher dimension. Its soundness is an immediate consequence of Theorem 7, and we give the short proof below.

▶ **Theorem 8.** *For the same constant c_{KM} from Theorem 7 it holds that if f is ε-far from $\mathsf{RM}[n, q, d]$, then $\Pr_{U\subset\mathbb{F}_q^n}[\deg(f|_U) > d] \geq q^{-c_{\mathsf{KM}}} \cdot \min\{1, q^{t_{q,d}}\varepsilon\}$, where U is a random* linear *subspace of dimension $t_{q,d}$.*

Proof. We can sample a random $t_{q,d}$-space by first choosing a $(t_{q,d}-1)$-flat U', and then choosing a random linear subspace U containing U'. Note that U is a uniformly random $t_{q,d}$-space. It follows that,

$$\Pr_{U \subset \mathbb{F}_q^n}[\deg(f|_U) > d] = \mathop{\mathbb{E}}_{U' \subset \mathbb{F}_q^n}\left[\Pr_{U \supset U'}[\deg(f|_{U'}) > d]\right]$$
$$= \Pr_{U' \subset \mathbb{F}_q^n}[\deg(f|_{U'}) > d] \geq q^{-c_{\mathsf{KM}}} \cdot \min\{1, q^{t_{q,d}}\varepsilon\}. \quad ◀$$

Lifted Affine-Invariant Codes. Lifted affine-invariant codes were introduced by Guo, Kopparty, and Sudan [15]. We formally define the codes below.

▶ **Definition 9** ([15, Definition 1.1]). *Let $\mathcal{C} \subseteq \{\mathbb{F}^r \to \mathbb{F}\}$ be a linear affine-invariant family of functions. Then the lifted code $\mathcal{C}^{r \nearrow n}$ is the following set of functions:*

$$\{f : \mathbb{F}^n \to \mathbb{F} \mid f|_U \in \mathcal{C}_r, \forall U \text{ affine }, \dim(U) = r\}.$$

When $\mathcal{C}$ in the above definition is clear, we will often use $\mathcal{C}_n$ to refer to the code $\mathcal{C}^{r \nearrow n}$. The code $\mathcal{C}$ is often referred to as the base code of $\mathcal{C}^{r \nearrow n}$ and $\mathcal{C}^{r \nearrow n}$ is referred to as the n-dimensional lift of $\mathcal{C}$. Lifted affine-invariant codes generalize Reed-Muller codes, indeed it can be seen that $\mathsf{RM}[n, q, d]$ is the n-dimensional lift of $\mathsf{RM}[t_{q,d}-1, q, d]$. In [20], it was shown that a natural test based on restrictions to r-dimensional affine subspaces is a local test for $\mathcal{C}^{r \nearrow n}$ with optimal soundness. Their result also implies similar soundness for a test based on restricting to $(r+1)$-dimensional linear subspaces. We state this version of their result below as it is more convenient for us.

▶ **Theorem 10** ([20, Theorem 1.4]). *Given $f : \mathbb{F}^n \to \mathbb{F}$, the local test obtained by choosing a uniformly random linear subspace of dimension $r+1$ has the following guarantees:*

- *Completeness: If $f \in \mathcal{C}_n$ then* $\Pr_U[f|_U \in \mathcal{C}_{r+1}] = 1$.
- *Soundness: If f is ε-far from $\mathcal{C}_n$, then,* $\Pr_U[f|_U \notin \mathcal{C}_{r+1}] \geq q^{-O(1)} \min\{1, q^r\varepsilon\}$.

Distinguishing all members of a function family.

▷ Claim 11. Fix $f : \mathbb{F}_q^n \to \mathbb{F}_q$ and let $\mathcal{F} \subseteq \{\mathbb{F}_q^n \to \mathbb{F}_q\}$ be a family of functions with size $|\mathcal{F}| = N$ and distance $\delta(\mathcal{F}) = \delta$. Then choosing a set $S \subseteq \mathbb{F}_q^n$ of size M uniformly at random with replacement, with probability at least, $1 - (1-\delta)^M \cdot \binom{N}{2} \geq 1 - e^{-\delta M} \cdot N^2$. every pair of distinct functions in $\mathcal{F}$ disagree on at least one point in $\mathcal{F}$. As a consequence, there is at most one $F \in \mathcal{F}$ such that $f(x) = F(x), \forall x \in S$.

Proof. Fix a pair of functions $F_1, F_2 \in \mathcal{F}$. Then the probability that they agree on all points in S is at most $(1-\delta)^M$. As $|\mathcal{F}| = N$, we can union bound over all pairs of functions in $\mathcal{F}$ to get that with probability at most $(1-\delta)^M \cdot \binom{N}{2}$, there exist two distinct functions in $\mathcal{F}$ that agree on S. The result follows. ◁

We conclude that whenever $M >> \frac{\log N}{\delta}$ the probability of S not satisfying Claim 11 is negligible. For example, let $\mathcal{F} \subset \mathbb{F}_2^n \to \mathbb{F}_2$ is the class of degree at most d functions over $\mathbb{F}_2$ then, $\delta = \delta_d = 2^{-d}$ and $N = 2^{\binom{n}{\leq d}}$ to have a good set with probability at least 0.9 it is enough to take $M \approx 2^d \binom{n}{\leq d}$ samples.

3 The Hyperplane Agreement Lemma

In this section, we show, for any affine-invariant family of functions $\mathcal{C}_n$, if f is close to some function in $\mathcal{C}_n$ on sufficiently many hyperplanes, then f must in fact be close to $\mathcal{C}_n$. The quality of our bound on how many hyperplanes are needed depends on the distance of the code as well as the local testability of $\mathcal{C}_n$.

Before stating the main lemma of the section, let us describe some necessary notation regarding the family $\mathcal{C}_n$. Recall that $\mathcal{G}_j$ denotes the set of j-dimensional subspaces of $\mathbb{F}_q^n$. Fix $\mathcal{C}_n \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$ to be some affine-invariant family of functions and for each dimension k, let $\mathcal{C}_k \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$ be the subset of functions given by, $\mathcal{C}_k = \{f|_U \mid U \subseteq \mathbb{F}^n, U \in \mathcal{G}_k, f \in \mathcal{C}_n\}$. Suppose that there is a dimension t with the following properties:

▶ **Definition 12.** *We say that a code $\mathcal{C}_n$ is locally testable via the t-space test with soundness function $c(\cdot)$, if for any function $f : \mathbb{F}^n \to \mathbb{F}$ it holds that $\Pr_{U \in \mathcal{G}_t}[f|_U \notin \mathcal{C}_t] \geq c(\mathsf{dist}(f, \mathcal{C}_n))$, for some function $c : \mathbb{R} \to \mathbb{R}$ that is continuous, non-decreasing, and satisfies $c(0) = 0$.*

We fix this t for the remainder of the section and let δ_0 be the smallest minimum distance of the codes $\mathcal{C}_t, \mathcal{C}_{t+1}, \ldots, \mathcal{C}_n$. That is,

$$\delta_0 := \min_{k : t \leq k \leq n} \delta(\mathcal{C}_k). \tag{1}$$

The main lemma of this section states that if f is far from $\mathcal{C}_n$, then $f|_W$ must also be far from $\mathcal{C}_{n-1}$ on almost all hyperplanes of $W \subset \mathbb{F}^n$. Formally,

▶ **Lemma 13** (Hyperplane Agreement). *Let $\mathcal{C}_n$ be an affine invariant family of functions over $\mathbb{F}_q^n$ that is locally testable via the t-space test with soundness function $c(\cdot)$ and let δ_0 as in (1). Set $\varepsilon < \delta_0/6$ and let $M \geq \max\left\{\frac{10^5 q^4}{\varepsilon}, \frac{20q^t}{c(\varepsilon)}\right\}$. If there are M distinct hyperplanes $W_1, \ldots, W_M \subset \mathbb{F}^n$, such that for each $i = 1, \ldots, M$, the restriction of f to the hyperplane W_i, denoted by $f|_{W_i}$, is ε-close to $\mathcal{C}_{n-1}$ on W_i, then f is 4ε-close to $\mathcal{C}_n$.*

In the lemma above, M is set sufficiently large with respect to the soundness function, $c(\cdot)$, of the t-space test on $\mathcal{C}_n$, so that $\frac{20q^t}{M} \leq c(\varepsilon)$. In other words, M is sufficiently large so that if $\mathsf{dist}(f, \mathcal{C}_n) > \varepsilon$, then the testing guarantee of Definition 12 gives $\Pr_{U \in \mathcal{G}_t}[f|_U \notin \mathcal{C}_t] \geq \frac{20q^t}{M}$.

The remainder of the section is dedicated to showing Lemma 13. We follow the strategy of [31] in the plane versus point analysis. Suppose that there are many hyperplanes W_i with local degree d functions $f_i : W_i \to \mathbb{F}_q$ such that f_i and f are close. We will first show that many – in fact at least a constant fraction – of these f_i's are actually consistent with each other. Using these consistent f_i's we can then define a function F on the whole space and argue that 1) F is close to f and 2) F passes the $t_{q,d}$-space test with high probability. For both 1) and 2) we are crucially using the fact that F is defined using many consistent local functions on hyperplanes, so for nearly every point x in the domain, there are, in fact, many of these local degree d functions defined on x.

3.1 The Sampling Lemma

In this section, we show a sampling lemma regarding hyperplanes and points. The proof is similar to that of [29, Lemma 5.1], but we strengthen the analysis. In particular, Lemma 14 below includes an additive "error" term of $\sim q/M$, which improves upon a similar factor with $\sqrt{M}$ in the denominator in [29, Lemma 5.1].

We next describe the sampling process and the lemma. Let $\mathcal{W} = \{W_1, \ldots, W_M\}$ be a collection of hyperplanes in $\mathbb{F}_q^n$. Denote by μ the uniform measure over $x \in \mathbb{F}_q^n$, i.e., $\mu(x) = q^{-n}$, and by $\nu_{\mathcal{W}}$ the probability measure over $\mathbb{F}_q^n$ generated by the following sampling procedure:

1. Choose $W_i \in \mathcal{W}$ uniformly at random.
2. Choose a point $x \in W_i$ uniformly at random.

We will omit the subscript $\mathcal{W}$ and use ν when it is clear from the context. Furthermore, we naturally extend the measure on point to sets by $\mu(S) = \sum_{x \in S} \mu(x)$ and similarly for ν. Let us define for $x \in \mathbb{F}_q^n$ the set of hyperplanes containing it

$$N(x) := |\{W_i \in \mathcal{W} \mid W_i \ni x\}|. \tag{2}$$

Since Item 1 samples $W_i \ni x$ with probability $\frac{N(x)}{M}$, and conditioned on this event, Item 2 chooses x with probability $\frac{1}{|W_i|} = q^{-(n-1)}$, we get that by definition

$$\nu(x) = \frac{N(x)}{M} \cdot q^{-(n-1)}. \tag{3}$$

We turn to show the following lemma, which asserts that the measures μ and $\nu_{\mathcal{W}}$ are close whenever $\mathcal{W}$ contains sufficiently many hyperplanes.

▶ **Lemma 14** (Sampling). *For any collection $\mathcal{W} = \{W_1, \ldots, W_M\}$ of M distinct hyperplanes in $\mathbb{F}_q^n$, and any set of points $S \subseteq \mathbb{F}_q^n$, the probability measures $\mu, \nu_{\mathcal{W}}$ satisfy*

$$\frac{1}{2}\left(\mu(S) - \frac{4q}{M}\right) \le \nu_{\mathcal{W}}(S) \le 2\mu(S) + \frac{8q}{M}.$$

The key to proving the sampling lemma is analyzing the concentration of $N(x)$.

▷ Claim 15. For any $c > 0$ it holds that $\Pr_{x \in \mathbb{F}_q^n}\left[\left|N(x) - \frac{M}{q}\right| \ge \frac{c \cdot M}{q}\right] \le \frac{q}{c^2 M}$.

Proof. We analyze the random variable $N(x)$ where x is uniformly chosen from $\mathbb{F}_q^n$, showing that $\mathbb{E}[N(x)] = \frac{M}{q}$ and $\mathrm{var}[N(x)] \le \frac{M}{q}$. The lemma then follows by applying Chebyshev's inequality.

For the expectation, we write $N(x)$ as a sum of indicators, and use linearity of expectation

$$\mathbb{E}_x[N(x)] = \sum_{i=1}^{M} \mathbb{E}_x[\mathbb{1}_{x \in W_i}] = \sum_{i=1}^{M} \Pr_x[x \in W_i] = \frac{M}{q},$$

where the last equality follows from the cardinality of a hyperplane in $\mathbb{F}_q^n$.

For the variance, writing $N(x)$ as a sum, we get

$$\mathrm{var}[N(x)] = \sum_{i,j \in [M]} cov(\mathbb{1}_{x \in W_i}, \mathbb{1}_{x \in W_j}) = \sum_{i \in [M]} \mathrm{var}[\mathbb{1}_{x \in W_i}] + \sum_{i \ne j \in [M]} cov(\mathbb{1}_{x \in W_i}, \mathbb{1}_{x \in W_j}).$$

The first summation is at most M/q, as for each $i \in [M]$ we have $\mathrm{var}[\mathbb{1}_{x \in W_i}] = 1/q - 1/q^2 \le 1/q$. We next show the second summation equals 0, using the premise that all hyperplanes are distinct. Indeed, for any $i \ne j$ the intersection of $W_i \ne W_j$ is a subspace of dimension $n-2$ and cardinality M/q^2. Thus, for any $i \ne j$ we get

$$\begin{aligned} cov(\mathbb{1}_{x \in W_i}, \mathbb{1}_{x \in W_j}) &= \mathbb{E}[\mathbb{1}_{x \in W_i} \cdot \mathbb{1}_{x \in W_j}] - \mathbb{E}[\mathbb{1}_{x \in W}] \cdot \mathbb{E}[\mathbb{1}_{x \in W'}] \\ &= \Pr_x[x \in W_i \cap W_j] - \Pr_x[x \in W_i] \cdot \Pr_x[x \in W_j] = 1/q^2 - (1/q)^2 = 0. \end{aligned}$$

Overall, we get that $\mathrm{var}[N(x)] \le M/q$, as required. ◁

We are now ready to prove the sampling lemma.

Proof of Lemma 14. For each integer i, let

$$m_i := \left|\left\{x \in \mathbb{F}_q^n \mid \frac{2^i \cdot M}{q} \leq N(x) < \frac{2^{i+1} \cdot M}{q}\right\}\right|.$$

By Claim 15, we get that $\frac{m_i}{q^n} \leq \frac{q}{(2^i-1)^2 \cdot M}$.

To complete the picture, let $m_0 = \left|\left\{x \in \mathbb{F}_q^n \mid N(x) \leq M/q\right\}\right|$. Towards the upper bound, we expand $\nu_{\mathcal{W}}(S)$ and perform a dyadic partitioning with $m_0 \leq |S|$:

$$\nu_{\mathcal{W}}(S) = \sum_{x \in S} \nu_{\mathcal{W}}(x) = \sum_{x \in S} \frac{N(x)}{M \cdot q^{n-1}} \leq \frac{1}{M \cdot q^{n-1}} \left(\frac{2 \cdot M \cdot |S|}{q} + \sum_{i=1}^{\infty} \frac{m_i \cdot 2^{i+1} \cdot M}{q} \right).$$

Once we multiply, the first term in the parenthesis contributes $2\mu(S)$, whereas the contribution of the second summand can be upper bounded by

$$\frac{1}{q^n} \cdot \sum_{i=1}^{\infty} m_i \cdot 2^{i+1} \leq \frac{q}{M} \cdot \sum_{i=1}^{\infty} \frac{2^{i+1}}{(2^i-1)^2} \leq \frac{8q}{M},$$

where the last inequality is since for all $i \geq 1$, it holds that $\frac{2^{i+1}}{(2^i-1)^2} \leq \frac{8}{2^i}$, and the infinite sum of the latter is at most 8.

For the lower bound, recall $\nu_{\mathcal{W}}(x) = \frac{N(x)}{M \cdot q^{n-1}}$ from (3). We expand

$$\nu_{\mathcal{W}}(S) = \sum_{x \in S} \nu_{\mathcal{W}}(x) \geq \sum_{\{x \in S : N(x) \geq M/(2q)\}} \nu_{\mathcal{W}}(x) \geq |\{x \in S \mid N(x) \geq M/(2q)\}| \cdot \frac{1}{2q^n}.$$

We next use Claim 15 with the parameters above, which gives $\Pr_x\left[N(x) < \frac{M}{2q}\right] \leq \frac{4q}{M}$.

It follows that $|\{x \in S \mid N_{\mathcal{W}}(x) \geq M/(2q)\}| \geq q^n \left((\mu(S) - \frac{4q}{M}\right)$. Putting everything together, we get the desired lower bound: $\nu_{\mathcal{W}}(S) \geq \frac{1}{2}\left(\mu(S) - \frac{4q}{M}\right)$. ◀

3.2 Proof of Lemma 13

Suppose $\mathcal{W} = \{W_1, \dots, W_M\}$ is the set of distinct hyperplanes on which f is ε-close $\mathcal{C}_{n-1}$. For each W_i, let $f_i \in \mathcal{C}_{n-1}$ be this function, so $\mathsf{dist}(f_i, f|_{W_i}) \leq \varepsilon$, for each $1 \leq i \leq M$. Note that if $\varepsilon \geq 1/4$ then the result is trivial, so for the remainder of the proof we suppose $\varepsilon < \frac{1}{4}$.

We now construct the following graph $G = (V, E)$, typically referred to as the *consistency graph*. The vertex set is $V = \{1, \dots, M\}$ while the set of edges E consists of all pairs (i, j) such that the functions f_i and f_j are consistent: $f_i|_{W_i \cap W_j} = f_j|_{W_i \cap W_j}$.

If $W_i \cap W_j$ is empty, then we also consider f_i and f_j consistent and have $(i, j) \in E$. We will first show that G contains a large clique. Call a graph *transitive* if it is an edge disjoint union of cliques. Define the following "measure" of non-transitivity:

$$\beta(G) = \max_{(i,j) \notin E} \Pr_{k \in V}[(i, k), (j, k) \in E].$$

Clearly, a graph G is transitive if and only if $\beta(G) = 0$. The following lemma from [31] gives an approximate version of this, quantifying how close is G to be transitive in terms of $\beta(G)$.

▶ **Lemma 16** ([31, Lemma 2]). *A graph G can be made transitive by removing at most $3\sqrt{\beta(G)}|V|^2$ edges.*

In our next two lemmas we show that G has many edges and then bound $\beta(G)$. Along with Lemma 16, these steps will establish that G contains a large clique. The large clique corresponds to many local degree d functions f_i that are consistent with one another, and using these functions we will extrapolate a degree d function that is close to f.

▶ **Lemma 17.** *The graph G has at least $0.4M^2$ edges.*

Proof. For each W_i, let $S_i \subseteq W_i$ be the set of points on which f_i and f differ. We have that $\frac{|S_i|}{|W_i|} \leq \varepsilon$, by assumption.

Now fix a W_i. We will apply Lemma 14 inside of W_i. For each j, let $W_j' = W_i \cap W_j$. Let $\mathcal{W}' = \left\{W_j' \mid \frac{W_j' \cap S_i}{W_j'} \geq 3\varepsilon\right\}$, let $\mathcal{I} = \{j \in [M] \mid W_j' \in \mathcal{W}'\}$, and say $M' = |\mathcal{W}'|$. Now consider the size of S_i inside of W_i under the measure $\nu_{\mathcal{W}'}$. We have, $\nu_{\mathcal{W}'}(S_i) \geq 3\varepsilon$.

By Lemma 14 $\nu_{\mathcal{W}'}(S_i) \leq 2\varepsilon + 5q/M'$.

Combining both bounds, it follows that, $M' \leq \frac{8q}{\varepsilon}$.

Finally, we can bound $|\mathcal{I}| \leq 8q^3/\varepsilon$ by noting that for any $W' \in \mathcal{W}'$, there can be at most q^2 hyperplanes W such that $W \cap W_i = W'$.

Thus, for a fixed i,

$$\Pr_{W_j}\left[\frac{|W_i \cap W_j \cap S_i|}{|W_i \cap W_j|} \geq 3\varepsilon\right] \leq \Pr_{j \in [M]}[j \in \mathcal{I}] \leq \frac{8q^3}{\varepsilon M}.$$

Now choose W_i, W_j uniformly at random. By a union bound, it follows that, with probability at least $1 - \frac{16q^3}{\varepsilon M}$, we have that both

$$\frac{|W_j \cap W_i \cap S_i|}{|W_j \cap W_i|} \leq 3\varepsilon \quad \text{and} \quad \frac{|W_j \cap W_i \cap S_j|}{|W_j \cap W_i|} \leq 3\varepsilon.$$

For such i, j, we get that

$$\Pr_{x \in W_i \cap W_j}[f_i(x) = f_j(x)] \geq 1 - \Pr_{x \in W_i \cap W_j}[x \in S_i \cup S_j] \geq 1 - 6\varepsilon > 1 - \delta_0,$$

so $f_i|_{W_i \cap W_j} = f_j|_{W_i \cap W_j}$ by (1). Thus, choosing i, j randomly, we get that they are adjacent in G with probability at least $1 - \frac{16q^3}{\varepsilon M} \geq 0.9$. ◀

▶ **Lemma 18.** *Suppose W_i, W_j are distinct hyperplanes such that $f_i|_{W_i \cap W_j} \neq f_j|_{W_i \cap W_j}$. Then,*

$$\Pr_{W_k \in \mathcal{W}}\left[f_i|_{W_i \cap W_k} \equiv f_k|_{W_i \cap W_k} \bigwedge f_j|_{W_j \cap W_k} \equiv f_k|_{W_j \cap W_k}\right] \leq \frac{1}{100}.$$

Proof. Let $W = W_i \cap W_j$. For each k let $W_k' = W_k \cap W$, let $f_i' = f_i|_W$, and let $f_j' = f_j|_W$. Also let

$$\mathcal{W}' = \{W \cap W_k \mid W_k \in \mathcal{W} \setminus \{W_i, W_j\}, f_i'|_{W_k'} = f_j'|_{W_k'}\},$$

and let $\mathcal{I} = \{k \in [M] \mid W \cap W_k \in \mathcal{W}'\}$. Note that the probability of interest is at most,

$$\Pr_{W_k \in \mathcal{W}}[f_i|_{W_i \cap W_j \cap W_k} = f_j|_{W_i \cap W_j \cap W_k}] = \frac{|\mathcal{W}'|}{|\mathcal{W}|} = \frac{|\mathcal{I}|}{M} \leq \frac{q^3|\mathcal{W}'|}{M},$$

and therefore we may focus on bounding $M' := |\mathcal{W}'|$. Above we are using the fact that for each $W' \in \mathcal{W}'$, there can be at most q^3 values k such that $W' = W \cap W_k$. This follows by noting that the dual space of W' has size q^3 and the one-dimensional dual space of W_k must be inside this space.

To bound M', we will consider W as our ambient space and apply Lemma 14 inside of W, where the set of sampling hyperplanes is $\mathcal{W}'$. Let $S \subseteq W$ be $S = \{x \in W \mid f_i'(x) \neq f_j'(x)\}$.

The key observation is that sampling $W' \in \mathcal{W}'$ and then a point $x \in W'$, the probability that $x \in S$ is actually 0. On the other hand, S is rather large since f_i', f_j' are both from a code with non-trivial distance, so together with Lemma 14, these two observations give an upper bound on the size of S.

Indeed, by the definition of $\mathcal{W}'$, we have that $\nu_{\mathcal{W}'}(S) = 0$, while on the other hand, $|S|/|W| \geq \delta_0$ by the assumption that $f'_i \neq f'_j$ combined with the definition of δ_0 from (1). Now by Lemma 14, we have $0 \geq \frac{1}{2}\left(\delta_0 - \frac{4q}{M'}\right)$. Thus, $M' \leq \frac{4q}{\delta_0}$. Since $\varepsilon \leq \delta_0/6$ and $M \geq 10^5 q^4/\varepsilon \geq 10^5 q^4/\delta_0$, we conclude that $q^3 M'/M \leq 1/100$ and the lemma follows. ◀

Lemma 18 shows that $\beta(G) \leq 1/100$, and applying Lemma 16, G can be made transitive by removing at most $0.3M^2$ edges. Combined with Lemma 17, the remaining transitive subgraph of G has at least $0.1M^2$ edges. Let $\mathcal{K}_1, \ldots, \mathcal{K}_J$ be the edge disjoint union cliques of this transitive subgraph and say $\mathcal{K}_1$ is the largest one. It follows that, $0.1M^2 \leq \sum_{I=1}^{J} |\mathcal{K}_i|^2/2 \leq |\mathcal{K}_1| \cdot M/2$, and G contains a clique of size $M_1 := |\mathcal{K}_1| = 0.2M$. From now on we let $\mathcal{K}$ denote this clique and without loss of generality let us suppose the members are $\mathcal{K} = \{1, \ldots, M_1\}$. We define the following function F using the functions $f_1, \ldots, f_{M_1}$: set $F(x) = f_j(x)$ if there is some $j \in \mathcal{K}$ such that $x \in W_j$, and $F(x) = 0$ otherwise.

We are now ready to prove the main lemma of this section.

Proof of Lemma 13. We will show that

1. F is close to f.
2. F is close to a degree d function.

Using the triangle inequality, this will establish that f is close to a degree d function.

Item 1: Proof that F is 3ε-close to f. Let $S = \{x \in \mathbb{F}^n \mid f(x) \neq F(x)\}$. Since f is ε-close to F on all hyperplanes in $\mathcal{H} := \{W_i\}_{i \in \mathcal{K}}$, we have $\nu_{\mathcal{H}}(S) \leq \varepsilon$.

By Lemma 14, it follows that

$$\varepsilon \geq \frac{1}{2}\left(\mu(S) - \frac{4q}{0.6M}\right),$$

and by our assumption on M $\mu(S) \leq 3\varepsilon$.

Item 2: Proof that F is ε-close to $\mathcal{C}_n$. We will show that F passes the t-space test with high probability. Recall that this test operates by choosing a random subspace L of dimension t and accepting if and only if $F|_L \in \mathcal{C}_t$. By Definition 12 we have,

$$\Pr_{L \in \mathcal{G}_t}[F|_L \notin \mathcal{C}_t] \geq c(\mathsf{dist}(F, \mathcal{C}_n)). \tag{4}$$

Next, we upper bound the rejection probability to bound the distance of F from the family of functions $\mathcal{C}_n$. Let $\mathcal{L} = \{L \in \mathcal{G}_t \mid \nexists i \in \mathcal{K}, W_i \supseteq L\}$.

We have

$$\Pr_{L \in \mathcal{G}_t}[F|_L \notin \mathcal{C}_t] \leq \Pr_{L \in \mathcal{G}_t}[\nexists i \in \mathcal{K}, W_i \supseteq L].$$

We will follow similar calculations to those of Claim 15. Here, for any $L \in \mathcal{G}_t$, we let

$$N(L) = |\{i \in \mathcal{K} \mid W_i \supseteq L\}|.$$

We can now express the upper bound using $N(\cdot)$.

$$\Pr_{L \in \mathcal{G}_t}[\nexists i \in \mathcal{K}, W_i \supseteq L] = \Pr_{L \in \mathcal{G}_t}[N(L) = 0].$$

To further bound the probability, we next bound the first and second moments of the random variable $N(L)$ for uniform random $L \in \mathcal{G}_t$. Let $p_1 := \Pr_{L \in \mathcal{G}_t}[L \subseteq W]$. For the expectation, note that by the linearity of expectation $\mathbb{E}_{L \in \mathcal{G}_t}[N(L)] = \sum_{i=1}^{M} \Pr_{L \in \mathcal{G}_t}[L \subseteq W_i] = p_1 M$.

For the variance analysis, similarly to p_1 we define the probability of a uniform subspace to be contained in a fixed $n-2$ dimensional subspace, which we will use as an intersection of two distinct hyperplanes W, W'.

$$p_2 := \Pr_{L \in \mathcal{G}_t}[L \subseteq W \cap W'].$$

We next bound the $\mathbb{E}_{L\in\mathcal{G}_t}\left[N(L)^2\right]$:

$$\begin{aligned}\mathbb{E}_{L\in\mathcal{G}_t}\left[N(L)^2\right] &= \mathbb{E}_{L\in\mathcal{G}_t}\left[\left(\sum_{i=1}^{M} \mathbb{1}_{L\subseteq W_i}\right)^2\right] \\ &\le M\cdot \Pr_{L\in\mathcal{G}_t}[L\subseteq W_i] + M^2\cdot \Pr_{L\in\mathcal{G}_t}[L\subseteq W_i\cap W_{i'}] = p_1 M + p_2 M^2 \le p_1 M + p_1^2 M^2.\end{aligned}$$

The last inequality follows from

$$p_2 = \frac{\begin{bmatrix} n-2 \\ t\end{bmatrix}_q}{\begin{bmatrix} n \\ t\end{bmatrix}_q} \le \left(\frac{\begin{bmatrix} n-1 \\ t\end{bmatrix}_q}{\begin{bmatrix} n \\ t\end{bmatrix}_q}\right)^2 = p_1^2.$$

Now,

$$\Pr_{L\in\mathcal{G}_t}[N(L)=0] \le \frac{\operatorname{var}(N(L))}{\mathbb{E}[N(L)^2]} = 1 - \frac{\mathbb{E}[N(L)]^2}{\mathbb{E}[N(L)^2]} \le 1 - \frac{p_1^2M^2}{p_1M + p_1^2M^2} \le \frac{1}{1+p_1M} \le \frac{1}{p_1 M}.$$

Note that,

$$p_1 = \frac{\begin{bmatrix} n-1 \\ t\end{bmatrix}_q}{\begin{bmatrix} n \\ t\end{bmatrix}_q} = \frac{\prod_{i=0}^{t-1}(q^{n-1}-q^i)}{\prod_{i=0}^{t-1}(q^n-q^i)} = q^{-t}\frac{\prod_{i=1}^{t}(q^n-q^i)}{\prod_{i=0}^{t-1}(q^n-q^i)} = q^{-t}\frac{q^n-q^t}{q^n-1} \ge \frac{1}{2q^t}.$$

By our assumption on the size of M, $\frac{1}{p_1\cdot M} \le \frac{2q^t}{M} \le c(\varepsilon)$.

Combining with (4), and recall that $c(\cdot)$ is non-decreasing, we conclude $\mathsf{dist}(f, \mathcal{C}_n) \le \varepsilon$. ◀

4 Sample-Based Testers for All Affine Invariant Families

In this section, we describe our sample-based tester that works for any fixed family of functions $\mathcal{C}_n \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$. The tester samples a certain number of points and makes a decision based on this view. Since we wish the tester to have a one-sided error (always accept inputs from the family $\mathcal{C}_n$), the algorithm does the only thing possible. It rejects only when there is no codeword that can explain the sampled view.

The $\mathcal{T}$ for $\mathcal{C}_n$.

1. Sample s points from $\mathbb{F}^n$ uniformly with replacement. Let S denote the set of points sampled.
2. Accept if there is a function $g \in \mathcal{C}_n$ satisfying $g|_S = f|_S$. Otherwise, Reject.

▶ **Lemma 19.** *Let f be ε-far from $\mathcal{C}_n$, if $s \ge \frac{\ln(2|\mathcal{C}_n|)}{\varepsilon}$ then $\mathcal{T}$ rejects with probability at least $1/4$.*

Proof. For any $g \in \mathcal{C}_n$, we have $\mathsf{dist}(f,g) \geq \varepsilon$ and therefore the probability g is consistent with f on s_n uniform random samples is at most $(1-\varepsilon)^{s_n}$. Union bounding over all members of $\mathcal{C}_n$, the probability that there exists $g \in \mathcal{C}_n$ such that $f|_S \equiv g|_S$, which is the case of acceptance, is at most $(1-\varepsilon)^s|\mathcal{C}_n| \leq \exp(-\varepsilon s)|\mathcal{C}_n| \leq 1/2$. The last inequality is by the bound on s. ◀

Though the result of Lemma 19 is easy and quite general, it suits well when the tester receives the proximity parameter ε as an input and can decide how many samples to make based on the proximity parameter.

Since we are going to apply the sample-based tester on functions with varying distances, and yet we need to squeeze all the probability of rejecting non-codewords. We will use an oblivious (to the proximity parameter) algorithm; the only difference is the setting of the number of samples it uses. We show that such a tester is optimal in the sense discussed in Section 1.1.

Our approach to the analysis now is based on the *learn-and-test* paradigm. The tester operates by choosing twice the amount of points needed to interpolate a unique function from $\mathcal{C}_n$; crucially, this amount is independent of the proximity parameter ε. The first half of the points are used to reduce to one candidate function $g \in \mathcal{C}_n$, and the second half of the points are used to check if the tested function is indeed g.

▶ **Lemma 20.** *Let f be ε-far from $\mathcal{C}_n$, if $s \geq \frac{10 \ln |\mathcal{C}_n|}{\delta}$, where $\delta = \delta(\mathcal{C}_n)$ is the minimal distance of the code $\mathcal{C}_n$, then $\mathcal{T}$ rejects with probability at least* $\min\{1/4, s\varepsilon/8\}$.

Proof. Let S_1 be the first $\frac{5 \ln |\mathcal{C}_n|}{\delta}$ points sampled in S, and let S_2 be the remaining points. Let E be the event that there is at most one $g \in \mathcal{C}_n$ satisfying $g|_{S_1} = f|_{S_1}$. Then the probability that the tester rejects is at least $\Pr[E] \cdot \Pr_S[g|_{S_2} \neq f|_{S_2} \mid E]$, where in the probability g refers to the unique function in $\mathcal{C}_n$ agreeing with f on S_1. By Claim 11, $\Pr_S[\overline{E}] \leq \exp(-\delta|S_1|)|\mathcal{C}_n|^2 \leq 1/2$. Furthermore, by the assumption that f is ε-far from $\mathcal{C}_n$, we have, $\Pr_S[g|_{S_2} \neq f|_{S_2} \mid E] \geq 1-(1-\varepsilon)^{s/2} \geq \min\left\{\frac{1}{2}, \frac{s\varepsilon}{4}\right\}$. The last inequality is proved in the proof of [5, Lemma 3.1]. The result follows. ◀

We will use Lemma 20 in order to build our semi-sample-based tester. Here we note a quick application of it, improving on the sample complexity given by Lemma 19 whenever ε is much smaller than δ_0.

▶ **Theorem 21.** *There is a sample-based tester for $\mathcal{C}_n$ with a one-sided error that uses at most $O(s+1/\varepsilon)$ samples. Where $s = O(\ln|\mathcal{C}_n|/\delta)$ and $\delta = \delta(\mathcal{C}_n)$ is the minimal distance between two distinct codewords in $\mathcal{C}_n$.*

Proof. We claim that repeating $\mathcal{T}$ (with $s = O(\ln|\mathcal{C}_n|/\delta)$) enough times and rejecting if any of the instances rejects is a valid tester. Indeed, if f is a codeword, all samples will be consistent with f, and we always accept. If f is ε-far, then by Lemma 20 each instance accepts with probability at most $1 - \min\{1/4, s\varepsilon/8\}$. Since the instances are independent, by repeating $O\left(\frac{1}{s\varepsilon}+1\right)$ times, the probability that they all accept is at most $1/3$. For the sample complexity, we have any instance use s samples, which in total become $O\left(s\left(\frac{1}{s\varepsilon}+1\right)\right) = O(s+1/\varepsilon)$. ◀

5 Semi-Sample-Based Testers

The semi-sample-based tester for $\mathcal{C}_n$ on input function $f : \mathbb{F}_q^n \to \mathbb{F}_q$ works by first choosing a random k-dimensional subspace, A, and then running the sample-based tester on the smaller, restricted function $f|_A$.

The tester $\mathcal{T}_k(Q, \mathcal{C}_n)$.

1. Choose a k-dimensional subspace $A \subset \mathbb{F}^n$ uniformly at random.
2. Sample Q points from A uniformly with replacement. Let S denote the set of points sampled.
3. Accept if there is a function $g \in \mathcal{C}_k$ satisfying $g|_S = f|_S$. Otherwise, Reject.

We will show that with large enough k, choosing Q sufficiently large with respect to the rate and distance of $\mathcal{C}_k$, the tester $\mathcal{T}_k(Q)$ is indeed a valid tester with perfect completeness and non-trivial soundness (or even optimal soundness for certain families).

▶ Remark 22. We remark that while our focus is on query complexity rather than time complexity, the tester $\mathcal{T}_k(Q, \mathcal{C}_n)$ does have good running time when $\mathcal{C}_n$ is a linear code. In this case, checking if there exists $g \in \mathcal{C}_k$ such that $g|_S = f|_S$ can be done by solving a system of $|S|$ linear equations with $\log_q |\mathcal{C}_k| \leq |S|$ variables. In particular, the time complexity is $O(|S|^3)$.

In this section, we show our main theorem regarding semi-sample-based testers for affine invariant families of functions. At a high level, our tester for *every* affine invariant family is simply the tester $\mathcal{T}_k(Q)$ except one must choose the quantity k as well as the number of queries Q depending on the specific code. For families of functions such as Reed-Muller codes or lifted affine-invariant codes, we are able to give explicit values of k and Q, and show that the semi-sample-based tester achieves optimal soundness. For general affine invariant families however, the known soundness of current subspace based testers is not strong enough for us to give explicit values of k and Q for a valid semi-sample-based tester. We state our main theorems for testing Reed-Muller codes, lifted affine-invariant codes, and finally general affine invariant families below.

We state our main result for Reed-Muller codes and subsequently prove it in Section 5.1.

▶ **Theorem 23.** *Let $f : \mathbb{F}^n \to \mathbb{F}$, let $k \geq 8d + 3c_{\mathsf{KM}} + 150$ be a dimension parameter, let $\delta_0 := q^{-d/(q-1)}$ and set*

$$s_k := \left\lceil \frac{100 \ln |\mathsf{RM}[k, q, d]|}{\delta_0} \right\rceil + 1$$

Then $\mathcal{T}_k^f(s_k, \mathsf{RM}[n, q, d])$ is a valid tester for the degree d Reed-Muller code over $\mathbb{F}^n$. Specifically, the tester satisfies the following:

- *Completeness: If $f \in \mathsf{RM}[n, q, d]$, then the tester accepts with probability 1.*
- *Soundness: If f is ε-far from $\mathsf{RM}[n, q, d]$ then the tester rejects with probability at least* $\min\{s_k\varepsilon/8, 1/128\}$

Furthermore, the tester has query complexity s_k, which is bounded by

$$s_k = O\left(q^{t_{q,d}} \cdot \left(\frac{e(d+k)}{k}\right)^k\right) \quad . \tag{5}$$

5.1 Reed-Muller Codes: Proof of Theorem 23

As is usual in property testing, the completeness is clear. To bound the query complexity, s_k, we recall that $t_{q,d} = \lceil \frac{d}{q - q/p} \rceil + 1$, and the bounds from Fact 1, which in particular imply $\ln |\mathsf{RM}[k, q, d]| \leq \binom{d+k}{k} \cdot \ln q$, and $1/\delta_0 \leq q^{t_{q,d}-1} \leq q^{t_{q,d}} / \ln q$. Overall, this bounds the query complexity by

$$s_k = \left\lceil \frac{100 \ln |\mathsf{RM}[k, q, d]|}{\delta_0} \right\rceil + 1 \leq 100 q^{t_{q,d}} \cdot \binom{d+k}{k}. \tag{6}$$

Equation (5) follows from standard bounds on binomial coefficients.

From here on we focus on the proof of soundness. Fix any dimension parameter k such that

$$k \geq 8d + 3c_{\mathsf{KM}} + 150 \tag{7}$$

as in Theorem 23. To simplify notation, we also drop the subscript from s_k, writing $s = s_k$ instead, and write $\mathcal{T}_k^f$ in place of $\mathcal{T}_k^f(s_k, \mathsf{RM}[n,q,d])$. Also let $\varepsilon_f := \mathsf{dist}(f, \mathsf{RM}[n,q,d])$ be the actual distance from f to $\mathsf{RM}[n,q,d]$ and we assume that f is ε-far from $\mathsf{RM}[n,q,d]$, meaning $\varepsilon_f \geq \varepsilon$. Finally, we remark that by Fact 1, δ_0 is the minimum distance of $\mathsf{RM}[t,q,d]$ for all $t_{q,d} \leq t \leq n$.

Similar to the soundness analysis of the standard subspace test for Reed-Muller codes in [10], we split our soundness analysis into two cases based on if the magnitude of ε_f is small or large. The analyses of these two cases together will show the desired soundness for Theorem 23.

Before going into the proof, we state and prove two key claims. The first claim shows that s_k is sufficiently smaller than q^k.

▷ Claim 24. Fix any constant $c > 0$. For any $k \geq 8d + 3c + 24$, it holds that $\frac{s_k}{q^k} \leq q^{-c}$.

Proof. Define $B_k := 100q^{t_{q,d}} \cdot \binom{d+k}{d}$, the upper bound on s_k from (6). We focus on the expression B_k/q^k instead, bound it for a proper dimension k_0, and show it decreases multiplicatively by a significant factor as k increases until for certain $k(c)$ the whole quotient is small enough, immediately implying the lemma. One "step" changes the quotient by a factor of

$$\frac{\frac{B_{k+1}}{q^{k+1}}}{\frac{B_k}{q^k}} = \frac{1}{q} \cdot \frac{B_{k+1}}{B_k} = \frac{1}{q} \cdot \left(1 + \frac{d}{k+1}\right) , \tag{8}$$

where the second equality is due to quotient of "consecutive" binomial coefficients.

From here on, We split into two cases.

The case $q = 2$. Here, it holds that $t_{q,d} = d + 1$. We start from dimension $k_0 = 2d$ and plug in q and $t_{q,d}$ to get:

$$\frac{B_{2d}}{q^{2d}} \leq 100q^{t_{q,d}-2d} \cdot \binom{3d}{d} \leq 100q^{1-d} \cdot 2^{3d} \leq 2^{2d+8} .$$

By (8), for any choice of $k \geq 2d$ the quotient diminishes by a factor of at most

$$\frac{\frac{B_{k+1}}{q^{k+1}}}{\frac{B_k}{q^k}} = \frac{1}{q} \cdot \left(1 + \frac{d}{k+1}\right) < 2^{-1/3} .$$

Applying the above for $z \geq 3(2d + c + 8)$ steps, we get

$$\frac{B_{2d+z}}{q^{2d+z}} \leq \frac{B_{2d}}{q^{2d}} \cdot \left(2^{-1/3}\right)^z \leq 2^{2d+8-z/3} \leq 2^{-c} .$$

That is, for any dimension $k = 2d + z \geq 8d + 3c + 24$ the assertion holds for $q = 2$.

The case $q > 2$. Here we have $t_{q,d} \leq d$. Plugging this and $q \geq 3$, we directly calculate for $k_0 = d$:

We next bound the quotient for $k = d$ (for $q > 2$ we have $t_{q,d} \leq d$):

$$\frac{B_d}{q^d} \leq 100q^{t_{q,d}-d} \cdot \binom{2d}{d} \leq 100 \cdot 4^d \leq q^{2d+5} .$$

By (8), for any $k \geq d$, the quotient diminishes by a factor of

$$\frac{\frac{B_{k+1}}{q^{k+1}}}{\frac{B_k}{q^k}} = \frac{1}{q} \cdot \left(1 + \frac{d}{k+1}\right) < \frac{2}{q} < q^{-1/3} .$$

Applying the above iteratively for $z \geq 3(2d + c + 5)$ steps, we get $\frac{B_{d+z}}{q^{d+z}} \leq \frac{B_d}{q^d} \cdot q^{-z/3} \leq q^{2d+5-z/3} \leq q^{-c}$, concluding that for any dimension $k = d + z \geq 7d + 3c + 15$ the assertion holds for $q > 2$. ◁

The second claim is a specialization of Lemma 13 to degree d Reed-Muller codes with the parameters that we will use.

▷ Claim 25. Suppose there are $M = 10^6 q^{\max\{4, c_{\mathsf{KM}}\}} \cdot s$ distinct hyperplanes $W_1, \ldots, W_M \subset \mathbb{F}^n$ such that for each $i = 1, \ldots, M$, $f|_{W_i}$ is $\frac{1}{8s}$-close to a degree d function on W_i. Then f is $\frac{1}{2s}$-close to a degree d function.

Proof. We apply Lemma 13 with the ε therein set to $\frac{1}{8s}$ (which indeed satisfies $\frac{1}{8s} \leq \frac{\delta_0}{6}$). The conclusion suffices for our proof. It therefore remains to verify that M is large enough to satisfy the premise of Lemma 13. In particular, we check that

$$M \overset{?}{\geq} \max\left\{\frac{10^5 q^4}{1/(8s)}, \frac{20 q^{t_{q,d}}}{c\left(\frac{1}{8s}\right)}\right\}, \tag{9}$$

where $c(\cdot)$ is the soundness function of Definition 12 for Reed-Muller codes.

Clearly, we have $M \geq \frac{10^5 q^4}{1/(8s)}$, which leaves the need to verify M is larger than the second term. By Theorem 8, the $t_{q,d}$-flat test has soundness function $c(\varepsilon) = q^{-c_{\mathsf{KM}}} \cdot \min\{1, q^{t_{q,d}}\varepsilon\}$ for Reed-Muller codes, and in particular $c\left(\frac{1}{8s}\right) = q^{-c_{\mathsf{KM}}} \cdot \frac{q^{t_{q,d}}}{8s}$, since $8s \geq q^{t_{q,d}}$. This implies

$$\frac{20 q^{t_{q,d}}}{c\left(\frac{1}{8s}\right)} = 20 q^{t_{q,d}} \cdot q^{c_{\mathsf{KM}}} \cdot \frac{8s}{q^{t_{q,d}}} = 160 q^{c_{\mathsf{KM}}} \cdot s \leq M ,$$

which shows our choice of M indeed satisfies (9), concluding the proof. ◁

5.1.1 Small distance

► **Lemma 26.** *If* $\varepsilon_f \leq \frac{1}{2s}$ *then* $\Pr\left[\mathcal{T}_k^f \text{ rejects}\right] \geq \frac{s\varepsilon_f}{8} \geq \frac{s\varepsilon}{8}$.

Proof. Let $g^* \in \mathsf{RM}[n, q, d]$ be such that $\mathsf{dist}(f, g) = \varepsilon_f$. Let S_1 be the first $1/10$ of the points chosen in S and S_2 be the remaining $9/10$ of the points chosen in S, so $|S_2| = 9s/10$.

Let $\Delta = \{x \in \mathbb{F}^n \mid f(x) \neq g^*(x)\}$. The tester proceeds by choosing a linear subspace $U \subset \mathbb{F}^n$ of dimension k and then sampling $S_1, S_2 \subseteq U$.

For our analysis, it will be convenient to think of S as being chosen as follows.

1. Fix $U \subseteq \mathbb{F}^n$ to be a generic k-dimensional subspace of dimension U and choose $S_1' \subseteq U$ of size $|S_1|$.
2. Choose $S_2' \subseteq U$ of size $|S_2|$.
3. Choose a random, full-rank linear transformation $T : \mathbb{F}^n \to \mathbb{F}^n$.
4. Set $S_i = \{T(x) \mid x \in S_i\}$ for $i = 1, 2$.
5. Set $S = S_1 \cup S_2$.

In the above view of $\mathcal{T}_k^f$, the k-dimensional subspace fixed by the tester is $T \circ U$, and the points queried are $S \subseteq T \circ U$. To analyze the rejection probability, we define the following events:

- E_1: the event that every pair of degree d functions on U (these correspond to $\mathsf{RM}[k,q,d]$) disagree on at least one point in S_1'.
- E_2: the event that f and g^* agree on S_1, i.e. $f|_{S_1} \equiv g^*|_{S_1}$.
- E_3: the event that $S_2 \cap \Delta \neq \emptyset$.

One can check that if all three events above occur, then the tester rejects. Indeed, if E_1 and E_2 both occur, then $g^*|_{T\circ U} \in \mathsf{RM}[k,q,d]$ is the unique degree d function over $T \circ U$ that agrees with f on S_1. The tester will then reject if g^* and f disagree on any point on S_2 – in other words, if E_3 occurs. Thus,

$$\Pr[\mathcal{T}_k^f \text{ rejects}] \geq \Pr_{S_1,S_2}[E_1 \wedge E_2 \wedge E_3] = \Pr_{S_1'}[E_1] \cdot \Pr_{S_1',S_2',T}[E_2 \wedge E_3 \mid E_1].$$

We now bound each term appearing above. First note that E_1 does not depend on the identity of the subspace U. By Claim 11, we have

$$\Pr_{S_1'}[E_1] \geq 1 - e^{-\delta_0|S_1|}|\mathsf{RM}[k,q,d]|^2 \geq \frac{1}{2}.$$

For the second term, we have

$$\Pr_{S_1',S_2',T}[E_2 \wedge E_3 \mid E_1] \geq \Pr_{S_1',S_2',T}[E_3 \mid E_1] - \Pr_{S_1',S_2',T}[\overline{E_2} \mid E_1].$$

We will bound each term on the right-hand side. Let us suppose S_1' and S_2' have already been sampled, and now bound the probabilities over random T.

For the first term, we use the principle of inclusion-exclusion to get a lower bound:

$$\begin{aligned}\Pr_T[E_3 \mid E_1] &\geq \sum_{x \in S_2' \setminus \{0\}} \Pr_T[T(x) \in \Delta] - \sum_{x,y \in S_2' \setminus \{0\}, x \neq y,} \Pr_T[T(x), T(y) \in \Delta] \\ &\geq \frac{9s/10 - 1}{2}\varepsilon_f - 2\frac{(9s/10-1)^2}{4}\varepsilon_f^2 \geq \frac{3s\varepsilon_f}{8}\ .\end{aligned}$$

In the first transition we are using the fact that $T(x)$ is uniformly random over random T and the fact that $T(x), T(y)$ are nearly independent for $x \neq y$. In the last transition we are using the fact that $s\varepsilon_f/4 \leq 1/8$ by assumption.

To upper bound the second term (again as a probability over T, fixing S_1', S_2'), we use a union bound:

$$\Pr_{S_1',S_2',T}[\overline{E_2} \mid E_1] \leq \sum_{x \in S_1'} \Pr_T[T(x) \in \Delta] \leq \frac{s\varepsilon_f}{10}.$$

Putting everything together, we get

$$\Pr[\mathcal{T}_k^f \text{ rejects}] \geq \Pr_{S_1'}[E_1] \cdot \Pr_{S_1',S_2',T}[E_2 \wedge E_3 \mid E_1] \geq \frac{1}{2} \cdot \frac{s\varepsilon_f}{4} \geq \frac{s\varepsilon_f}{8}\ . \qquad \blacktriangleleft$$

5.1.2 Large distance

We have already shown that the tester has the desired soundness when $\varepsilon_f \leq \frac{1}{2s}$. To complete the proof of Theorem 23, it remains to show the desired soundness when $\varepsilon_f > \frac{1}{2s}$. In particular, the following lemma suffices.

▶ **Lemma 27.** *Set $\eta := \frac{1}{128}$ and $\gamma := 10^7 q^{\max\{4, c_{\mathsf{KM}}\}} \cdot s$. The following holds for every $n \geq k$. Let $f\colon \mathbb{F}_q^n \to \mathbb{F}_q$ be such that $\mathsf{dist}(f, \mathsf{RM}[n,q,d]) = \varepsilon_f \geq 16\eta/s$. Then*

$$\Pr\left[\mathcal{T}_k^f \text{ rejects}\right] \geq \eta + \frac{\gamma}{q^n}.$$

We will analyze the large distance case by induction of n. Consider the following tester $\mathcal{T}'_k$: the tester first picks a (linear) hyperplane W and then simulates $\mathcal{T}_k^{f|_W}$. That is, the tester chooses a random k-dimensional subspace $U \subseteq W$ and then performs the sample-based tester inside of U. We note that U is still a uniformly random k-dimensional subspace in $\mathbb{F}^n$, so it follows that

$$\Pr[\mathcal{T}_k^f \text{ rejects}] = \mathbb{E}_W\left[\Pr[\mathcal{T}_k^{f|_W} \text{ rejects}]\right]. \tag{10}$$

Therefore, in order to conclude that the rejection probability is high, we will invoke Claim 25 to say that for nearly all of the hyperplanes W, the distance of $f|_W$ to $\mathcal{C}_{n-1}$ is still within a constant factor of ε_f. As $f|_W$ and $\mathcal{C}_{n-1}$ are one dimension down, we can conclude that $\mathcal{T}_k^{f|_W}$ rejects with good probability. Since Claim 25 allows us to conclude $\mathsf{dist}(f|_W, \mathcal{C}_{n-1})$ can only increase by more than a constant factor for a number of W that is independent of the dimension n, the above argument is sufficient for the inductive step.

Proof of Lemma 27. The proof is by induction on the ambient dimension, n. In the base case, $n = k$, and $\mathcal{T}_k^f$ is simply the sample-based tester of Lemma 20. Therefore, $\mathcal{T}_k^f$ rejects with probability at least $\min\{1/4, s\varepsilon_f/8\} \geq 2\eta$, using the premise and the definition of η. Next, apply Claim 24 with $c = \max\{4, c_{\mathsf{KM}}\} + 31$, noting that by (7), we have large enough k, yielding

$$\frac{\gamma}{q^k} = \frac{10^7 q^{\max\{4, c_{\mathsf{KM}}\}} \cdot s_k}{q^k} \leq 10^7 q^{\max\{4, c_{\mathsf{KM}}\}} q^{-c} < \frac{1}{128} = \eta,$$

which implies rejection probability of at least $2\eta \geq \eta + \gamma/q^n$, concluding the base case.

For the induction step, suppose that the lemma holds for testing all functions over $\mathbb{F}_q^{n'}$ that are sufficiently far from $\mathsf{RM}[n', q, d]$, for all $k \leq n' < n$. We will show that the lemma also holds for $f : \mathbb{F}_q^n \to \mathbb{F}_q$ with $\varepsilon_f \geq 16\eta/s$.

Let $\mathcal{H}$ denote the set of all hyperplanes in $\mathbb{F}^n$ and let $N = |\mathcal{H}| = (q^n - 1)/(q - 1)$. Define

$$\mathcal{H}^* = \{W \in \mathcal{H} \ : \ \mathsf{dist}(f|_W, \mathcal{C}_{n-1}) < 16\eta/s\},$$

and let $K = |\mathcal{H}^*|$. Using (10) and the induction hypothesis, we can decompose the rejection probability as follows:

$$\begin{aligned}\Pr[\mathcal{T}_k^f \text{ rejects}] &= \mathbb{E}_{W \in \mathcal{H}}\left[\Pr[\mathcal{T}_k^{f|_W} \text{ rejects}]\right] \\ &\geq \mathbb{E}_{W \in \mathcal{H} \setminus \mathcal{H}^*}\left[\Pr[\mathcal{T}_k^{f|_W} \text{ rejects}]\right] - \Pr_{W \in \mathcal{H}}[W \in \mathcal{H}^*] \geq \eta + \frac{\gamma}{q^{n-1}} - \frac{K}{N},\end{aligned}$$

where we are applying the induction hypothesis on $f|_W$ in the third transition.

Case 1: $K \leq \gamma \cdot \frac{q^n - 1}{q^n}$**.** In this case, $\frac{K}{N} \leq \gamma \cdot \frac{q-1}{q^n}$. Then, $\Pr[\mathcal{T}_k^f \text{ rejects}] \geq \eta + \frac{\gamma}{q^{n-1}} - \frac{K}{N} \geq \eta + \frac{\gamma}{q^n}$, and we are done.

Case 2: $K > \gamma \cdot \frac{q^n - 1}{q^n}$**.** In this case $K > \gamma/2 > 10^6 q^{\max\{4, c_{\mathsf{KM}}\}} \cdot s$, which means there are at least this many hyperplanes $W \in \mathcal{H}$ on which $f|_W$ is $\frac{1}{8s}$-close to degree d (by the definition of $\mathcal{H}^*$ and $\eta = 1/128$). This allows us to invoke Claim 25 and deduce that $\varepsilon_f \leq 1/(2s)$. Finally, applying Lemma 26 shows that $\Pr[\mathcal{T}_k^f \text{ rejects}] \geq \frac{s\varepsilon_f}{8} \geq \eta + \frac{\gamma}{q^n}$, where the second inequality uses the same argument as in the induction base case. ◀

6 Application to Testing with an Online-Adversary

We show that Theorem 23 lead to a tester for Reed-Muller codes in the online-adversary model. In the full version, we give a formal statement and proof of Theorem 4 and then show how it leads to a tester for lifted affine-invariant codes. Also in the full version, for more general affine-invariant families, we state and prove an analogue of Theorems 1 and 4, making the requisite modifications to handle general affine-invariant families that admit an efficient flat tester of bounded dimension. We further show that this result also gives a tester in the online adversary model. However, we are not aware of any such families for which the currently known testers have strong enough soundness for our theorem to apply.

The online-manipulation-resilient testing model was originally proposed by Kalemaj, Raskhodnikova, and Varma [19], and with extended variants was later formally defined by Ben-Eliezer, Kelman, Meir, and Raskhodnikova [5]. We follow the definitions of the latter.What matters for our discussion is that the model provides oracle access to the data via an online adversarial oracle. Specifically, given an erasure-rate parameter t, an adversary is allowed to erase (or, in a more challenging variant of the model, corrupt) up to t data points in the input data *after* each query is answered.

6.1 Optimal Online-Manipulation-Resilient Tester for Reed-Muller

By setting the dimension of the semi-sample-based tester large enough relative to the erasure rate t, we obtain testers for Reed-Muller codes in the t-online erasure model.

▶ **Theorem 28.** *Let ε be a distance parameter, for any erasure rate $t \le \varepsilon^2 q^{n-O(d)}$ and set $k_{\mathsf{adv}} = O(d) + \log_q(t/\varepsilon^2)$ sufficiently large. Then given an input function $f : \mathbb{F}^n \to \mathbb{F}$, there is a tester for $\mathsf{RM}[n,q,d]$ with the following parameters in the t-online erasure model satisfying:*

- *Completeness: If $f \in \mathsf{RM}[n,q,d]$, the tester always accepts.*
- *Soundness: If f is ε-far from $\mathsf{RM}[n,q,d]$, then $\Pr[\mathcal{T}_k^f \text{ rejects}] \ge \frac{2}{3}$.*
- *Query Complexity:*

$$Q = \max\left\{\frac{1}{\varepsilon}, 128 s_{k_{\mathsf{adv}}}\right\} = O\left(\frac{1}{\varepsilon} + q^{2d}\left(1 + O\left(\frac{\log_q(t/\varepsilon)}{d}\right)\right)^d\right)$$

Proof. The tester is the semi-sample-based tester $\mathcal{T}_k^f(Q, \mathsf{RM}[n,q,d])$ from Section 5 with $k = k_{\mathsf{adv}}$. For the online erasure model, we also make the following slight modification: if any of the queried points is erased (i.e., answered with $\perp$), the tester will stop and accept. This modification maintains completeness while degrading the soundness by an amount that we will see is negligible.

By Theorem 23, we can repeat the tester $O\left(\frac{1}{s_k \varepsilon} + 1\right)$ times and reject with probability at least $9/10$ if no erasure is ever queried, so it remains for us to upper bound the probability of querying an erasure. The total number of queries made is $Q = O(1/\varepsilon + s_k)$ queries.

It remains to bound the probability that the tester queries a point that is erased. We show that this probability is at most $1/5$ completing the proof. Indeed, since the erasure rate is t, the total number of erasure made during the run of the algorithm is at most tQ. Since every query is uniform within a k-space. the probability for each query to be erased is at most tQ/q^k, by a union bound, the probability that any of the Q queries is erased is at most

$$tQ^2/q^k = O\left(\frac{t}{q^k}\left(\frac{1}{\varepsilon^2} + s_k^2\right)\right). \tag{11}$$

Note that for $k > d$ we have $s_k \leq q^{2d}(10k/d)^d$. For large enough $k_{\mathsf{adv}} = O(d + \log_q(t/\varepsilon))$ we have that this quantity is at most,

$$O\left(\frac{t}{q^k}\left(\frac{1}{\varepsilon^2} + s_k^2\right)\right) \leq O\left(\frac{t}{q^k\varepsilon^2}\right) + O\left(\frac{t}{q^k}\left(\frac{10qk}{d}\right)^{4d}\right) < 1/5.$$

It follows that an ε-far function is accepted with probability at most $1/10 + 1/5 \leq 1/3$, as required. ◀

Using an observation by [19], we get a similar theorem for the online corruption model, wherein the adversary can change the value of f at a point rather than just erase the value at that point. Now, with a two-sided error, an error may occur whenever the tester queries a manipulated point in either case. The formal statement and its proof are provided in the full version.

7 Further Directions

Our work leaves a number of directions open for future work. We list these below.

- Improved Soundness for [22]: As mentioned in our introduction, there is a general result which shows any local characterization for an affine invariant family also yields a local tester [22]. The soundness shown for these testers is sub-optimal however, so it would be interesting to see if our techniques could be useful towards improving the soundness of the general testers of [22].
- Testing Other Affine Invariant Families: It would be interesting to see if our techniques can yield new testing results for other affine invariant families of functions, particularly the nonlinear ones considered in [8]. As mentioned, our analysis falls short in that the hyperplane agreement lemma currently relies on having an r-space tester with decent soundness, as opposed to the inexplicit soundness given in [8].

References

1 Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing Reed-Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005. doi:10.1109/TIT.2005.856958.

2 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs; A new characterization of NP. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, October 24-27, 1992*, pages 2–13. IEEE Computer Society, 1992. doi:10.1109/SFCS.1992.267824.

3 Vipul Arora, Esty Kelman, and Uri Meir. On optimal testing of linearity. In Ioana Oriana Bercea and Rasmus Pagh, editors, *2025 Symposium on Simplicity in Algorithms, SOSA 2025, New Orleans, LA, USA, January 13-15, 2025*, pages 65–76. SIAM, 2025. doi:10.1137/1.9781611978315.5.

4 Paul Beame, Shayan Oveis Gharan, and Xin Yang. On the bias of reed-muller codes over odd prime fields, 2018. arXiv:1806.06973.

5 Omri Ben-Eliezer, Esty Kelman, Uri Meir, and Sofya Raskhodnikova. Property testing with online adversaries. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ITCS.2024.11.

6 Amey Bhangale, Irit Dinur, and Inbal Livni Navon. Cube vs. cube low degree test. *CoRR*, abs/1612.07491, 2016. arXiv:1612.07491.

7 Amey Bhangale and Silas Richelson. Plane vs. plane low degree test. Technical Report 126, Electronic Colloquium on Computational Complexity (ECCC), 2025.

8 Arnab Bhattacharyya, Eldar Fischer, Hamed Hatami, Pooya Hatami, and Shachar Lovett. Every locally characterized affine-invariant property is testable. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 429–436. ACM, 2013. `doi:10.1145/2488608.2488662`.

9 Arnab Bhattacharyya, Eldar Fischer, and Shachar Lovett. Testing low complexity affine-invariant properties. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1337–1355. SIAM, 2013. `doi:10.1137/1.9781611973105.97`.

10 Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, Las Vegas, Nevada, USA, October 23-26, 2010*, pages 488–497, 2010. `doi:10.1109/FOCS.2010.54`.

11 Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993. `doi:10.1016/0022-0000(93)90044-W`.

12 Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998. `doi:10.1145/285055.285060`.

13 Oded Goldreich and Dana Ron. On sample-based testers. *ACM Trans. Comput. Theory*, 8(2), April 2016. `doi:10.1145/2898355`.

14 Alexander Golovnev, Zeyu Guo, Pooya Hatami, Satyajeet Nagargoje, and Chao Yan. Hilbert functions and low-degree randomness extractors, 2024. `doi:10.48550/arXiv.2405.10277`.

15 Alan Guo, Swastik Kopparty, and Madhu Sudan. New affine-invariant codes from lifting. In Robert D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 529–540. ACM, 2013. `doi:10.1145/2422436.2422494`.

16 Elad Haramaty, Noga Ron-Zewi, and Madhu Sudan. Absolutely sound testing of lifted codes. *Theory Comput.*, 11:299–338, 2015. `doi:10.4086/toc.2015.v011a012`.

17 Elad Haramaty, Amir Shpilka, and Madhu Sudan. Optimal testing of multivariate polynomials over small prime fields. *SIAM Journal on Computing*, 42(2):536–562, 2013. `doi:10.1137/120879257`.

18 Prahladh Harsha, Mrinal Kumar, Ramprasad Saptharishi, and Madhu Sudan. An improved line-point low-degree test. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 1883–1892. IEEE, 2024. `doi:10.1109/FOCS61266.2024.00113`.

19 Iden Kalemaj, Sofya Raskhodnikova, and Nithin Varma. Sublinear-time computation in the presence of online erasures. *Theory of Computing*, 19(1):1–48, 2023. `doi:10.4086/TOC.2023.V019A001`.

20 Tali Kaufman and Dor Minzer. Improved optimal testing results from global hypercontractivity. *SIAM Journal on Computing*, 54(3):625–663, 2025. `doi:10.1137/23M1571022`.

21 Tali Kaufman and Dana Ron. Testing polynomials over general fields. *SIAM Journal on Computing*, 36(3):779–802, 2006. `doi:10.1137/S0097539704445615`.

22 Tali Kaufman and Madhu Sudan. Algebraic property testing: the role of invariance. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 403–412, 2008. `doi:10.1145/1374376.1374434`.

23 Esty Kelman, Ephraim Linder, and Sofya Raskhodnikova. Online versus offline adversaries in property testing. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, January 7-10, 2025, Columbia University, New York, NY, USA*, volume 325 of *LIPIcs*, pages 65:1–65:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. `doi:10.4230/LIPIcs.ITCS.2025.65`.

24 Esty Kelman, Uri Meir, Debanuj Nayak, and Sofya Raskhodnikova. Homomorphism testing with resilience to online manipulations. *arXiv preprint arXiv:2511.23363*, 2025. doi:10.48550/arXiv.2511.23363.

25 Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. *Annals of Mathematics*, 198(1):1–92, 2023.

26 Dor Minzer and Kai Zheng. Approaching the soundness barrier: A near optimal analysis of the cube versus cube test. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2761–2776. SIAM, 2023. doi:10.1137/1.9781611977554.CH104.

27 Dor Minzer and Kai Zhe Zheng. Optimal testing of generalized reed-muller codes in fewer queries. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 206–233. IEEE, 2023. doi:10.1109/FOCS57990.2023.00022.

28 Dor Minzer and Kai Zhe Zheng. Adversarial low degree testing. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 4395–4409. SIAM, 2024. doi:10.1137/1.9781611977912.154.

29 Dor Minzer and Kai Zhe Zheng. Near optimal alphabet-soundness tradeoff pcps. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 15–23. ACM, 2024. doi:10.1145/3618260.3649606.

30 Dana Moshkovitz and Ran Raz. Sub-constant error low degree test of almost-linear size. *SIAM Journal on Computing*, 38(1):140–180, 2008. doi:10.1137/060656838.

31 Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 475–484. ACM, 1997. doi:10.1145/258533.258641.

32 Noga Ron-Zewi and Madhu Sudan. A New Upper Bound on the Query Complexity of Testing Generalized Reed-Muller Codes. *Theory of Computing*, 9(25):783–807, 2013. doi:10.4086/TOC.2013.V009A025.

33 Alek Westover, Edward Yu, and Kai Zheng. New direct sum tests. *CoRR*, abs/2409.10464, 2024. doi:10.48550/arXiv.2409.10464.