A New Upper Bound on the Query Complexity for Testing Generalized Reed-Muller codes

Noga Ron-Zewi[1] * and Madhu Sudan[2]

[1] Department of Computer Science, Technion, Haifa. nogaz@cs.technion.ac.il
[2] Microsoft Research New England, Cambridge, MA. madhu@mit.edu

Abstract. Over a finite field $\mathbb{F}_q$ the (n, d, q)-Reed-Muller code is the code given by evaluations of n-variate polynomials of total degree at most d on all points (of $\mathbb{F}_q^n$). The task of testing if a function $f : \mathbb{F}_q^n \to \mathbb{F}_q$ is close to a codeword of an (n, d, q)-Reed-Muller code has been of central interest in complexity theory and property testing. The query complexity of this task is the minimal number of queries that a tester can make (minimum over all testers of the maximum number of queries over all random choices) while accepting all Reed-Muller codewords and rejecting words that are δ-far from the code with probability $\Omega(\delta)$. (In this work we allow the constant in the Ω to depend on d.)

For codes over a prime field $\mathbb{F}_q$ the optimal query complexity is well-known and known to be $\Theta(q^{\lceil (d+1)/(q-1) \rceil})$, and the test consists of testing if f is a degree d polynomial on a randomly chosen $(\lceil (d+1)/(q-1) \rceil)$-dimensional affine subspace of $\mathbb{F}_q^n$. If q is not a prime, then the above quantity remains a lower bound, whereas the previously known upper bound grows to $O(q^{\lceil (d+1)/(q-q/p) \rceil})$ where p is the characteristic of the field $\mathbb{F}_q$. In this work we give a new upper bound of $(cq)^{(d+1)/q}$ on the query complexity, where c is a universal constant. Thus for every p and sufficiently large q this bound improves over the previously known bound by a polynomial factor.

In the process we also give new upper bounds on the "spanning weight" of the dual of the Reed-Muller code (which is also a Reed-Muller code). The spanning weight of a code is the smallest integer w such that codewords of Hamming weight at most w span the code. The main technical contribution of this work is the design of tests that test a function by *not* querying its value on an entire subspace of the space, but rather on a carefully chosen (algebraically nice) subset of the points from low-dimensional subspaces.

* Research conducted in part while this author was an intern at Microsoft Research New-England, Cambridge, MA, and supported in part by the Israel Ministry of Science and Technology.

1 Introduction

In this work we present new upper bounds on the query complexity of testing Reed-Muller codes, the codes obtained by evaluations of multivariate low-degree polynomials, over general fields. In the process we also give new upper bounds on the spanning weight of Reed-Muller codes. We explain these terms and our results below.

We start with the definition of Reed-Muller codes. Let $\mathbb{F}_q$ denote the finite field on q elements. Throughout we will let $q = p^s$ for prime p and integer s. The Reed-Muller codes have two parameters in addtion to the field size, namely the degree d and number of variables n. The (n, d, q)-Reed-Muller code $\mathrm{RM}[n, d, q]$ is the set of functions from $\mathbb{F}_q^n$ to $\mathbb{F}_q$ that are evaluations of n-variate polynomials of total degree at most d.

1.1 Testing Reed-Muller Codes

We define the notion of testing the "Reed-Muller" property as a special case of property testing. We let $\{\mathbb{F}_q^n \to \mathbb{F}_q\}$ denote the set of all functions mapping $\mathbb{F}_q^n$ to $\mathbb{F}_q$. A property $\mathcal{F}$ is simply a subset of such functions. For $f, g : \mathbb{F}_q^n \to \mathbb{F}_q$ we say the distance between them $\delta(f, g)$ is the fraction of points of $\mathbb{F}_q^n$ where they disagree. We let $\delta(f, \mathcal{F})$ denote the minimum distance between f and a function in $\mathcal{F}$. We say f is δ-close to $\mathcal{F}$ if $\delta(f, \mathcal{F}) \leq \delta$ and δ-far otherwise.

A (k, ϵ)-tester for the property $\mathcal{F} \subseteq \{\mathbb{F}_q^n \to \mathbb{F}_q\}$ is a randomized algorithm that makes at most k queries to an oracle for a function $f : \mathbb{F}_q^n \to \mathbb{F}_q$ and accepts if $f \in \mathcal{F}$ and rejects $f \notin \mathcal{F}$ with probability at least $\epsilon\delta(f, \mathcal{F})$.

For fixed d and q, we consider *query complexity* of testing the property of being a degree d multivariate polynomial over $\mathbb{F}_q$. Specifically, the query complexity $k = k(d, q)$, is the minimum integer such that there exists an ϵ such that for all n there is a (k, ϵ)-tester for the $\mathrm{RM}[n, d, q]$ property. (So the error ϵ of the tester is allowed to depend on q and d, but not on n.)

The query complexity of low-degree testing is a well-studied question and has played a role in many results in computational complexity including in the PCP theorem ([ALM$^+$98] and subsequent works), and in the works of Viola and Wigderson [VW08] and Barak et al. [BGH$^+$11]. Many of these results depend not only on a tight analysis of $k(d, q)$ but also a tight analysis of the parameter ϵ, but in this work we only focus on the first quantity. Below we describe what was known about these quantities.

For the case when d is (sufficiently) smaller than the field size, the works of Rubinfeld and Sudan [RS96] and Friedl and Sudan [FS95] show that $k(d, q) = d + 2$ (provided $d < q - q/p$). For the case when $q = 2$ and d is arbitrary, this quantity was analyzed in the work of Alon et al [AKK$^+$05] who show that $k(d, 2) = 2^{d+1}$ (exactly). Jutla et al [JPRZ09] and Kaufman and Ron [KR06] explored this question for general q and d (the former only considered prime q) and showed that $k(d, q) \leq q^{\lceil (d+1)/(q-q/p) \rceil}$. In [KR06] it is also shown that the bound is tight (to within a factor of q) if q is a prime. However for the non-prime case the only known lower bound on the query complexity was $k(d, q) \geq q^{(d+1)/(q-1)}$ (which is roughly the upper bound raised to the power of $(p-1)/p$). (In the following sections we describe the conceptual reason for this gap in knowledge.)

In this work we give a new upper bound on $k(d,q)$ which is closer to the lower bound when p is a constant and d and q are going to infinity. We state our main theorem below.

Theorem 1 (Main). *Let $q = p^s$ for prime p and positive integer s. Then there exists a constant $c_q \le 3q^4$ such that for every d and n, the Reed-Muller code $\mathrm{RM}[n,d,q]$ has a $(k, \Omega(1/k^2))$-local tester, for $k = k(d,q) \le c_q \cdot \left(2^{p-1}+p-1\right)^{(d+1)/(q(p-1))} q^{(d+1)/q}$. In particular $k(d,q) \le 3q^4 \cdot (3q)^{(d+1)/q}$.*

We note that when p goes to infinity the bound on $k(d,q)$ tends to $c_q \cdot (3q)^{(d+1)/q}$. We also note that the constant c_q is not optimized in our proofs and it seems quite plausible that it can be improved using more careful analysis. The more serious factor (especially when one considers a constant q and $d \to \infty$) is the constant factor multiplying q in the base of the exponent. Our techniques do seem to be unable to improve this beyond $(2^{p-1}+p-1)^{1/(p-1)}$ which is always between 2 and 3 (while the lower bounds suggest a constant which is close to 1).

We note that the above result does not compare well with previous bounds if one take the "soundness" parameter (ϵ) into account. Previous results by Bhattacharyya et al. [BKS+10] for $q = 2$ and Haramaty et al. [HSS11] for general q give a (k', ϵ_0)-local tester for ϵ_0 depending only on q (but independent of d) and $k' = q^{\lceil \frac{d+1}{q-q/p} \rceil}$. To get such a soundness independent of d, Theorem 1 yields a (k^3, ϵ_1)-local tester for ϵ_1 being some universal constant. Thus for small q and growing d this is worse than the results of [BKS+10,HSS11]. However for d and q growing at the same rate (for instance) our result does give the best bounds even if we want the soundness to be some absolute constant.

Theorem 1 is proved by proving that the Reed-Muller code $\mathrm{RM}[n,d,q]$ has a "k-single-orbit characterization" (a notion we will define later, see Definition 2 and Theorem 3). This will imply the testing result immediately by a result of Kaufman and Sudan [KS07].

1.2 Spanning weight

It is well-known (cf. [BHR05]) that the query complexity of testing a linear code C is lower bounded by the "minimum distance" of its dual, where the minimum distance of a code is the minimum weight of a non-zero codeword. (The weight of a word is simply the number of non-zero coordinates.) Applied to the Reed-Muller code $\mathrm{RM}[n,d,q]$ this suggests a lower bound via the minimum distance of its dual, which also turns out to be a Reed-Muller code. Specifically the dual of $\mathrm{RM}[n,d,q]$ is $\mathrm{RM}[n, n(q-1)-d-1, q]$. The minimum distance of the latter is well-known and is (roughly) $q^{(d+1)/(q-1)}$ and this leads to the tight analysis of the query complexity of Reed-Muller codes over prime fields.

Over non-prime fields however this bound has not been matched, so one could turn to potentially stronger lower bounds. A natural such bound would be the "spanning weight" of the dual code, namely the minimum weight w such that codewords of the dual of weight at most w span the dual code. It is easy to show that to achieve any positive ϵ (even going to 0 as $n \to \infty$) a (k, ϵ)-local tester must make at least w queries (on some random choices), where w is the spanning weight of the dual. Somewhat surprisingly, the spanning weight of the Reed-Muller code does not seem well-understood. (Some partial understanding comes from [DK00].). Since for a linear code, the spanning weight of its dual code is a lower bound on the query complexity of the code, our result gives new upper bounds on this spanning weight. Specifically, we have

Corollary 1. *Let $q = p^s$ for prime p and positive integer s. Then there exists a constant $c_q \leq 3q^4$ such that for every d and n, the Reed-Muller code* $\mathrm{RM}[n, n(q-1) - d - 1, q]$ *has a spanning weight of at most* $c_q \cdot \left(2^{p-1} + p - 1\right)^{(d+1)/(q(p-1))} \cdot q^{(d+1)/q} \leq 3q^4 \cdot (3q)^{(d+1)/q}$.

1.3 Qualitative description and techniques

Our tester differs from previous ones in some qualitative ways. All previously analyzed testers for low-degree testing roughly worked as follows: They picked a large enough dimension t (depending on q and d, but not n) and verified that the function to be tested was a degree d polynomial on a random t-dimensional affine subspace. The final aspect was verified by querying the function on the entire t-dimensional space, thus leading to a query complexity of q^t. The minimal choice of the dimension t that allows this test to detect functions that are not degree d polynomials with positive probability is termed the "testing dimension" (see, for instance, [HSS11]), and this quantity is well-understood, and equals $t_{q,d} = \lceil (d+1)/(q - q/p) \rceil$.

Any improvement to the query complexity of the test above requires two features: (1) For some choices of the tester's randomness, the set of queried points should span a $t_{q,d}$ dimensional space. (2) For all choices of the tester's randomness, it should make $o(q^{t_{q,d}})$ queries. Finding such a useful subset of $\mathbb{F}_q^n$ turns out to be a non-trivial task. The fortunate occurence that provides the basis for our tester is that such sets of points can indeed be found, and even (in retrospect) systematically.

To illustrate the central idea, consider the setting of $n = 2$, $d = q - 1$ and $q = 2^s$ for some large s. While the naive test would query the given function $f : \mathbb{F}_q^2 \to \mathbb{F}_q$ at all q^2 points, we wish to query only $O(q)$ points. Our test, for this simple setting is the following: We pick a random affine-transformation $T : \mathbb{F}_q^2 \to \mathbb{F}_q^2$ and test that the function $f \circ T$ has a zero "inner-product" with the function $g : \mathbb{F}_q^2 \to \mathbb{F}_q$ given by $g(x, y) = \frac{1}{y}((x + y)^{q-1} - x^{q-1})$. Here "inner-product" is simply the quantity $\sum_{\alpha,\beta \in \mathbb{F}_q} (f \circ T)(\alpha, \beta) g(\alpha, \beta)$. It can be verified that the function g is zero very often and indeed takes on non-zero values on at most $3q = O(q)$ points in $\mathbb{F}_q^2$. So querying $f(\alpha, \beta)$ at these $O(q)$ points suffices. The more interesting question is: Why is this test complete and sound?

Completeness is also easy to verify. It can be verified, by some simple manipulations that any monomial of the form $x^i y^j$ with $i + j < q$ has a zero inner product with g and by linearity of the test it follows that all polynomials of total degree at most d have a zero inner product with g. Since the degree of functions is preserved under affine-transformations, it then follows that $f \circ T$ also has zero inner product with g for every polynomial f of total degree at most d.

Finally, we turn to the soundness. Here we appeal to the emerging body of work on affine-invariant linear properties (linear properties that are preserved under affine-transformations), which allows us to focus on very specific monomials and to verify that their inner product with g is non-zero. In particular, we use a "monomial extraction" lemma (from [KS07]) which allows us to focus on the behavior of our tests only on monomials, as opposed to general polynomials. Further the theory also allows us to focus on specific monomials due to a "monomial spread" lemma which we use to prove that every affine-invariant family which contains some monomials of degree greater than d also contains some canonical monomials of degree slightly larger than d. In the special case of polynomials of degree at most $q-1$, these lemmas allow us to focus on only bivariate monomials of degree q, namely the monomials $x^i y^{q-i}$ for $1 \leq i \leq q-1$ and for these monomials one can again verify that their inner product with g is non-zero. Using the general methods in the theory of affine-invariant property testing, one can conclude that all polynomials of degree greater than d are rejected with positive probability.

Extending the above result to the general case turns out relatively clean, again using methods from the study of testing of affine-invariant linear properties. The extension to general n is immediate. Extending to other degrees involves some intuitive ways of combining tests, with analysis that get simplified by the emerging theory. These combinations yield the query complexity of roughly $(3q)^{(d+1)/q}$. We however attempt to reduce the constant in front of q in the base of this expression and manage to get an expression that tends to 2 when p goes to infinity. In order to do so we abstract the function g as being the derivative of the function x^{q-1} in direction y, and extend it to use iterative derivatives. This yields the best tests we give in the paper.

Organization In Section 2 we introduce some of the standard background material from the study of affine-invariant linear properties and use the theory to provide restatements of our problem. In Section 3 we introduce the main novelty of our work, which provides a restricted version of our test while achieving significant savings over standard tests. In the full version of this paper [RS12] we show how to build on the test from Section 3 to get a tester for the general case. Due to space limitations many proofs are omitted from this version. They may also be found in the full version.

2 Background and restatement of problem

We start by introducing some of the background material that leads to some reformulations of the main theorem we wish to prove. We first introduce the notions of "constraints" and "(single-orbit) characterizations", which leads to a first reformulation of our main theorem (see Theorem 3). We then give some sufficient conditions to recognize such characterizations, and this leads to a second reformulation of our main theorem (see Theorem 4).

2.1 Single-orbit characterizations

In this section we use the fact that Reed-Muller codes form a "linear, affine-invariant property". We recall these notions first. Given a finite field $\mathbb{F}_q$ a property is a set of functions $\mathcal{F}$ mapping $\mathbb{F}_q^n$ to $\mathbb{F}_q$. The property is said to be *linear* if it is an $\mathbb{F}_q$-vector space, i.e., $\forall f, g \in \mathcal{F}$ and $\alpha \in \mathbb{F}_q$ we have $\alpha f + g \in \mathcal{F}$. The property is said to be *affine-invariant* if it is invariant under affine-transformations of the domain, i.e., $\forall f \in \mathcal{F}$ it is the case that $f \circ T$ is also in $\mathcal{F}$ for every affine-transformation $T : \mathbb{F}_q^n \to \mathbb{F}_q^n$ given by $T(x) = A \cdot x + \beta$ for $A \in \mathbb{F}_q^{n \times n}$, $\beta \in \mathbb{F}_q^n$. [3] It can be easily verified that $\mathrm{RM}[n, d, q]$ is linear and affine-invariant for every n, d, q.

The main tool used so far for constructing testers for affine-invarinat linear properties is a structural theorem which shows that every linear affine-invariant property that is k-single characterizable is also k-locally testable. In order to describe the notion of single-orbit characterizability we start with a couple of definitions.

Definition 1 (k-constraint, k-characterization). *A* k-constraint $C = (\overline{\alpha}, \{\overline{\lambda}_i\}_{i=1}^r)$ *on* $\{\mathbb{F}_q^n \to \mathbb{F}_q\}$ *is given by a vector* $\overline{\alpha} = (\alpha_1, \ldots, \alpha_k) \in (\mathbb{F}_q^n)^k$ *together with* r *vectors* $\overline{\lambda}_i = (\lambda_{i,1}, \ldots \lambda_{i,k}) \in \mathbb{F}_q^k$ *for* $1 \leq i \leq r$. *We say that the constraint* C accepts *a function* $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ *if* $\sum_{j=1}^k \lambda_{i,j} f(\alpha_j) = 0$ *for all* $1 \leq i \leq r$. *Otherwise we say that* C rejects f.

Let $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ *be a linear property. A* k-characterization *of* $\mathcal{F}$ *is a collection of* k*-constraints* $C_1, \ldots, C_m$ *on* $\{\mathbb{F}_q^n \to \mathbb{F}_q\}$ *such that* $f \in \mathcal{F}$ *if and only if* C_j *accepts* f*, for every* $j \in \{1, \ldots, m\}$.

It is well-known [BHR05] that every k-locally testable linear property must have a k-characterization. In the case of affine-invariant linear families some special characterizations are known to lead to k-testability. We describe these special characterizations next.

Definition 2 (k-single-orbit characterization). *Let* $C = \left(\overline{\alpha}, \{\overline{\lambda}_i\}_{i=1}^r\right)$ *be a* k*-constraint on* $\{\mathbb{F}_q^n \to \mathbb{F}_q\}$. *The* orbit *of* C *under the set of affine-transformations is the set of* k*-constraints* $\{T \circ C\}_T = \left\{\left(\left(T(\alpha_1), \ldots, T(\alpha_k)\right), \{\overline{\lambda}_i\}_{i=1}^r\right) \mid T : \mathbb{F}_q^n \to \mathbb{F}_q^n \text{ is an affine-transformation}\right\}$. *We say that* C *is a* k-single-orbit characterization *of* $\mathcal{F}$ *if the orbit of* C *forms a* k*-characterization of* $\mathcal{F}$.

The following theorem, due to Kaufman and Sudan [KS07], says that k-single-orbit characterization implies local testability.

[3] We note that as in [KS07] we do not require A to be non-singular. Thus the affine-transformations we consider are not necessarily permutations from $\mathbb{F}_q^n$ to $\mathbb{F}_q^n$.

Theorem 2 (Single-orbit characterizability implies local testability, [KS07, Lemma 2.9]). *Let $\mathcal{F} \subseteq \left\{\mathbb{F}_q^n \to \mathbb{F}_q\right\}$ be an affine-invariant linear family. If $\mathcal{F}$ has a k-single-orbit characterization, then $\mathcal{F}$ has a $(k, \Omega(1/k^2))$-local tester.*

In view of the above theorem, it suffices to find a single-orbit characterization of $\mathrm{RM}[n, d, q]$ to test it. The following theorem, which we prove in the rest of this paper, thus immediately implies Theorem 1.

Theorem 3. *Let $q = p^s$ for prime p, and let n, d be arbitrary positive integers. Then the Reed-Muller code $\mathrm{RM}[n, d, q]$ has a k-single-orbit characterization for $k \leq c_q \cdot \left(2^{p-1} + p - 1\right)^{(d+1)/(q(p-1))} \cdot q^{(d+1)/q}$ where $c_q \leq 3q^4$.*

2.2 Constraints vs. Monomials

One of the main simplifications derived from the study of affine-invariant linear properties is that it suffices to analyze the performance of constraints on "monomials" as opposed to general polynomials. This allows us to rephrase our target (a single-orbit characterization of $\mathrm{RM}[n, d, q]$) in somewhat simpler terms. Below we describe some of the essential notions, namely the "degree set", the "border set" and the relationship of these to single-orbit characterizations. This leads to a further reformulation of our main theorem as Theorem 4. Variations of most of the results and notions presented in this section appeared in previous works [KS07,GKS09,BS11,BGM$^+$11]. In all the above works, with the exception of [KS07], the notions were specialized to the case of univariate funcions mapping $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$ that are invariant over the set of affine-transformations over $\mathbb{F}_{q^n}$. In this work we focus on these notions in the context of affine-invariant linear properties over the domain $\mathbb{F}_q^n$.

Let $\mathcal{F} \subseteq \{\mathbb{F}_q^n \to \mathbb{F}_q\}$ be a linear affine-invariant family of functions. Note that every member of $\{\mathbb{F}_q^n \to \mathbb{F}_q\}$ can be written uniquely as a polynomial in $\mathbb{F}_q[x_1, x_2, \ldots, x_n]$ of degree at most $q - 1$ in each variable. For a monomial $\prod_{i=1}^n x_i^{d_i}$ over n variables, we define its degree to be the vector $\overline{d} = (d_1, d_2, \ldots, d_n)$ and we define its *total degree* to be $\sum_{i=1}^n d_i$. For a function $f : \mathbb{F}_q^n \to \mathbb{F}_q$ we denote its *support*, denoted $\mathrm{supp}(f)$, to be the set degrees in the support of the associated polynomial. I.e., $\mathrm{supp}(f) = \{\overline{d} \in \{0, \ldots, q-1\}^n | c_{\overline{d}} \neq 0\}$ where $f(x) = \sum_{\overline{d}} c_{\overline{d}} x^{\overline{d}}$. The *degree set* $\mathrm{Deg}(\mathcal{F})$ of $\mathcal{F}$ is simply the union of the supports of the functions in $\mathcal{F}$, i.e., $\mathrm{Deg}(\mathcal{F}) = \cup_{f \in \mathcal{F}} \mathrm{supp}(f)$.

While the degree set of the Reed-Muller codes are natural to study, they are also natural in more general contexts. The following lemma from [KS07] says that every affine-invariant linear property from $\mathbb{F}_q^n$ to $\mathbb{F}_q$ is uniquely determined by its degree set.

Lemma 1 (Monomial extraction lemma, [KS07, Lemma 4.2]). *Let $\mathcal{F} \subseteq \left\{\mathbb{F}_q^n \to \mathbb{F}_q\right\}$ be an affine-invariant linear property. Then $\mathcal{F}$ has a monomial basis,* *that is, $\mathcal{F}$ is the set of all polynomials supported on monomials of the form $x^{\overline{d}}$ where $\overline{d} \in \mathrm{Deg}(\mathcal{F})$.* [4]

One main structural feature of the degree sets of affine-invariant linear properties is that they are *p-shadow-closed*. Before giving the definition of a shadow-closed set of degrees we need to introduce a bit of notation. For a pair of integers a, b let $a = \sum_j a_j p^j$, $b = \sum_j b_j p^j$ be their base-p representation, respectively. We say that b is in the *p-shadow* of a, and denote this $b \leq_p a$, if $b_j \leq a_j$ for all j. For a pair of integer vectors $\overline{d} = (d_1, d_2, \ldots, d_n)$, $\overline{e} = (e_1, e_2, \ldots, e_n)$ we say that $\overline{e} \leq_p \overline{d}$ if $e_i \leq_p d_i$ for every i.

Definition 3 (Shadow-closed set of degrees). *For a vector of integers $\overline{d} = (d_1, d_2, \ldots, d_n)$ of length n, the* p-shadow *of $\overline{d}$ is the set* $\mathrm{Shadow}_p(\overline{d}) = \{\overline{e} = (e_1, e_2, \ldots, e_n) \mid \overline{e} \leq_p \overline{d}\}$. *For a subset S of integer vectors of length n we let* $\mathrm{Shadow}_p(S) = \bigcup_{\overline{d} \in S} \mathrm{Shadow}_p(\overline{d})$. *Finally, we say that S is* p-Shadow-closed *if* $\mathrm{Shadow}_p(S) = S$.

Lemma 4.6 in [KS07] says that the degree set of every affine-invariant linear property over $\mathbb{F}_q^n$ is p-shadow-closed. This motivates the notion of a "border" set, the set of minimal elements (under $\leq_p$) that are not in $\mathrm{Deg}(\mathcal{F})$.

Definition 4 (Border). *For an affine-invariant linear family $\mathcal{F} \subseteq \{\mathbb{F}_q^n \to \mathbb{F}_q\}$, its* border set*, denoted* $\mathrm{Border}(\mathcal{F})$*, is the set*

$$Border(\mathcal{F}) = \{\overline{e} \in \{0, \ldots, q-1\}^n | \overline{e} \notin \mathrm{Deg}(\mathcal{F}) \text{ but } \forall \overline{e}' \leq_p \overline{e}, \overline{e}' \neq \overline{e}, \overline{e}' \in \mathrm{Deg}(\mathcal{F})\}.$$

The relationship between the degree set and the border set of an affine-invariant linear family and single-orbit characterizability is given by the following lemma. This lemma says that for an affine-invariant linear family, in order to establish k-single-orbit characterizability it suffices to exhibit a k-constraint whose orbit accepts all monomials of the form $x^{\overline{d}}$ for $\overline{d} \in \mathrm{Deg}(\mathcal{F})$ and rejects all monomials of the form $x^{\overline{b}}$ for $\overline{b} \in \mathrm{Border}(\mathcal{F})$. It is similar in spirit to Lemma 3.2 of [BGM+11] which shows that a similar result holds for affine-invariant linear properties over $\mathbb{F}_{q^n}$.

Lemma 2. *Let $\mathcal{F} \subseteq \{\mathbb{F}_q^n \to \mathbb{F}_q\}$ be an affine-invariant linear property and let C be a constraint. Then C is a single-orbit characterization of $\mathcal{F}$ if the orbit of C accepts every monomial $x^{\overline{d}}$ for $\overline{d} \in \mathrm{Deg}(\mathcal{F})$ and rejects every monomial $x^{\overline{b}}$ for $\overline{b} \in \mathrm{Border}(\mathcal{F})$.*

Proof omitted in this version.

In order to describe the border of the Reed-Muller family we shall use the following definition.

[4] Our language is somewhat different from that of [KS07]. After translation, their lemma says that all monomials $x^{\overline{d}}$ are contained in $\mathcal{F}$. The other direction saying $\mathcal{F}$ is contained in the span of such monomials is immediate from the definition of $\mathrm{Deg}(\mathcal{F})$.

Definition 5. *For integer d, let $d_0, d_1, \ldots,$ be its expansion in base-p, i.e., d_j's satisfy $0 \leq d_j < p$ and $d = \sum_{j=0}^{\infty} d_j p^j$. Let $b_i(d) = p^i + \sum_{j=i}^{\infty} d_j p^j$.*

Note that $b_i(d) > d$ for every i and conversely, for every integer $e > d$ there exists an i such that $b_i(d) \leq_p e$. The $b_i(d)$'s are useful in describing the border monomials of the Reed-Muller family, as formalized below.

Proposition 1. *For every n, d, q, where $q = p^s$ for a prime p, we have*

$$\mathrm{Deg}(\mathrm{RM}[n,d,q]) = \left\{ \overline{d} = (d_1, \ldots, d_n) \in \{0, \ldots, q-1\}^n \mid \sum_{j=1}^{n} d_j \leq d \right\} \text{ and}$$

$$\mathrm{Border}(\mathrm{RM}[n,d,q]) \subseteq$$

$$\left\{ \overline{e} = (e_1, \ldots, e_n) \in \{0, \ldots, q-1\}^n \mid \sum_{j=1}^{n} e_j = b_i(d) \text{ for some } 0 \leq i \leq s \right\}.$$

Proof omitted.

Combining Lemma 2 and Proposition 1 we have that Theorem 3 follows immediately from Theorem 4 below.

Theorem 4. *Let $q = p^s$ for a prime p. Then there exists a k-constraint C whose orbit accepts all monomials of total degree at most d and rejects all monomials of total degree $b_i(d)$ for $0 \leq i \leq s$, for $k \leq 3q^4 \cdot \left(2^{p-1} + p - 1\right)^{(d+1)/(q(p-1))} \cdot q^{(d+1)/q}$.*

The rest of this paper will be devoted to proving Theorem 4.

3 Canonical monomials and a new constraint

In this section we introduce the notion of "canonical monomials" of a given degree — very simplified monomials that appear in every affine-invariant linear property containing monomials of a given degree. We then give a constraint that rejects canonical monomials of some special degrees, while accepting all monomials of lower degrees. In the full version of this paper [RS12], we show how to use this to build a constraint whose orbit accepts all monomials of total degree at most d while rejecting all monomials of total degree $b_i(d)$, which suffices to get Theorem 4.

Definition 6 (Canonical monomials). *Let $q = p^s$ for a prime p. The canonical monomial of (total) degree d over $\mathbb{F}_q$ is the monomial $\prod_{i=1}^{\ell} x_i^{d_i}$ which satisfies $\sum_{i=1}^{\ell} d_i = d$, $d_i = q - q/p$ for all $2 \leq i \leq \ell$, $0 \leq d_1 \leq q-1$ and $d_1 + q - q/p > q - 1$.*

We note that [HSS11] used a different canonical monomial (cf. Definition 4.1., [HSS11]) for the construction of their improved tester for the Reed-Muller codes. Our different choice of canonical monomial is needed to construct single-orbit characterizations which improve on those given in [HSS11] in terms of the number of queries. The main property of the canonical monomial, that we will use in the full version of tihs paper to prove Theorem 4 is that every affine-invariant linear family that contains any monomial of total degree d also contains the canonical monomial of degree d. This will imply in turn that if we can find constraints that *reject* this canonical monomial their orbit will reject every monomial of total degree d.

3.1 A new constraint on monomials of total degree $< p(q - q/p)$

The main technical novelty in our paper is a k-constraint C that accepts all monomials of total degree strictly less than $p(q - q/p)$ in p variables but rejects the canonical monomial of degree $p(q-q/p)$ (note that the latter monomial also has p variables) for $k = (2^{p-1}+p-1)q^{p-1}$. We state the lemma below and devote the rest of this section to proving this lemma.

Lemma 3 (Main technical lemma). *For every q which is a power of a prime p there exists a k-constraint C which accepts all monomials of total degree smaller than $p(q-q/p)$ in p variables and rejects the canonical monomial (in p variables) of degree $p(q - q/p)$ over $\mathbb{F}_q$, where $k = (2^{p-1} + p - 1)q^{p-1}$.*

It will be convenient for us to represent the constraint C as a p-variate polynomial over $\mathbb{F}_q$. More precisely, suppose that $g(x)$ is a p-variate polynomial $g(x) \in \mathbb{F}_q[x_1, x_2, \ldots, x_p]$ that is non-zero on at most k points in $\mathbb{F}_q^p$. We associate with $g(x)$ the k-constraint $C = (\overline{\alpha}, \overline{\lambda})$, $\overline{\alpha} = (\alpha_1, \ldots, \alpha_k) \in (\mathbb{F}_q^p)^k$, $\overline{\lambda} = (\lambda_1, \ldots, \lambda_k) \in \mathbb{F}_q^k$, where the vector $\overline{\alpha}$ consists of all points in $\mathbb{F}_q^p$ on which $g(x)$ is non-zero and $\lambda_j = g(\alpha_j)$ for all $1 \leq j \leq k$. Clearly, for every function $f : \mathbb{F}_q^p \to \mathbb{F}_q$ it holds that

$$\sum_{j=1}^{k} \lambda_j f(\alpha_j) = \sum_{\beta_1,\ldots,\beta_p \in \mathbb{F}_q} g(\beta_1, \ldots, \beta_p) \cdot f(\beta_1, \ldots, \beta_p) \qquad (1)$$

Thus we reduce the task of finding a k-constraint which accepts all monomials of total degree smaller than $p(q - q/p)$ and rejects the canonical monomial of degree $p(q - q/p)$ to the task of finding a p-variate polynomial $g(x) \in \mathbb{F}_q[x_1, x_2, \ldots, x_p]$ with at most k non-zero points in $\mathbb{F}_q^p$ such that $\sum_{\beta_1,\ldots,\beta_p \in \mathbb{F}_q} g(\beta_1, \ldots, \beta_p) \cdot M(\beta_1, \ldots, \beta_p) = 0$ for every monomial in p variables of total degree smaller than $p(q - q/p)$ and $\sum_{\beta_1,\ldots,\beta_p \in \mathbb{F}_q} g(\beta_1, \ldots, \beta_p) \cdot M(\beta_1, \ldots, \beta_p) \neq 0$ when $M(x)$ is the canonical monomial of degree $p(q - q/p)$.

We start by describing a polynomial $P(x)$ that will satisfy the conditions we expect in g above. The best way to describe this polynomial is via the notion of *directional derivatives*. Let $f : \mathbb{F}_q \to \mathbb{F}_q$ be a function. Define the derivative of f in direction $y \in \mathbb{F}_q$ as $f_y(x) = f(x + y) - f(x)$. Define the iterated derivatives as

$$f_{y_1,\ldots,y_d}(x) = (f_{y_1,\ldots,y_{d-1}})_{y_d}(x) = \sum_{I \subseteq [d]} (-1)^{|I|+1} f\left(x + \sum_{i \in I} y_i\right).$$

Let $f(x)$ be the polynomial $f(x) = x_p^{q-1}$. Our polynomial $P(x)$ will be defined as follows.

$$P(x) = \frac{f_{x_1,\ldots,x_{p-1}}(x_p)}{x_1 \cdots x_{p-1}} = \frac{\sum_{I \subseteq [p-1]} (-1)^{|I|+1} (x_p + \sum_{i \in I} x_i)^{q-1}}{x_1 \cdots x_{p-1}}. \tag{2}$$

To see that $P(x)$ is indeed a polynomial we need to show that $f_{x_1,\ldots,x_{p-1}}(x_p)$ is divisible by $x_1 \cdots x_{p-1}$. We omit the proof here.

In order to prove our main technical Lemma 3 it suffices to show that the number of non-zero points of $P(x)$ in $\mathbb{F}_q^p$ is at most $(2^{p-1}+p-1)q^{p-1}$, that it accepts all monomials in p variables of total degree smaller $p(q-q/p)$, and that it rejects the canonical monomial of degree $p(q-q/p)$. We assert these three claims in Lemmas 4, 5 and 6 below, respectively. Given these three lemmas our main technical Lemma 3 is immediate. We start with bounding the number of non-zeros of $P(x)$.

Lemma 4. *The number of non-zero points of $P(x)$ in $\mathbb{F}_q^p$ is at most $(2^{p-1}+p-1)q^{p-1}$.*

Lemma 5. *Let C be the constraint associated with $P(x)$. Then C accepts all monomials in p variables of total degree smaller than $p(q-q/p)$.*

Lemma 6. *Let C be the constraint associated with $P(x)$. Then C rejects the canonical monomial of degree $p(q-q/p)$ over $\mathbb{F}_q$.*

Given Lemmas 4, 5 and 6 the proof of Lemma 3 is immediate.

Proof (Proof of Lemma 3). Let $P(x)$ be the polynomial given in (2), and let C be the constraint on $\left\{\mathbb{F}_q^p \to \mathbb{F}_q\right\}$ associated with $P(x)$. From Lemma 4 we have that the number of non-zero points of $P(x)$ in $\mathbb{F}_q^p$ is at most $(2^{p-1}+p-1)q^{p-1}$, and hence C is a $\left((2^{p-1}+p-1)q^{p-1}\right)$-constraint. Lemma 5 implies that C accepts all monomials of total degree smaller than $p(q-q/p)$, while Lemma 6 implies that C rejects the canonical monomial of degree $p(q-q/p)$.

Acknowledgements

We would like to thank Amir Shpilka for suggesting that our tests are related to directional derivatives.

References

[AKK+05] 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.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, May 1998.

[BGH$^+$11] Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. Making the long code shorter, with applications to the unique games conjecture. *CoRR*, abs/1111.0405, 2011.

[BGM$^+$11] Eli Ben-Sasson, Elena Grigorescu, Ghid Maatouk, Amir Shpilka, and Madhu Sudan. On sums of locally testable affine invariant properties. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, volume 6845 of LNCS*, pages 400–411, 2011.

[BHR05] Eli Ben-Sasson, Prahladh Harsha, and Sofya Raskhodnikova. Some 3CNF properties are hard to test. *SICOMP: SIAM Journal on Computing*, 35, 2005.

[BKS$^+$10] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *FOCS*, pages 488–497. IEEE Computer Society, 2010.

[BS11] Eli Ben-Sasson and Madhu Sudan. Limits on the rate of locally testable affine-invariant codes. In Leslie Ann Goldberg, Klaus Jansen, R. Ravi, and José D. P. Rolim, editors, *APPROX-RANDOM*, volume 6845 of *Lecture Notes in Computer Science*, pages 412–423. Springer, 2011.

[DK00] Peng Ding and Jennifer D. Key. Minimum-weight codewords as generators of generalized Reed-Muller codes. *IEEE Transactions on Information Theory*, 46(6):2152–2158, 2000.

[FS95] Katalin Friedl and Madhu Sudan. Some improvements to total degree tests. In *Proceedings of the 3rd Annual Israel Symposium on Theory of Computing and Systems*, pages 190–198, Washington, DC, USA, 4-6 January 1995. IEEE Computer Society. Corrected version available online at `http://people.csail.mit.edu/madhu/papers/friedl.ps`.

[GKS09] Elena Grigorescu, Tali Kaufman, and Madhu Sudan. Succint representation of codes with applications to testing. In *Proceedings of RANDOM-APPROX 2009*, volume 5687 of *Lecture Notes in Computer Science*, pages 534–547. Springer, 2009.

[HSS11] Elad Haramaty, Amir Shpilka, and Madhu Sudan. Optimal testing of multivariate polynomials over small prime fields. In Rafail Ostrovsky, editor, *FOCS*, pages 629–637. IEEE, 2011.

[JPRZ09] Charanjit S. Jutla, Anindya C. Patthak, Atri Rudra, and David Zuckerman. Testing low-degree polynomials over prime fields. *Random Struct. Algorithms*, 35(2):163–193, 2009.

[KR06] Tali Kaufman and Dana Ron. Testing polynomials over general fields. *SIAM Journal of Computing*, 36(3):779–802, 2006.

[KS07] Tali Kaufman and Madhu Sudan. Algebraic property testing: The role of invariance. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(111), 2007.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, April 1996.

[RS12] Noga Ron-Zewi and Madhu Sudan. A new upper bound on the query complexity for testing generalized Reed-Muller codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:46, 2012.

[VW08] Emanuele Viola and Avi Wigderson. Norms, xor lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(1):137–168, 2008.