Electronic Colloquium on Computational Complexity, Report No. 108 (2010)

Limits on the rate of locally testable affine-invariant codes

Eli Ben-Sasson*
Department of Computer Science
Technion — Israel Institute of Technology
Haifa, Israel
eli@cs.technion.ac.il

Madhu Sudan
Microsoft Research
Cambridge, Massachusetts, USA
madhu@microsoft.com

July 9, 2010

Abstract

A linear code is said to be affine-invariant if the coordinates of the code can be viewed as a vector space and the code is invariant under an affine transformation of the coordinates. A code is said to be locally testable if proximity of a received word to the code can be tested by querying the received word in a few coordinates. Locally testable codes have played a critical role in the construction of probabilistically checkable proofs and most such codes originate from simple affine invariant codes (in particular the Reed-Muller codes). Furthermore it turns out that the local testability of these codes can really be attributed to their affine-invariance. It was hoped that by studying broader classes of affine-invariant codes, one may find nicer, or simpler, locally testable codes, and in particular improve (significantly) on the rate achieved by Reed-Muller codes.

In this work we show that low-rate is an inherent limitation of affine-invariant codes. We show that any k-query affine-invariant binary code is contained in an 2^k-query testable Reed-Muller code. In fact our result shows that any affine-invariant code that has a k-local constraint (i.e., a weight k codeword in its dual), a necessary condition for k-query local testability, is contained in a Reed-Muller code that is 2^k–locally characterized (i.e., its dual is spanned by words of weight at most 2^k).

While the structure of affine-invariant codes has been explored quite extensively in the recent past Kaufman and Sudan [2008], Grigorescu et al. [2008, 2009], relating the locality of constraints in such codes to the rate has been a non-trivial challenge. Such questions lead to the task of showing that a class of systems of multivariate polynomial equations have no common zeroes over finite fields. Our result effectively shows a certain class of such systems (subclasses of diagonal systems) that have no common zeroes.

*The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258 and from the the US-Israel Binational Science Foundation under grant number 2006104.

ISSN 1433-8092

1 Introduction

In this work we consider an interesting subclass of "locally testable codes", namely "affine-invariant" codes, and prove upper bounds (limitations) on their rate. In the process we also introduce techniques of relevance to "algebraic property testing".

1.1 Locally testable codes, affine-invariance, and main result

Locally testable codes are error-correcting codes for whom membership can be tested extremely efficiently, probabilistically. Specifically, a linear code $\mathcal{C} \subseteq \Sigma^N$ is k-locally testable if there exists an algorithm T that accesses a word $w \in \Sigma^N$ as an oracle, queries the value of w on k coordinates, and accepts with probability one if $w \in \mathcal{C}$ and rejects with constant probability if w is "far" from all codewords of $\mathcal{C}$. ("Far" here refers to the relativized Hamming distance between words.)

Locally testable codes have implicitly been a subject of active study ever since the work of Blum, Luby, and Rubinfeld [1990] that showed that (effectively) the Hadamard code is 3-locally testable. They play a major role in the construction of PCPs [Arora and Safra, 1998, Arora et al., 1998] from the early days of this theorem and continuing through the recent work of Dinur [2007]. Their first systematic investigation started in [Goldreich and Sudan, 2006] and yet most basic questions about their limits remain unanswered (e.g., is there an asymptotically good family of locally testable codes?).

A particularly interesting class of locally testable codes are the affine-invariant ones. Here the code is a linear code over some finite field $\mathbb{F}$ and the coordinates of the code are themselves vector spaces over some finite extension field $\mathbb{K}$ of $\mathbb{F}$. Thus words in such codes can be viewed as functions from $\mathbb{K}^m$ to $\mathbb{F}$ and the code is a subfamily of such functions. The code is said to be affine invariant if it is invariant under affine-transformations of the domain. Specifically if $A : \mathbb{K}^m \to \mathbb{K}^m$ is an affine transformation and $f : \mathbb{K}^m \to \mathbb{F}$ is a function in $\mathcal{C}$, then so is $f \circ A$ where $f \circ A(x) = f(A(x))$.

Affine-invariant codes form a natural class of algebraic codes which have natural local tests under minimal conditions. Specifically, it is well known that for a linear code to be testable it must have a low weight codeword in its dual, or equivalently a local "constraint" (see, for instance, Ben-Sasson et al. [2005]). In the notation used above for affine invariant codes, a k-local "constraint" is a collection of points $\alpha_1, \ldots, \alpha_k \in \mathbb{K}^m$ and values $\lambda_1, \ldots, \lambda_k \in \mathbb{F}$ such that for every function $f \in \mathcal{C}$, it is the case that $\sum_{i=1}^k \lambda_i f(\alpha_i) = 0$. For affine-invariant codes the presence of one local constraint immediately implies many local constraints by affine "rotations": For every affine map A, the set of points $A(\alpha_1), \ldots, A(\alpha_k)$ also define a constraint on $\mathcal{C}$. This abundance of constraints leads to the hope that affine-invariant codes may be locally testable, and indeed Kaufman and Sudan [2008] show that if the code is characterized by the set of constraints derived from the affine rotations of a single constraint, then it is also locally testable (by the natural test).

We point out that it is the *abundance* of local constraints, not their mere existence, that seems to be essential for obtaining locally testable codes. In extreme cases where there is no abundance of local constraints, such as for low-density-parity-check (LDPC) codes based on random expanders, or for codes that have the very minimal number of local constraints needed to characterize them, there cannot be any hope for local testability [Ben-Sasson et al., 2005, 2009]. But, all things considered, abundance of local constraints should reduce the rate of the code, unless the constraints are carefully chosen in an algebraically consistent way. The class of affine invariant codes offered a promising approach to balance the need for abundance of local constraints with maintaining high rate.

This leads to the question: Which affine invariant codes have local constraints (and characterizations), and in particular how local can the constraints be, given other parameters of the code, most notably, its rate.

One, somewhat optimistic hope, was that affine-invariance might lead to simpler constructions of codes matching the best known parameters (the current constructions are immensely complicated [Ben-Sasson and Sudan, 2005, Dinur, 2007]), or even improve on them, since the scope is wider than just the class of Reed-Muller codes. This question however, resisted attacks till now, since the question of determining when a low-weight constraint can exist in an affine-invariant code leads to questions about the zeroes of certain systems of multivariate polynomial equations and these are challenging to analyze.

Here we take some first steps in this direction, though unfortunately to give a negative answer to the optimistic hope above. Specifically, we give a full analysis of a certain class of polynomial equations that arise in this setting to get a rate upper bound on affine invariant codes. For simplicity of exposition we describe our result for the case of prime fields $\mathbb{F} = \mathbb{F}_p$. The statement for the case fields of size $p^r, r > 1$ is somewhat more technical but the rate bounds we get for this case are similar to that of prime fields (cf. Theorem 2.9 and Corollary 2.10). Our main theorem (Theorem 2.9) shows that if $\mathbb{K}$ is an extension field of $\mathbb{F}$ and $\mathcal{C}$ is a k-locally testable code, then $\mathcal{C}$ is contained in a p^k-locally testable Reed-Muller code. If k and p are constants (which is the desired setting of parameters) then it says that going to general affine-invariance only buys (at best) a constant difference in the locality, when compared to the Reed-Muller codes. Since Reed-Muller codes with constant locality over constant field sizes are known to have exponentially low-rate, this rules out the hope described in the previous paragraph, by a long margin.

Notice there is an exponential gap between the query complexity of affine-invariant codes with a k-local constraint and the query complexity of the Reed-Muller code which we show contains them, which is p^k. Getting a full characterization of affine-invariant codes with a k-local constraint, even over specific fields (like $\mathbb{F}_{2^n}$ for prime n, a field which contains no subfields other than $\mathbb{F}_2$) seems to us like an interesting question for future research.

1.2 Algebraic property testing

Property testing considers the task of testing if a function f from a large domain D to a small range R satisfies some given property, where the property itself is given by the set of functions $\mathcal{F} \subseteq \{g : D \to R\}$ that satisfy the property. Again the interest here is in "quick and dirty" tests, i.e., probabilistic tests that query the given function f on few inputs, and accept if $f \in \mathcal{F}$ and reject with constant probability if f is far from $\mathcal{F}$. (Note that a locally testable code is just property testing where we view the set of functions $\mathcal{F}$ as an error-correcting code.)

Property testing also emerged in the work of Blum et al. [1990], was formally defined by Rubinfeld and Sudan [1996], and was systematically explored (in particular in non-algebraic contexts) by Goldreich et al. [1998]. Subsequently the study of combinatorial property testing, and in particular, graph property testing has developed into a rich study and by now we have almost complete understanding (at least in the dense-graph model) of which graph properties are locally testable [Alon et al., 2006, Borgs et al., 2006].

In contrast algebraic properties have not been understood as well, despite the overwhelming applications in complexity, and indeed till recently even an understanding of what makes a property algebraic was missing. The concept of affine-invariance was introduced by Kaufman and Sudan [2008] to propose such a notion, and when the domain is a vector space over a small field $\mathbb{K}$ (of constant size) they manage to characterize locally testable properties completely. Such codes are constant-locally testable if and only if they admit a constant local constraint, and the size of the constraint can be related loosely to the highest degree of polynomials in the family.

This naturally leads to the question: What about affine invariant codes over large fields. In particular for the extreme (and essentially most general) case when $m = 1$ and the functions of interest map $\mathbb{K}$ to a prime subfield $\mathbb{F}_p$, there was no interesting relationships known between the degrees of the functions in the family and the locality of the test. And such understanding is essential to get a characterization of affine-invariant locally testable codes that would be analogous to the characterizations of graph properties of [Alon et al., 2006, Borgs et al., 2006].

Our work takes a step in this direction by giving non-trivial lower bounds on the locality of tests for affine-invariant properties in the general case. Below we describe the main technical question resolved in this paper (which has a self-contained description).

1.3 Zeroes of some special polynomial systems

The existing work on affine invariant testing [Kaufman and Sudan, 2008, Grigorescu et al., 2008, 2009] already converts questions about local testability of such codes into questions about solutions of certain systems of multivariate polynomial equations. (A solution of a system of equations is a common root, i.e., an assignment that sets all polynomials in the system to 0.) A non-trivial solution of the system leads to a non-trivial test. Consequently, non-existence of a solution implies no test is possible. Unfortunately, analyzing conditions under which a system of polynomials has a solution is nontrivial. In our case we manage to find a new class of polynomial systems for whom we can describe the zeroes exactly, and believe this result may be of independent interest.

The system we deal with is easy to describe: A *diagonal* polynomial $P(X_1, \ldots, X_k)$ over a field $\mathbb{F}$ is a polynomial $P \in \mathbb{F}[X_1, \ldots, X_k]$ in which each monomial depends on only one variable. A *homogenous* polynomial is one in which each monomial has the same degree. Thus each homogenous diagonal polynomial P_d is given by k coefficients $\lambda_1, \ldots, \lambda_k$ and the degree d and $P_d(X_1, \ldots, X_k) = \sum_{i=1}^{k} \lambda_i X_i^d$. Now a system of homogenous diagonal polynomials is *uniform* if each polynomial in the system is given by the same coefficients. To summarize, a *uniform system of homogenous diagonal polynomials* (UHD) is given by $\Lambda = (\lambda_1, \ldots, \lambda_k)$ and the set of degrees D and has the form

$$\mathcal{P}(D, \Lambda) = \left\{ \sum_{i=1}^{k} \lambda_i X_i^d \,\middle|\, d \in D \right\}. \tag{1}$$

It was shown by Kaufman and Sudan [2008] that every affine invariant family $\mathcal{F} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ induces a UHD $\mathcal{P}(D, \Lambda)$ over the algebraic closure of $\mathbb{F}$ (actually, the UHD is over the finite field of size $|\mathbb{K}|^m$). This UHD has a certain property that is crucial to our proof. Let p denote the characteristic of $\mathbb{F}$. If $d \in D$ has the base-p expansion $d = \sum_i d_i \cdot p^i$ where $d_i \in \{0, \ldots, p-1\}$ and $e = \sum_i e_i \cdot p^i$ satisfies $0 \leq e_i \leq d_i$ for all i, then e also belongs to D. We say that e lies in the *p-shadow* of d and that D is *p-shadow closed*. We also call $\sum_i d_i$ the *p-weight* of d.

The main mathematical contribution of this paper (Theorem 3.6) is that a UHD system $\mathcal{P}(D, \Lambda)$ that is p-shadow closed and such that D contains an exponent $d \in D$ of p-weight greater than k, has no *nontrivial* solution over the algebraic closure of $\mathbb{F}$. Regarding trivial solutions, 0 is always one. Another family of trivial solutions is obtained whenever Λ can be partitioned such that the elements of each partition sum up to 0. In this case we can assign the same value (which can be arbitrary) to all variables in a partition to obtain a solution, which we call trivial because it does not depend on D. To repeat, our main theorem says that no other solutions exist, and this result represents a rare case where one is able to fully describe the set of zeroes of a system of polynomials.

Let us now sketch how our result is related to the problem we started with, that of bounding the rate of affine invariant locally testable codes, by showing they are subcodes of Reed-Muller codes. Suppose $\mathcal{F} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ is an affine invariant family of functions that has a k-local constraint. As mentioned earlier, Kaufman and Sudan [2008] showed this implies that a UHD system $\mathcal{P}(D, \Lambda)$ has a solution of the form $(\alpha_1, \ldots, \alpha_k)$ where all α_i are distinct. In particular, this solution is nontrivial according to our definition of the term. So our main theorem implies D must contain elements of p-weight less than k. It turns out that this implies that $\mathcal{F}$ is isomorphic to a family of polynomials over $\mathbb{F}$ of degree at most k, i.e., it is a subcode of $\mathrm{RM}[p, n, k]$. (The case of fields of prime-power size is slightly more delicate, cf. Theorem 2.9.)

A brief word about how we prove the mathematical statement Theorem 3.6. We assume $\alpha = (\alpha_1, \ldots, \alpha_k)$ is a nontrivial solution to $\mathcal{P}(D, \Lambda)$. This assumption implies that the two operations of *(i)* Λ-weighted-summation and *(ii)* raising-to-power-d, commute with respect to $\alpha_1, \ldots, \alpha_{k-1}$. I.e., for certain powers $d \in D$ we have $\left(\sum_{i=1}^{k-1} \lambda_i \alpha_i\right)^d = \sum_{i=1}^{k-1} \lambda_i \alpha_i^d$. Equalities of this type are then manipulated to obtain our main result.

Organization of the rest of the paper Next we formally state our main results. In Section 3 we give an overview of the proof by reducing the problem of affine invariance rate to that of bounding the solution space of a UHD. Section 4 discusses algebraic properties of affine invariant codes that are essential to our proof. Section 5 completes the proof of our main theorem and we conclude in Section 6 with a characterization of the set of solutions to certain UHD systems of polynomials that is a consequence of our work.

2 Definitions and Main Results

2.1 Preliminaries — Locally testable, and Reed-Muller codes

Notation We use $[n]$ to denote the set $\{1, \ldots, n\}$. Throughout we let $\mathbb{F}, \mathbb{K}, \mathbb{L}$ denote fields. The q element finite field is denoted by $\mathbb{F}_q$. An $[N, K, D]_{\mathbb{F}}$-(linear) code is a K-dimensional subspace $\mathcal{C} \subseteq \mathbb{F}^N$ of Hamming distance D. Elements of $\mathcal{C}$ are referred to as codewords (of $\mathcal{C}$). Two vectors $u, w \in \mathbb{F}^N$ are said to be δ-close if they are within Hamming distance $\leq \delta N$ of each other, otherwise they are said to be δ-far. A vector u is said to be δ-close to $\mathcal{C}$ if it is δ-close to some codeword $w \in \mathcal{C}$, otherwise we say w is δ-far from $\mathcal{C}$. We define $\langle u, w \rangle \triangleq \sum_{i=1}^N u_i w_i$. Let $\mathcal{C}^\perp = \left\{u \in \mathbb{F}^N \;\middle|\; \langle u, w \rangle = 0 \text{ for all } w \in \mathcal{C}\right\}$ denote the space that is dual to $\mathcal{C}$ (it is also known as the *dual code* of $\mathcal{C}$).

We recall the standard definitions of a tester for a linear code and and a linear locally testable code. All codes considered in this paper are linear so from here on we drop further reference to this linearity (of testers, codes, etc.).

Definition 2.1 (Tester)**.** Suppose $\mathcal{C}$ is a $[N, K, D]_{\mathbb{F}}$-code. A *k-query tester* for $\mathcal{C}$ is a probabilistic oracle algorithm T that makes at most k oracle queries to a word $w \in \mathbb{F}^N$ and outputs an accept/reject verdict. The tester is said to have completeness c and ϵ-soundness s if it accepts every codeword of $\mathcal{C}$ with probability at least c and accepts words that are ϵ-far from $\mathcal{C}$ with probability at most s.

Definition 2.2 (Locally Testable Code (LTC))**.** An $[N, K, D]_{\mathbb{F}}$-code $\mathcal{C}$ is said to be a (k, ϵ, ρ)-*Locally Testable Code* (LTC) if there exists a k-query tester that has completeness c and ϵ-soundness $c - \rho$.

We are typically interested in infinite family of codes. If an infinite family of codes is a (k, ϵ, ρ)-LTC for absolute constants k and $\epsilon, \rho > 0$, then we simply refer to this (family of) code(s) as an LTC. For linear LTCs the nature of tests can be simplified significantly, due to a result of Ben-Sasson, Harsha, and Raskhodnikova [2005], to get them to a canonical form, which has perfect completeness ($c = 1$), and is *non-adaptive* (while the soundness parameter ρ changes by a constant factor). This leads to the following definition.

Definition 2.3 (Canonical tester)**.** A *canonical k-query test* for $\mathcal{C}$ is given by an element $u \in \mathcal{C}^\perp$ that has support size at most k, i.e., $|\{i \mid u_i \neq 0\}| \leq k$, where the test accepts $w \in \mathbb{F}^n$ if and only if $\langle u, w \rangle = 0$. A *k-query canonical tester T* for $\mathcal{C}$ is defined by a distribution μ over canonical k-query tests. Invoking the tester T on a word $w \in \mathbb{F}^n$ is done by sampling a test u according to the distribution μ and outputting accept if the canonical test given by u accepts.

The following proposition of Ben-Sasson et al. [2005] — stated as Theorem 3.3 there — shows that tests may always be assumed to be canonical (up to a constant factor change in soundness).

Proposition 2.4. *For every $\epsilon, \rho > 0$ and positive integer k, there exist $\rho' > 0$ such that every (k, ϵ, ρ)-LTC has a canonical k-query tester with perfect completeness and ϵ-soundness $1 - \rho'$.*

Our main theorem compares the performance of affine-invariant locally testable codes to that of Reed-Muller codes, which we define next. In what follows let $\binom{n}{\leq k} = \sum_{i=0}^{k} \binom{n}{i}$.

Definition 2.5 (Reed-Muller codes)**.** For $\mathbb{F}$ a finite field of size q and m, k integers, the m-variate Reed-Muller code of degree k over $\mathbb{F}$, denoted $\mathrm{RM}[q, m, k]$ is the $[N = q^m, K = \binom{m}{\leq k}, d = (q^m - q^{m-k})]_{\mathbb{F}}$-code whose codewords are all evaluations of m-variate polynomials over $\mathbb{F}$ of degree at most k.

These codes have also been studied for the testability properties (see, e.g., Rubinfeld and Sudan [1996], Alon et al. [2005], Samorodnitsky [2007], Kaufman and Ron [2004], Jutla et al. [2004], and Bhattacharyya et al. [2009]) and the case most relevant to us is that of constant q and k and arbitrarily large m. For this choice of parameters the codes are known to be $(q^{O(k)}, \epsilon, \rho)$-locally testable for some constants $\epsilon, \rho > 0$ that may depend on q and k [Alon et al., 2005, Kaufman and Ron, 2004].

2.2 Affine invariant codes

The main concept of interest to us is that of affine-invariance. We borrow some of the main definitions related to this concept from [Kaufman and Sudan, 2008].

From here on we associate a code with a family of functions. Let p be a prime, $\mathbb{F} = \mathbb{F}_q$ for $q = p^r$ be a finite field and let $\mathbb{K} = \mathbb{F}_Q$ for $Q = q^n$ be an extension of $\mathbb{F}$. For integer m we can consider $[N = Q^m, k, d]_{\mathbb{F}}$-codes whose entries are indexed by elements of $\mathbb{K}^m$. In other words, from here on a code will be identified with an $\mathbb{F}$-linear subspace of $\{\mathbb{K}^m \to \mathbb{F}\}$, the space of all functions from $\mathbb{K}^m$ to $\mathbb{F}$.

Definition 2.6 (Affine invariant codes)**.** Let $\mathbb{K}$ be a finite degree extension of $\mathbb{F}$. A code $\mathcal{C} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ is said to be *affine invariant* if it is invariant under the action of the affine semi-group[1] over $\mathbb{K}^m$. In other words, for every $f \in \mathcal{C}$ and every affine transformation $A : \mathbb{K}^m \to \mathbb{K}^m$, the function $f \circ A$ defined by $(f \circ A)(x) = f(A(x))$ belongs to $\mathcal{C}$ as well.

The work of Ben-Sasson et al. [2005] shows that in order for a linear property to be testable, it must have some "local constraints" (low-weight words in its dual). For affine invariant codes, Kaufman and Sudan [2008] show that when $\mathbb{K}$ is small, then the existence of such constraints is also a sufficient condition. (Our main result will show that the existence of such constraints imposes a bound on the rate of a code, over any field $\mathbb{K}$, not just over fields of constant size.) We recall the following definition from Kaufman and Sudan [2008].

[1] The set of all affine maps from $\mathbb{K}^m$ to itself forms a semi-group under composition. If one restricted this set to full rank maps, then one gets a group.

Definition 2.7 (k-local constraint)**.** A k-local constraint is given by k points in $\mathbb{K}^m$ $\alpha = (\alpha_1, \ldots, \alpha_k) \in (\mathbb{K}^m)^k$. We say that a code $\mathcal{C} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ satisfies (or, has) a *k-local constraint α* if there exists nonzero $\Lambda = (\lambda_1, \ldots, \lambda_k) \in \mathbb{K}^k$ such that $\sum_{i=1}^k \lambda_i f(\alpha_i) = 0$ for every $f \in \mathcal{C}$.

The following statement is the main result of Kaufman and Sudan [2008] regarding the local testability of affine invariant codes, and is stated as Theorem 2.10 there.

Theorem 2.8 (Affine invariant codes satisfying a k-local constraint over a small field are locally testable)**.** *For fields $\mathbb{F} \subseteq \mathbb{K}$ with $|\mathbb{F}| = q$ and $|\mathbb{K}| = Q$, let $\mathcal{F} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ be an affine-invariant code satisfying a k-local constraint. Then for any $\delta > 0$, the code $\mathcal{F}$ is*

$$\left(k' = (Q^2k)^{Q^2}, \delta, \frac{\delta}{2(2k'+1)(k'+1)}\right)\text{-locally testable.}$$

Notice the above theorem implies local testability only when the field $\mathbb{K}$ is relatively small, and is of interest only when $m \to \infty$. When m is small (and $\mathbb{K}$ large) no general bounds were known on the locality of the tests. Grigorescu et al. [2008] show that it is possible to have affine invariant families with one 8-local constraint that is not $O(1)$-locally characterized. And all this previous work leaves open the possibility that there may exist other affine-invariant families that are $O(1)$-locally characterized, perhaps even $O(1)$-testable (say, over fields of growing size and $m = 1$), and do have large rate. Our work rules this out.

We can now state our main theorem which bounds the rate of affine invariant codes containing a k-local constraint.

Theorem 2.9 (Affine invariant families with a local constraint are contained in low-degree Reed-Muller codes)**.** *Let p be a prime and r, n, m be positive integers and let $q = p^r$ and $Q = q^n$. For $\mathbb{F} = \mathbb{F}_q$ and $\mathbb{K} = \mathbb{F}_Q$ a degree-n extension of $\mathbb{F}$, let $\mathcal{C} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ be an affine-invariant family that satisfies a k-local constraint. Then*

1. *The dimension of $\mathcal{C}$ as a vector space over $\mathbb{F}_q$ is at most $(mrn)^{k-1}$. Since the blocklength of $\mathcal{C}$ is $Q = p^{rmn}$ we get*
$$\dim(\mathcal{C}) \leq (\log_p Q)^{k-1}.$$
2. *$\mathcal{C}$ is isomorphic to a subcode of* $\mathrm{RM}[q, nm, (k-1)q/p]$.[2] *In particular, for $q = p$ we get that $\mathcal{C}$ is isomorphic to a subcode of* $\mathrm{RM}[p, nm, k-1]$.

Note that the when $q = p$, Part (1) of the theorem above follows from Part (2), since the dimension of $\mathrm{RM}[p, nm, k]$ is at most $(mn)^k$. When $q = p^r$ for $r > 1$, this is not true, and the dimension of the code $\mathrm{RM}[q, nm, kq/p]$ is much larger. In this case Part (2) is a weak description of our understanding of $\mathcal{C}$. A somewhat better understanding of affine-invariant codes over $\mathbb{F}_{p^r}$, for $r > 1$ can be obtained if we use a broader class of codes. In particular, by viewing a code over $\mathbb{F}_{p^r}$ as a code over the vector space $\mathbb{F}_p^r$, or as an r-tuple of codes over $\mathbb{F}_p$, one gets a more strict inclusion for such codes. Specifically, let $\mathrm{RM}[p, n, k-1]^r$ denote codes obtained by evaluations of $f = \langle f_1, \ldots, f_r \rangle : \mathbb{F}_p^n \to \mathbb{F}_p^r$, where each $f_i : \mathbb{F}_p^n \to \mathbb{F}_p$ is an n-variate polynomial over $\mathbb{F}_p$ of degree at most $k-1$. We then have the following Corollary of Theorem 2.9.

Corollary 2.10 (Affine invariant families with a local constraint over fields of prime powers)**.** *Let p be a prime and r, n, m be positive integers and let $q = p^r$ and $Q = q^n$. For $\mathbb{F} = \mathbb{F}_q$ and $\mathbb{K} = \mathbb{F}_Q$ the degree-n extension of $\mathbb{F}$, let $\mathcal{C} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ be an affine-invariant family that satisfies a k-local constraint. The for every $\mathbb{F}_p$-linear bijection $\psi : \mathbb{F}_q \to \mathbb{F}_p^r$, the code $\mathcal{C}' = \{\psi \circ f | f \in \mathcal{C}\} \subseteq \{\mathbb{K}^m \to \mathbb{F}_p^r\}$ is isomorphic to a subcode of* $\mathrm{RM}[p, nmr, k-1]^r$.

[2] In other words, there exists an isomorphism $\phi : \mathbb{F}^{nm} \to \mathbb{K}^m$ such that for every $f \in \mathcal{C}$, the function $(f \circ \phi) \in \{\mathbb{F}^{nm} \to \mathbb{F}\}$ defined by $(f \circ \phi)(x) = f(\phi(x))$ belongs to $\mathrm{RM}[q, nm, (k-1)q/p]$.

Proof. For $i \in [r]$, let $\mathcal{C}_i$ be the projection of $\mathcal{C}'$ to the ith coordinate. Then $\mathcal{C}_i$ is a $\mathbb{F}_p$-linear, affine-invariant code (over the domain $\mathbb{F}_Q^m$). By Theorem 2.9 we get that it is isomorphic to a subcode of $\mathrm{RM}[p, nmr, k-1]$. It follows that $\mathcal{C}'$ is a subcode of $\mathrm{RM}[p, nmr, k-1]^r$. □

3 Proof of Main Theorem

In this section we prove Theorem 2.9 modulo some technical lemmas. We give an overview of the proof first. Our first step (Lemma 3.1) is a "reduction" to the univariate case. This step follows easily from the fact that a vector space over some field can be embedded into an extension of the field, with affine transformations over the extension being affine transformations over the vector space. Thus affine invariance over the vector space implies affine invariance over the large field.

Next, we exploit the (known) structure of univariate affine invariant families to reduce the question of the availability of local constraints to a question of finding a non-trivial root for a special class of multivariate polynomial equations (see Lemmas 3.3 and 3.5). Our main technical contribution is to show that this class of polynomial equations do not have non-trivial roots and this is done in Theorem 3.6. The combination of Lemmas 3.3, 3.5 and Theorem 3.6 give structural limitations on affine-invariant families with local constraints and in Lemma 3.7 we show that these limitations imply that the codes are contained in Reed-Muller codes, allowing us to conclude the proof of Theorem 2.9.

3.1 Reducing to univariate functions

Lemma 3.1. *If Theorem 2.9 holds for the case $m = 1$ then it holds for all positive integers m.*

Since the proof is immediate, we include it here.

Proof. Consider a case where $m \neq 1$ and let $\mathbb{L}$ be the field of size $|\mathbb{K}|^m$. Fix a $\mathbb{K}$-linear bijection ϕ from $\mathbb{L}$ to $\mathbb{K}^m$. Note that every affine map $a(x) = \alpha x + \beta$ mapping $\mathbb{L}$ to $\mathbb{L}$, with $\alpha, \beta \in \mathbb{L}$, corresponds under ϕ to a linear map $A(y) = By + c$ with $B \in \mathbb{K}^{m \times m}$ and $c \in \mathbb{K}^m$ such that for every $x \in \mathbb{L}$, $\phi(a(x)) = A(\phi(x))$. It follows that if $\mathcal{C} \subseteq \{\mathbb{K}^m \to \mathbb{F}\}$ is affine-invariant then so is $\mathcal{C}' \subseteq \{\mathbb{L} \to \mathbb{F}\}$ defined as $\mathcal{C}' = \{f \circ \phi | f \in \mathcal{C}\}$. The theorem statement for $\mathcal{C}$ now follows from the theorem statement for $\mathcal{C}'$. □

3.2 Degree Sets of Affine-Invariant Families

From now one we consider only univariate functions, i.e., $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$. Recall that every function from $\mathbb{K} \to \mathbb{K}$ and hence from $\mathbb{K} \to \mathbb{F}$ is the evaluation of a polynomial in $\mathbb{K}[x]$ of degree at most $q^n - 1$. For a polynomial $p \in \mathbb{K}[x]$ given by $p(x) = \sum_d c_d x^d$, let $\mathrm{supp}(p)$ denote its support, i.e., $\mathrm{supp}(p) = \{d | c_d \neq 0\}$. The set of degrees in the support of the functions in $\mathcal{C}$ turns out be a central ingredient in understanding the structure of $\mathcal{C}$, motivating the following definition.

Definition 3.2 (Degree set of $\mathcal{C}$)**.** For a class of functions $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$, its degree set is the set $D(\mathcal{C}) = \cup_{p \in \mathcal{C}} \mathrm{supp}(p)$.

It turns out that the representations of elements of $D(\mathcal{C})$ in base p play a central role in the structure of affine-invariant families over fields of characteristic p. To this end we introduce some terminology.

For integer d, let $[d]_p = \langle d_0, d_1, \ldots \rangle$ denotes its representation in base p (i.e., $0 \leq d_i < p$ and $d = \sum_{i=0}^{\infty} d_i p^i$). The *$p$-weight* of d, denoted $\mathrm{wt}_p(d)$, is the quantity $\sum_{i=0}^{\infty} d_i$. We say e is in the p-shadow of d, denoted $e \leq_p d$, if $[e]_p = \langle e_0, e_1, \ldots \rangle$ and $e_i \leq d_i$ for all i. The set $\{e | e \leq_p d\}$ is called the p-shadow of d.

Lemma 3.3. *For every affine invariant family $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$ where $\mathbb{F}, \mathbb{K}$ are fields of characteristic p, $D(\mathcal{C})$ is closed under p-shadow, i.e., if $d \in D(\mathcal{C})$ and $e \leq_p d$ the $e \in D(\mathcal{C})$.*

We prove Lemma 3.3 in Section 4.

Aside: $D(\mathcal{C})$ also satisfies the additional property that it is $(q, q^n - 1)$-modular, i.e., if $d \in D(\mathcal{C})$ then also $q \cdot d \pmod{q^n - 1} \in D(\mathcal{C})$ (cf. Grigorescu et al. [2009]). With this additional property we get an exact characterization of degree sets. Specifically, if $D \subseteq \{0, \ldots, q^n - 1\}$ is closed under p-shadow and is $(q, q^n - 1)$-modular, then D is the degree set of an affine-invariant family. However we don't use these additional facts below.

3.3 Uniform Homogenous Diagonal Systems of Polynomial Equations

The task of finding the set of zeroes of a system of multivariate polynomial equations is a central theme in mathematics. (Linear algebra considers the special case where all equations are linear/affine and understanding the "variety" of a given system of (higher-degree) equations is a central theme in algebraic geometry.) In general of course, the set of zeroes may be too complex, even for degree two polynomials. Nevertheless, our quest to understand the locality of constraints in an affine-invariant property leads to such a question, where the set of polynomials has a reasonably clean description. Somewhat surprisingly, we are even able to describe the set of zeroes in a fairly precise way. We describe the class of polynomial systems that we consider next.

Definition 3.4 (Uniform Homogenous Diagonal (UHD) System)**.** Fix a system of polynomials $P_1, \ldots, P_m \in \mathbb{F}[X_1, \ldots, X_k]$.

- We say the system is *homogenous* if every polynomial in the system is homogenous.
- We say that the system is *diagonal* if every monomial in the support of every polynomial is a power of a single variable. I.e, a homogenous system is diagonal if for every $j \in [m]$, it is the case that $P_j(X_1, \ldots, X_k) = \sum_{i=1}^{k} \lambda_{ji} \cdot X_i^{d_j}$.
- We say a homogenous diagonal system is *uniform* if the coefficients are the same for every polynomial, i.e., λ_{ji} is independent of j.

We conclude that a uniform homogenous diagonal system is given by a sequence of coefficients $\Lambda = \langle \lambda_1, \ldots, \lambda_k \rangle \in \mathbb{F}^k$ and degrees $D = \{d_1, \ldots, d_m\}$ such that $P_j(X_1, \ldots, X_k) = \sum_{i=1}^{k} \lambda_i X_i^{d_j}$. We refer to such a system as the (D, Λ)-UHD system. We say that the (D, Λ)-system has a *pairwise-distinct solution* over some field $\mathbb{K}$ if there exist *distinct* values $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$ such that $P_j(\alpha_1, \ldots, \alpha_k) = 0$ for every $j \in [m]$.

The following lemma motivates the study of UHD systems in our setting.

Lemma 3.5. *If an affine-invariant property $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$ has a k-local constraint, then there exists a non-zero vector $\Lambda \in \mathbb{F}^k$ such that the $(D(\mathcal{C}), \Lambda)$-a UHD system has a pairwise-distinct solution over $\mathbb{K}$.*

We prove this lemma in Section 4.2.

We note that any homogenous system always has a trivial solution given by $\alpha_1 = \ldots = \alpha_k = 0$. Some UHD systems have other trivial solutions, for instance, if $\sum_i \lambda_i = 0$ then taking any $\alpha_1 = \ldots = \alpha_k$ gives a trivial solution. Other such "trivial" solutions can be found if some subset of the coefficients sum to zero. However all solutions involve assigning non-distinct values to the variables, and our lemma above says that such solutions are not of interest to us. This leads to the main focus of this paper, which is the set of pairwise-distinct solutions to UHD systems and the following lemma shows that if D is p-shadow-closed and contains a degree of high p-weight, then the UHD system has no pairwise-distinct solutions over any field of characteristic p.. We point out that the statement below is more general than needed for our purposes. In particular, to prove Theorem 2.9 it suffices to consider $\lambda_i \in \mathbb{F}$ and restrict our attention to solutions in $\mathbb{K}$. However, as the proof generalizes to any extension and because we believe this statement may be of interest in other settings we state and prove the more general form.

Theorem 3.6 (Shadow-closed UHD systems containing nontrivial solutions have bounded weight)**.** *Let $\mathbb{F}$ be any field of characteristic p and let D be a p-shadow-closed set of integers containing an element d with $\mathrm{wt}_p(d) \geq k$. Then for every $\Lambda = (\lambda_1, \ldots, \lambda_k) \in \mathbb{F}^k$ where not all λ_i's are zero, the (D, Λ)-UHD system has no pairwise-distinct solutions over $\mathbb{K}$ for any field $\mathbb{K}$ extending $\mathbb{F}$.*

We prove this theorem in Section 5. In Section 6, we show how Theorem 3.6 actually yields a fairly combinatorial characterization of all the zeroes (trivial or otherwise) of a (D, Λ)-UHD system where D is p-shadow closed, and contains an element of weight greater than k. See Theorem 6.1.

3.4 Containment in Reed-Muller Codes

The lemmas of the previous section effectively give us our main theorem. The only remaining ingredient is to show that families $\mathcal{C}$ with $D(\mathcal{C})$ containing only small weight elements are isomorphic to Reed-Muller codes.

Lemma 3.7. *Let $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$ be an affine invariant family, with $\mathbb{F}, \mathbb{K}$ being fields of characteristic p and respective sizes $q = p^r$ and $Q = q^n$. Suppose further that $D(\mathcal{C})$ contains elements of p weight less than k.*

1. *For every $\mathbb{F}$-linear bijection $\phi : \mathbb{K} \to \mathbb{F}^n$, the family $\mathcal{C}' = \{f \circ \phi^{-1} \mid f \in \mathcal{C}\} \subseteq \{\mathbb{F}^n \to \mathbb{F}\}$ is contained in* $\mathrm{RM}[q, n, (k-1)q/p]$.
2. *For every $\mathbb{F}$-linear bijection $\phi : \mathbb{K} \to \mathbb{F}_p^{nr}$ and every $\mathbb{F}_p$-linear map $\psi : \mathbb{F} \to \mathbb{F}_p$, the family $\mathcal{C}' = \{\psi \circ f \circ \phi^{-1} \mid f \in \mathcal{C}\} \subseteq \{\mathbb{F}^n \to \mathbb{F}_p\}$ is contained in* $\mathrm{RM}[p, nr, k-1]$.

We prove this lemma in Section 4.3.

3.5 Proof of Main Theorem

We are now ready to prove the main theorem assuming the lemmas claimed in the previous subsections.

Theorem 2.9. By Lemma 3.1, we know that it suffices to prove the theorem for the univariate case (i.e., $m = 1$). Let $D(\mathcal{C})$ be the degree set of $\mathcal{C}$. By Lemma 3.3, we know that $D(\mathcal{C})$ is p-shadow closed. Furthermore if $\mathcal{C}$ has a k-local constraint then, by Lemma 3.5 there exists a non-zero vector $\Lambda \in \mathbb{F}^{k'}$ such that the $(D(\mathcal{C}), \Lambda)$-UHD system has a pairwise-distinct solution. But then, by Theorem 3.6, we have that the weight of every element $d \in D(\mathcal{C})$ must be at most $k - 1$.

The dimension of $\mathcal{C}$, which is at most $|D(\mathcal{C})|$, can now be bounded from above by the number of integers $d \in \{0, \ldots, q^n - 1\}$ of p-weight less than k which is (crudely) at most $(rn)^{k-1}$ (where $q = p^r$). By Lemma 3.7, Part (1), it follows that $\mathcal{C}$ is isomorphic to a subcode of $\mathrm{RM}[q, nm, (k-1)q/p]$, thus concluding the proof of the theorem. □

4 Structure of affine-invariant families

In this section we prove Lemmas 3.3 and 3.5. Most of the notions and claims here are standard and in particular based on the work of Kaufman and Sudan [2008], Grigorescu et al. [2008, 2009], but some of our definitions are cleaner and proofs included for completeness.

4.1 Degree sets of affine-invariant families

We start with a lemma which is basic to all affine-invariant families. Two concepts are central to this section, the "Trace" function, and the "order" of an integer. We define these first.

The Trace function is a basic map from the extension of a field to the base field. In particular for $\mathbb{K} = \mathbb{F}_{q^n}$ extending $\mathbb{F} = \mathbb{F}_q$, the Trace function $\mathrm{Trace}_{\mathbb{K},\mathbb{F}} : \mathbb{K} \to \mathbb{F}$ is the function $\mathrm{Trace}_{\mathbb{K},\mathbb{F}}(x) = x + x^q + x^{q^2} + \cdots + x^{q^{n-1}}$. We often suppress the subscripts when the domain is $\mathbb{K}$ and range is $\mathbb{F}$. It is easily verified that $\mathrm{Trace}(\alpha)^q = \mathrm{Trace}(\alpha)$ for $\alpha \in \mathbb{K}$ and so the range of Trace is indeed $\mathbb{F}$. It can also be verified that $\mathrm{Trace}(x + y) = \mathrm{Trace}(x) + \mathrm{Trace}(y)$ and $\mathrm{Trace}(\lambda \cdot x) = \lambda \cdot \mathrm{Trace}(x)$ for $\lambda \in \mathbb{F}$, and so the Trace function is $\mathbb{F}$-linear. Finally (and this will be important below), we note that the Trace function is a q^{n-1}-to-1 map from $\mathbb{K}$ to $\mathbb{F}$ (i.e., every element of $\mathbb{F}$ has exactly q^{n-1} pre-images mapping to it).

The order of a non-negative integer d, denoted $\mathrm{ord}_{q,n}(d)$ (again we drop the subscripts since they will always be q and n), is the smallest positive integer b such that $d \cdot q^b = d (\mathrm{mod}\ q^n - 1)$. Note such a b always exists and is at most n. Below we mention some basic properties of the order function.

Proposition 4.1. *Let $b = \mathrm{ord}_{q,n}(d)$. Let $\mathbb{F} = \mathbb{F}_q$, $\mathbb{L} = \mathbb{F}_{q^b}$ and $\mathbb{K} = \mathbb{F}_{q^n}$. Then the following hold:*

1. *b divides n.*
2. *$\mathbb{L} = \{\sum_{\alpha\in\mathbb{K}} c_\alpha \cdot \alpha^d | c_\alpha \in \mathbb{F}\}$, i.e., the set of all $\mathbb{F}$-linear sums of dth powers of elements of $\mathbb{K}$.*
3. *For every $\lambda \in \mathbb{K}$, $\mathrm{Trace}_{\mathbb{K},\mathbb{F}}(\lambda x^d) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\lambda) x^d)$.*

Proof. Part (1) is standard and follows from the fact that $d \cdot q^{\gcd(b,n)} = d \pmod{q^n - 1}$ and so if b does not divide n, then it contradicts its minimality. To see the fact mentioned above write $\gcd(b, n) = ab - cn$ (the case of negative a and positive c is handled identically). We have $d \equiv d \cdot q^{ab}$ and $d \equiv d \cdot q^{cn} \pmod{q^n - 1}$ which implies $d \cdot q^{ab} \equiv d \cdot q^{cn} \pmod{q^n - 1}$. Divide both sides by q^{cn} to obtain $d \equiv d \cdot q^{\gcd(b,n)} \pmod{q^n - 1}$. This division uses the fact that $q^n - 1$ does not divide q^{cn}.

For Part (2), let $S = \{\sum_{\alpha\in\mathbb{K}} c_\alpha \cdot \alpha^d | c_\alpha \in \mathbb{F}\}$. The containment $S \subseteq \mathbb{L}$ is immediate from the fact that for every sequence $\langle c_\alpha \rangle$, we have $(\sum_\alpha c_\alpha \alpha^d)^{q^b} = \sum_\alpha c_\alpha \alpha^d$ (and $\mathbb{L} = \{\gamma \in \mathbb{K} | \gamma^{q^b} = \gamma\}$). In the other direction, note that S is closed under multiplication and addition and thus is a subfield of $\mathbb{K}$. If $S \neq \mathbb{L}$ then it must be of size q^c for some $c < b$. But then it must be case that for every element $\alpha \in \mathbb{K}$, $(\alpha^d)^{q^c} = \alpha^d$ and thus $(X^{d \cdot q^c} - X^d) \equiv 0 \pmod{X^{q^n} - X}$ which implies $d \cdot q^c = d \pmod{q^n - 1}$ contradicting the minimality of b. We conclude that S must equal $\mathbb{L}$.

Finally Part (3) follows immediately from the definition of the Trace function. □

The role of the Trace function in our study is given by the following lemma. The order function plays a central role in the proof.

Lemma 4.2. *Let $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$ be an affine-invariant family with degree set $D = D(\mathcal{C})$. Then $\mathcal{C} = \{\mathrm{Trace}(p(x)) | p(x) \in \mathbb{K}[x], \mathrm{supp}(p) \subseteq D\}$.*

Proof. Let $\mathcal{D} = \{\mathrm{Trace}(p(x)) | p(x) \in \mathbb{K}[x], \mathrm{supp}(p) \subseteq D\}$. We need to show $\mathcal{C} \subseteq \mathcal{D}$ and $\mathcal{D} \subseteq \mathcal{C}$ and we do so in order. First we make a structural observation about functions mapping $\mathbb{K}$ to $\mathbb{F}$. Consider a function $f : \mathbb{K} \to \mathbb{F}$ and let $f(x) = \sum_{d=0}^{q^n-1} c_d x^d$. Since $f(x)^q = f(x) (\bmod\ x^{q^n} - x)$ (since for every $x \in \mathbb{K}$ we have $f(x) \in \mathbb{F}$), it follows that $\sum_{d=0}^{q^n-1} (c_d x^d)^q = \sum_{d=0}^{q^n-1} c_d x^d (\bmod\ x^{q^n} - x)$. We conclude that for every $d \in [q^n - 1]$ we have $c_{qd (\bmod\ q^n-1)} = c_d^q$, In particular this implies that for every d, $(c_d)^{q^b} = c_d$ and so $c_d \in \mathbb{F}_{q^b}$ where $b = \mathrm{ord}(d)$.

We now prove that $\mathcal{C} \subseteq \mathcal{D}$. We do so by induction on the size of the support of $f \in \mathcal{C}$. Let $f(x) = \sum_e c_e x^e$ and let d be an index such that $c_d \neq 0$. Let $b = \mathrm{ord}(d)$ and $\mathbb{L} = \mathbb{F}_{q^b}$. Now consider the polynomial $g(x) = \sum_{j=0}^{b-1} (c_d x^d)^{q^j} = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(c_d x^d)$. Note that the support of g is a subset of the support of f and on the support of g the coefficients of g match the coefficients of f. Thus, we can use induction to claim that $f - g \in \mathcal{D}$ and so it suffices to show that $g(x) \in \mathcal{D}$. As noted earlier, $c_d \in \mathbb{L}$. We now use the property that the Trace function $\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(x) = x + x^{q^b} + \cdots + x^{q^{n-b}}$ is a q^{n-b}-to-1 map to claim that there must exist some $\alpha \in \mathbb{K}$ such that $\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\alpha) = c_d$. We now use Proposition 4.1, Part (3), to see that $g(x) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(c_d x^d) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\alpha) x^d) = \mathrm{Trace}_{\mathbb{K},\mathbb{F}}(\alpha x^d)$ which is clearly in $\mathcal{D}$ as required.

Next we show the other direction, namely that $\mathcal{D} \subseteq \mathcal{C}$. For this part note that it suffices to verify that for every $d \in D$ and $\lambda \in \mathbb{K}$, we have $\mathrm{Trace}(\lambda x^d) \in \mathcal{C}$. (The containment $\mathrm{Trace}(p(x)) \in \mathcal{C}$ will then follow from the linearity of the Trace function and of $\mathcal{C}$.) Fix such a d and λ. Since $d \in D$, we must have some function $f \in \mathcal{C}$ with $f(x) = \sum_e c_e x^e$ such that $c_d \neq 0$. If $d = 0$, then the function $f(0 \cdot x) = c_0 \in \mathcal{C}$, by affine-invariance. By $\mathbb{F}$-linearity of $\mathcal{C}$, we now have that every constant function including $\mathrm{Trace}(\lambda x^0) \in \mathcal{C}$. Now consider $d \neq 0$. First note that w.l.o.g. we may assume $c_0 = 0$, since otherwise we can simply work with the function $\tilde{f}(x) = f(x) - f(0 \cdot x)$ which is also contained in $\mathcal{C}$ and has d in its support. Again let $b = \mathrm{ord}(d)$ and $\mathbb{L} = \mathbb{F}_{q^b}$ as usual. Now consider the function $h(x) = \sum_{\alpha \in \mathbb{K}^*} \sum_{j=0}^{b-1} \alpha^{dq^j} f(\alpha^{-1} x)$. We claim (1) $h(x) \in \mathcal{C}$ and (2) $h(x) = -\mathrm{Trace}_{\mathbb{L},\mathbb{F}}(c_d x^d)$. For (1), recall that $\alpha^d \in \mathbb{L}$ (Proposition 4.1, Part (2)). So $\sum_{j=0}^{b-1} \alpha^{d \cdot q^j} = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\alpha^d) \in \mathbb{F}$. Thus $h(x) = \sum_{\alpha \in \mathbb{K}^*} \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\alpha^d) \cdot f(\alpha^{-1} x)$ is an $\mathbb{F}$-linear combination of functions in $\mathcal{C}$ and thus $h(x) \in \mathcal{C}$. For (2), note that

$$\begin{aligned} h(x) &= \sum_{\alpha \in \mathbb{K}^*} \sum_{j=0}^{b-1} \alpha^{dq^j} f(\alpha^{-1} x) = \sum_j \sum_\alpha \alpha^{dq^j} \sum_{e=1}^{q^n-1} c_e \alpha^{-e} x^e \\ &= \sum_j \sum_e c_e x^e \sum_\alpha \alpha^{dq^j - e} = -\sum_j c_{d \cdot q^j} x^{d \cdot q^j} \end{aligned}$$

where the final identity holds since $\sum_{\alpha \in \mathbb{K}^*} \alpha^i$ equals 0 if $i \neq 0$ and $|\mathbb{K}^*| \equiv -1 \pmod{p}$ when $i = 0$. Using the fact that $c_{d \cdot q^j} = c_d^{q^j}$, we have $\mathrm{Trace}_{\mathbb{L},\mathbb{F}}(c_d x^d) = \sum_{j=0}^{b-1} (c_d x^d)^{q^j} = -h(x) \in \mathcal{C}$ and (2) follows. By affine-invariance and $\mathbb{F}$-linearity of $\mathcal{C}$, we also have $\mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\sum_{\alpha \in \mathbb{K}} a_\alpha c_d (\alpha x)^d) \in \mathcal{C}$ for every sequence $\langle a_\alpha | a_\alpha \in \mathbb{F} \rangle$ and so we have $\mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\gamma x^d) \in \mathcal{C}$ for every $\gamma \in \{c_d \cdot \sum_{\alpha \in \mathbb{K}} a_\alpha \cdot \alpha^d | a_\alpha \in \mathbb{F}\}$. We now claim that this set is exactly $\mathbb{L} = \mathbb{F}_{q^b}$ since $c_d \in \mathbb{L}^*$ and the set $\{\sum_{\alpha \in \mathbb{K}} a_\alpha \cdot \alpha^d | a_\alpha \in \mathbb{F}\} = \mathbb{L}$ (Proposition 4.1, Part (2)).

We can now show easily that for every $\lambda \in \mathbb{K}$, $\mathrm{Trace}(\lambda x^d) \in \mathcal{C}$. By Proposition 4.1, Part (3), we have $\mathrm{Trace}(\lambda x^d) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\lambda) x^d) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(\gamma x^d)$ for $\gamma = \mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\lambda) \in \mathbb{L}$, which is in $\mathcal{C}$ as shown above. □

We now prove Lemma 3.3. We recall the statement below.

Lemma 3.3 (restatement): *For every affine invariant family $\mathcal{C} \subseteq \{\mathbb{K} \to \mathbb{F}\}$ where $\mathbb{F}, \mathbb{K}$ are fields of characteristic p, $D(\mathcal{C})$ is closed under p-shadow, i.e., if $d \in D(\mathcal{C})$ and $e \leq_p d$ the $e \in D(\mathcal{C})$.*

Proof. Let $b = \mathrm{ord}(d)$ and $\mathbb{L} = \mathbb{F}_{q^b}$. From Lemma 4.2 we have that $\mathrm{Trace}(\lambda x^d) \in \mathcal{C}$ for every $\lambda \in \mathbb{K}$ and in particular, there exists λ_0 such that $\mathrm{Trace}_{\mathbb{K},\mathbb{L}}(\lambda_0) = 1$ so that the function $h(x) = \mathrm{Trace}_{\mathbb{K},\mathbb{F}}(\lambda_0 x^d) = \mathrm{Trace}_{\mathbb{L},\mathbb{F}}(x^d) \in \mathcal{C}$. Now consider the functions $h_a(x) = h(x + a)$. We claim that there exists an $a \in \mathbb{K}$ such that the coefficient of x^e in $h_a(x)$ is non-zero. This suffices to prove the lemma since $h_a(x) \in \mathcal{C}$ and so $e \in \mathrm{supp}(h_a) \subseteq D(\mathcal{C})$.

To see the claim, note that the coefficient of x^e in $h_a(x)$ is itself some polynomial in a, say $g(a)$. We focus on the coefficient of a^{d-e} in $g(a)$ and note that the only term that contributes to this polynomial is from the expansion of $(x + a)^d$ (and any other term $(x + a)^{dq^j}$ for $j \neq 0$ contribute to the coefficient of a^{dq^j-e} in $g(a)$). Thus the coefficient of a^{d-e} in $g(a)$ equals $\binom{d}{e}$ which is non-zero if modulo p (and only if) $e \leq_p d$. We conclude that $g(a)$ is itself a non-zero polynomial, and so there exists a setting of a which makes $g(a)$ non-zero. □

4.2 Local constraints and UHD systems

We now prove Lemma 3.5 which asserts that the presence of a k-local constraint implies that a certain UHD family has a pairwise-distinct solution.

Lemma 3.5. Notice that the set $\{\langle f(\alpha_1), \ldots, f(\alpha_k)\rangle | f \in \mathcal{C}\}$ is a linear space strictly contained in $\mathbb{F}^k$ and so there must exist $\lambda_1, \ldots, \lambda_k \in \mathbb{F}$, not all zero, so that $\sum_{i=1}^k \lambda_i f(\alpha_i) = 0$ for every $f \in \mathcal{C}$. Fix $d \in D(\mathcal{C})$. Then, using the fact that $\mathrm{Trace}(\beta x^d) \in \mathcal{C}$ for every $\beta \in \mathbb{K}$ we see that $\sum_{i=1}^k \lambda_i \mathrm{Trace}(\beta \alpha_i^d) = 0$ for every β. By the linearity of the Trace function, we thus have $\mathrm{Trace}(\beta \cdot \sum_{i=1}^k \lambda_i \alpha_i^d) = 0$. But this implies $\gamma = \sum_{i=1}^k \lambda_i \alpha_i^d = 0$ since otherwise there will exist some $\beta \in \mathbb{K}$ such that $\mathrm{Trace}(\beta\gamma) \neq 0$. We conclude that for every $d \in D(\mathcal{C})$ we have $\sum_{i=1} \lambda_i \alpha_i^d = 0$. Restating in the language of polynomial systems, we have that the $(D(\mathcal{C}), \Lambda)$-UHD has a pairwise-distinct solution at $x_1 = \alpha_1, \ldots, x_k = \alpha_k$, for $\Lambda = \langle \lambda_1, \ldots, \lambda_k \rangle$. □

4.3 Reed-Muller codes

We include below the (simple) proof of Lemma 3.7 which asserts that families $\mathcal{C}$ whose degree sets only contain low-weight integers are effectively (by isomorphisms) contained in Reed-Muller families (i.e., are low-degree multivariate polynomials over the base field).

Lemma 3.7. For Part (1) it suffices to consider a monomial X^d with $d \in D$ of p-weight less than k and show that $\phi^{-1}(x_1, \ldots, x_n)^d$ is a polynomial over $\mathbb{F}$ of degree at most $(k-1)q/p$ (possibly with coefficients in $\mathbb{K}$). Let $d = \sum_j d_j q^j$. Note that since d has p-weight at most k, it follows that $\sum_j d_j \leq (k-1)q/p$. For $i \in [n]$, let $\zeta_i = \phi^{-1}(0^{i-1}10^{n-i}) \in \mathbb{K}$. By the $\mathbb{F}$-linearity of ϕ^{-1} we have $\phi^{-1}(x_1, \ldots, x_n) = \sum_{i=1}^n \zeta_i x_i$. We now have $\phi^{-1}(x_1, \ldots, x_n)^d = (\sum_{i=1}^n \zeta_i x_i)^{\sum_j d_j q^j} = \prod_j (\sum_i \zeta_i^{q^j} x_i)^{d_j}$ which is a polynomial of degree at most $\sum_j d_j \leq (k-1)q/p$ in $x_1, \ldots, x_n$, yielding Part (1).

Part (2) follows immediately from Part (1) and the observation that the family $\mathcal{C}' = \{\psi \circ f | f \in \mathcal{C}\}$ is an affine-invariant family mapping $\mathbb{K}$ to $\mathbb{F}_p$ with $D(\mathcal{C}') \subseteq \mathcal{D}(\mathcal{C})$. □

5 Proof of Theorem 3.6

We now prove our main technical theorem, Theorem 3.6. Recall that this theorem states that if $\lambda_1, \ldots, \lambda_k \in \mathbb{F}$ are not all zero and $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$ are distinct elements such that $\sum_{i=1}^k \lambda_i \alpha_i^d = 0$ for every $d \in \mathcal{D}$ for some p-shadow closed set D, then $\mathrm{wt}_p(d) < k$ for every $d \in D$. Notice that it suffices to prove the theorem for the case where $D = \text{shadow}_p(d)$ for some integer d. (Else we can simply take the element d of largest weight in D and work with the set $D' = \text{shadow}_p(d)$.)

We say that $\alpha_1, \ldots, \alpha_k$ is *d-good* if such $\lambda_1, \ldots, \lambda_k$ exist (in $\mathbb{K}$), i.e., if there exist $\lambda_1, \ldots, \lambda_k \in \mathbb{K}$, not all zero, such that $\sum_{i=1}^{k} \lambda_i \alpha_i^e = 0$ for every $e \in \text{shadow}_p(d)$.

The key to our analysis in this section is a notion that we refer to as the *d-rank* of the sequence $\alpha_1, \ldots, \alpha_k$. Given distinct elements $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$ and positive integer d, let $M = M[d; \alpha_1, \ldots, \alpha_k]$ be the $|\text{shadow}_p(d)| \times k$ matrix of elements from $\mathbb{K}$ whose rows are indexed by elements of $\text{shadow}_p(d)$ and columns by elements of $\{1, \ldots, k\}$ and where $M_{e,i} = \alpha_i^e$. We define the *d-rank* of $\alpha_1, \ldots, \alpha_k$ to be the rank (over $\mathbb{K}$) of $M = M[d; \alpha_1, \ldots, \alpha_k]$.

Our main lemma relates the rank of a sequence to its "good"ness, and as we will see shortly, immediately yields Theorem 3.6.

Lemma 5.1. *Let $\mathbb{K}$ be a field of characteristic p. If the sequence $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$ is d-good for integer d, then the d-rank of $\alpha_1, \ldots, \alpha_k$ is at least $\text{wt}_p(d) + 1$. In particular, $k > \text{wt}_p(d)$.*

We note first that Theorem 3.6 follows from Lemma 5.1.

Proof of Theorem 3.6. Consider a (D, Λ)-UHD system with a pairwise-distinct solution $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$. We wish to show that $k > \text{wt}_p(d)$ for every $d \in D$. We note that $\text{shadow}_p(d) \subseteq D$. It follows that $\alpha_1, \ldots, \alpha_k$ is a d-good sequence, since we have $\sum_{i=1}^{k} \lambda_i \alpha_i^e = 0$ for every $e \in \text{shadow}_p(d) \subseteq D$. From Lemma 5.1 it follows that $k > \text{wt}_p(d)$ as desired. □

We now prove Lemma 5.1

Proof of Lemma 5.1. We prove the lemma by induction on the weight of d. We show that if the lemma is not true for some d, then it is also not true for some integer of the form $d' = d - e$, where $e \in \text{shadow}_p(d)$. Since $\text{wt}_p(d') < \text{wt}_p(d)$ this allows for an inductive argument based on weight. The key to establishing the inductive step is showing that the d'-rank of $\alpha_1, \ldots, \alpha_k$ is actualler smaller than its d-rank and this is done in Claim 5.2 below. In turn, the key idea behind the proof of Claim 5.2 is that for every $e \in \text{shadow}_p(d)$ the operation of "raising to power e" commutes with addition, at least when we only consider summing up $\alpha_1, \ldots, \alpha_k$. This strange commutativity implies a multitude of equality constraints imposed on α. The tricky part is choosing the correct ones to work with and manipulate to obtain our claim.

For the base case, we consider the case where the p-weight of d is 1. In this case the matrix M contains as a submatrix the matrix

$$\begin{bmatrix} 1 & 1 \\ \alpha_1^d & \alpha_2^d \end{bmatrix}.$$

The first row comes from the 0th powers of the α's and the second row from the dth powers. This is a rank 2 matrix (note that $\alpha_1^d \neq \alpha_2^d$ for any d of p-weight 1) and we conclude that d-rank of $\langle \alpha_1, \ldots, \alpha_k \rangle$ is at least 2 as desired.

We now move to the inductive step. To get some intuition into this part, suppose we knew that the d-rank of $\alpha_1, \ldots, \alpha_k$ is exactly $k - 1$, and that this held with $\lambda_1 = \cdots = \lambda_{k-1} = -\lambda_k = 1$. Further, suppose $\text{shadow}_p(d)$ includes 1. Then, first note that $\alpha_k = \sum_{i<k} \alpha_i$. Also, for any e in the p-shadow of d, and any d' in the p-shadow of $d - e$, we have $\sum_{i<k} \alpha_i^{d'+e} = \alpha_k^{d'+e} = \alpha_k^{d'} \alpha_k^e = (\sum_{i<k} \alpha_i^{d'}) \cdot (\sum_{j<k} \alpha_j^e)$. Rearranging terms, this turns into a linear dependence among $\alpha_1^{d'}, \ldots, \alpha_{k-1}^{d'}$ where the coefficients are independent of d', and this holds for every d' in the shadow of $d - e$. Assuming this is a non-trivial linear dependence (and we do verify that this holds in the formal analysis below) this shows that the $(d - e)$-rank of $\alpha_1, \ldots, \alpha_k$ is at most $k - 2$ (and so we are making progress). In the analysis below, we have to work a little harder (with

messier notation, and extra care) to deal with the more general case of arbitrary rank $\leq k$. We now give the formal details.

Assume for contradiction that the lemma is not true for $\alpha = (\alpha_1, \ldots, \alpha_k)$ and let d be the smallest integer showing this contradiction, i.e., for which $\alpha_1, \ldots, \alpha_k$ are d-good, but the d-rank of $\alpha_1, \ldots, \alpha_k$ is smaller than $1 + \mathrm{wt}_p(d)$. Note that $\mathrm{wt}_p(d) \geq 2$ (since we already ruled out the case $\mathrm{wt}_p(d) = 1$ above).

Let $\lambda_1, \ldots, \lambda_k$ be such that $\sum_{i=1}^{k} \lambda_i \alpha_i^{d'} = 0$ for all $d' \in \mathrm{shadow}_p(d)$.

Let $e \in \mathrm{shadow}_p(d)$ have $\mathrm{wt}_p(e) = 1$. Notice that α is $(d-e)$-good because $\mathrm{shadow}_p(d-e) \subset \mathrm{shadow}_p(d)$. We prove below that the $(d-e)$-rank of a subset (possibly containing all) of $\alpha_1, \ldots, \alpha_k$ is less than $\mathrm{wt}_p(d)$, contradicting the minimality of k and d.

Let the d-rank of $\alpha_1, \ldots, \alpha_k$ be ℓ and assume without loss of generality that the following conditions hold:

1. All λ_i's are nonzero (otherwise remove from α those elements such that $\lambda_i = 0$ to obtain a d-good set whose d-rank is no bigger than that of α).

2. $\alpha_k \neq 0$ (note we use the convention $0^0 = 1$ and so the column corresponding to some $\alpha_i = 0$ is not all zero) and the last $k - \ell$ columns of the matrix $M[d; \alpha]$ are linearly dependent on the first ℓ columns (by reordering). Note that this implies that there exist $A_{i,j} \in \mathbb{K}$, for every $i \in [k]$ and $j \in [\ell]$ such that $\alpha_i^{d'} = \sum_{j=1}^{\ell} A_{i,j} \alpha_j^{d'}$ for every $d' \in \mathrm{shadow}_p(d)$.

3. $\lambda_k = -1$ (which can be achieved by dividing all the λ_i's by $-\lambda_k$ which is non-zero.

We show below that there is a (non-trivial) linear dependence among the first ℓ columns of $M[d; \alpha]$, when restricted to rows corresponding to $d' \in \mathrm{shadow}_p(d-e)$, and thus the $(d-e)$-rank of $\alpha_1, \ldots, \alpha_k$ is at most $\ell - 1$ (recall α is also $(d-e)$-good).

Claim 5.2. *For every $d' \in \mathrm{shadow}_p(d-e)$ we have $\sum_{j=1}^{\ell} \tau_j \alpha_j^{d'} = 0$ where $\tau_j = (\alpha_k^e - \alpha_j^e) \cdot \sum_{i=1}^{k-1} \lambda_i A_{ij}$. Furthermore τ_j's are not all zero.*

Before proving the claim let us complete the proof of the lemma. The sequence $\tau = (\tau_1, \ldots, \tau_\ell)$ shows a nontrivial linear dependence among the first ℓ columns of $M[d-e; \alpha]$. Thus, the $(d-e)$-rank of α is strictly less than the d-rank of α. But $\mathrm{wt}_p(d-e) = \mathrm{wt}_p(d) - 1$, so we conclude that α is $(d-e)$-good but the $(d-e)$-rank of α is smaller than $1 + \mathrm{wt}_p(d-e)$, contradicting the minimality of d and completing the proof of Lemma 3.6 (assuming Claim 5.2). □

We now prove Claim 5.2. We follow the steps described in the intuitive description of the proof of Lemma 5.1, with care to ensure that we allow the rank to be general, and the dependencies among the columns of M to be arbitrary.

Proof of Claim 5.2. We start by noticing that, since $d' + e \in \mathrm{shadow}_p(d)$, we have

$$\alpha_k^{d'+e} = \sum_{i=1}^{k-1} \lambda_i \alpha_i^{d'+e}. \qquad (2)$$

But the left-hand side can be expanded as follows:

$$\begin{aligned}
\alpha_k^{d'+e} &= \alpha_k^e \cdot \alpha_k^{d'} = \alpha_k^e \cdot \sum_{i=1}^{k-1} \lambda_i \alpha_i^{d'} \text{ (Since } d' \in \mathrm{shadow}_p(d)\text{)} \\
&= \alpha_k^e \cdot \sum_{i=1}^{k-1} \lambda_i \sum_{j=1}^{\ell} A_{ij} \alpha_j^{d'} \\
&= \sum_{j=1}^{\ell} \left(\alpha_k^e \cdot \sum_{i=1}^{k-1} \lambda_i A_{ij} \right) \cdot \alpha_j^{d'} \text{ (Rearranging)} \qquad (3)
\end{aligned}$$

Next we work on the right had side of Equation (2). We have:

$$\begin{aligned}
\sum_{i=1}^{k-1} \lambda_i \alpha_i^{d'+e} &= \sum_{i=1}^{k-1} \lambda_i \cdot \sum_{j=1}^{\ell} A_{ij} \alpha_j^{d'+e} \\
&= \sum_{j=1}^{\ell} \left(\alpha_j^e \cdot \sum_{i=1}^{k-1} \lambda_i A_{ij} \right) \alpha_j^{d'}. \qquad (4)
\end{aligned}$$

Combining Equations (2), (3), and (4), we get

$$\sum_{j=1}^{\ell} \left((\alpha_k^e - \alpha_j^e) \cdot \sum_{i=1}^{k-1} \lambda_i A_{ij} \right) \cdot \alpha_j^{d'} = 0,$$

as claimed.

We now show that the τ_j's are not all zero. Let $\gamma_j = \sum_{i=1}^{k-1} \lambda_i A_{ij}$ and so $\tau_j = (\alpha_k^e - \alpha_j^e) \cdot \gamma_j$. We first note that for every $j \in \{1, \ldots, \ell\}$,

$$\alpha_k^e - \alpha_j^e = (\alpha_k - \alpha_j)^e \neq 0.$$

The first equation follows since $\mathrm{wt}_p(e) = 1$ which implies $e = p^c$ for some integer c and the second holds because of the distinctness of the α_i's. We conclude that the τ_j's are all zero only if the γ_j's are all zero. But if the γ_j's are all zero, we have

$$\alpha_k^e = \sum_{i=1}^{k-1} \lambda_i \alpha_i^e = \sum_{i=1}^{k-1} \sum_{j=1}^{\ell} \lambda_i A_{ij} \alpha_j^e = \sum_{j=1}^{\ell} \gamma_j \alpha_j^e = 0,$$

which contradicts our assumption that $\alpha_k \neq 0$. This proves the claim. □

6 Characterizing the zeroes of shadow-closed UHD systems

We show below that the main technical theorem, Theorem 3.6, allows us to characterize all the zeroes of the class of UHD systems it is able to consider, in terms of the α's that are distinct and not. It does so in terms of some natural partitions of the set $[k]$. Recall that a partition $\Pi = \{\pi_1, \ldots, \pi_\ell\}$ of $[k]$ is simply a collection of disjoint subsets of $[k]$ whose union equals $[k]$.

A sequence $\alpha_1, \ldots, \alpha_k \in \mathbb{K}$ induces a natural partition of $[k]$ based on equality of α_i's. Formally, let $\pi_{\mathrm{eq}}(\alpha_1, \ldots, \alpha_k) = \{\pi_1, \ldots, \pi_\ell\}$ if there exist distinct $\beta_1, \ldots, \beta_\ell \in \mathbb{K}$ such that for every $j \in [\ell]$ and $i \in \pi_j$, it is the case that $\alpha_i = \beta_j$.

For a sequence $\Lambda = \langle \lambda_1, \ldots, \lambda_k \rangle$ let the "zero-sum" partition family be the set of partitions $P(\Lambda) = \{$ Partitions $\pi = \{\pi_1, \ldots, \pi_\ell\}$ of $[k] | \forall j \in [\ell] \sum_{i \in \pi_j} \lambda_i = 0\}$. (I.e., in each $\pi \in P(\Lambda)$ the λ's in each partitioned subset sum to zero.)

Theorem 6.1. *Let $\mathbb{F} = \mathbb{F}_{p^r}$ and let $\mathbb{K}$ be any extension of $\mathbb{F}$. Let D be a p-shadow-closed subset of integers with an element $d \in D$ of p-weight greater than k. Let $\Lambda = \langle \lambda_1, \ldots, \lambda_k \rangle \in \mathbb{F}^k$. Then $\alpha_1, \ldots, \alpha_k$ is a solution of the (D, Λ)-UHD if and only if $\pi_{\mathrm{eq}}(\alpha_1, \ldots, \alpha_k) \in P(\Lambda)$.*

Proof. For any assignment $\alpha_1, \ldots, \alpha_k$, let $\pi = \pi_{\mathrm{eq}}(\alpha_1, \ldots, \alpha_k) = \{\pi_1, \ldots, \pi_\ell\}$ with $\beta_1, \ldots, \beta_\ell$ being such that $\alpha_i = \beta_j$ for every $j \in [\ell]$ and $i \in \pi_j$. Let $\gamma_j = \sum_{i \in \pi_j} \lambda_i$ and $\Gamma = \langle \gamma_1, \ldots, \gamma_\ell \rangle$. Then we note that $\alpha_1, \ldots, \alpha_k$ is a solution to the (D, Λ)-UHD system if and only if $\beta_1, \ldots, \beta_\ell$ is a solution to the (D, Γ)-UHD system. Furthermore any solution $\beta_1, \ldots, \beta_\ell$ would be a pairwise-distinct solution, by definition.

If $\pi \in P(\Lambda)$, then $\Gamma = 0$ and every assignment is a solution, else by Theorem 2.9 $\beta_1, \ldots, \beta_\ell$ is not a solution and so neither is $\alpha_1, \ldots, \alpha_k$. □

References

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.

Noga Alon, Eldar Fischer, Ilan Newman, and Asaf Shapira. A combinatorial characterization of the testable graph properties: it's all about regularity. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 251–260. ACM, 2006. ISBN 1-59593-134-1.

Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, January 1998.

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.

Eli Ben-Sasson and Madhu Sudan. Simple PCPs with poly-log rate and query complexity. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 266–275. ACM, 2005. ISBN 1-58113-960-8. URL http://doi.acm.org/10.1145/1060590.1060631.

Eli Ben-Sasson, Prahladh Harsha, and Sofya Raskhodnikova. Some 3CNF properties are hard to test. *SIAM J. Comput*, 35(1):1–21, 2005. URL http://epubs.siam.org/SICOMP/volume-35/art_44544.html.

Eli Ben-Sasson, Venkatesan Guruswami, Tali Kaufman, Madhu Sudan, and Michael Viderman. Locally testable codes require redundant testers. In *CCC '09: Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity*, pages 52–61, Washington, DC, USA, 2009. IEEE Computer Society. ISBN 978-0-7695-3717-7. doi: http://dx.doi.org/10.1109/CCC.2009.6.

Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of reed-muller codes. *CoRR*, abs/0910.0641, 2009.

Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. In *STOC*, pages 73–83. ACM, 1990.

Christian Borgs, Jennifer T. Chayes, László Lovász, Vera T. Sós, Balázs Szegedy, and Katalin Vesztergombi. Graph limits and parameter testing. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 261–270. ACM, 2006. ISBN 1-59593-134-1.

Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12:1–12:44, June 2007. ISSN 0004-5411.

Oded Goldreich and Madhu Sudan. Locally testable codes and pcps of almost-linear length. *J. ACM*, 53(4): 558–655, 2006.

Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998.

Elena Grigorescu, Tali Kaufman, and Madhu Sudan. 2-transitivity is insufficient for local testability. In *IEEE Conference on Computational Complexity*, pages 259–267. IEEE Computer Society, 2008. ISBN 978-0-7695-3169-4.

Elena Grigorescu, Tali Kaufman, and Madhu Sudan. Succinct representation of codes with applications to testing. In Irit Dinur, Klaus Jansen, Joseph Naor, and José D. P. Rolim, editors, *APPROX-RANDOM*, volume 5687 of *Lecture Notes in Computer Science*, pages 534–547. Springer, 2009. ISBN 978-3-642-03684-2.

Charanjit S. Jutla, Anindya C. Patthak, Atri Rudra, and David Zuckerman. Testing low-degree polynomials over prime fields. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 423–432. IEEE Computer Society, 2004. ISBN 0-7695-2228-9.

Tali Kaufman and Dana Ron. Testing polynomials over general fields. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 413–422. IEEE Computer Society, 2004. ISBN 0-7695-2228-9.

Tali Kaufman and Madhu Sudan. Algebraic property testing: the role of invariance. In Richard E. Ladner and Cynthia Dwork, editors, *STOC*, pages 403–412. ACM, 2008. ISBN 978-1-60558-047-0.

Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.*, 25(2):252–271, 1996.

Alex Samorodnitsky. Low-degree tests at large distances. In David S. Johnson and Uriel Feige, editors, *STOC*, pages 506–515. ACM, 2007. ISBN 978-1-59593-631-8.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de