Some closure features of locally testable affine-invariant properties

by

Alan Xinyu Guo

B.S. in Mathematics, Duke University (2011)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2013

Author . . . .
Department of Electrical Engineering and Computer Science
December 12, 2012

Certified by . . . .
Madhu Sudan
Adjunct Professor
Thesis Supervisor

Accepted by . . . .
Leslie Kolodziejski
Chairman, Department Committee on Graduate Students

Some closure features of locally testable affine-invariant properties

by

Alan Xinyu Guo

Submitted to the Department of Electrical Engineering and Computer Science
on December 12, 2012, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

We prove that the class of locally testable affine-invariant properties is closed under sums, intersections and "lifts". The sum and intersection are two natural operations on linear spaces of functions, where the sum of two properties is simply their sum as a vector space. The "lift" is a less well-studied property, which creates some interesting affine-invariant properties over large domains, from properties over smaller domains.

Previously such results were known for "single-orbit characterized" affine-invariant properties, which are known to be a subclass of locally testable ones, and are potentially a strict subclass. The fact that the intersection of locally-testable affine-invariant properties are locally testable could have been derived from previously known general results on closure of property testing under set-theoretic operations, but was not explicitly observed before. The closure under sum and lifts is implied by an affirmative answer to a central question attempting to characterize locally testable affine-invariant properties, but the status of that question remains wide open.

Affine-invariant properties are clean abstractions of commonly studied, and extensively used, algebraic properties such linearity and low-degree. Thus far it is not known what makes affine-invariant properties locally testable — no characterizations are known, and till this work it was not clear if they satisfied any closure properties. This work shows that the class of locally testable affine-invariant properties are closed under some very natural operations. Our techniques use ones previously developed for the study of "single-orbit characterized" properties, but manage to apply them to the potentially more general class of all locally testable ones via a simple connection that may be of broad interest in the study of affine-invariant properties.

Thesis Supervisor: Madhu Sudan
Title: Adjunct Professor

Acknowledgments

This thesis is a result of joint work with Madhu Sudan. I am deeply grateful to Madhu, my advisor, for the collaboration, as well as for his support, advice, and encouragement, his help with technical matters, and many lively and enlightening discussions. Special thanks to Greg Aloupis, Andrea Campagna, Erik Demaine, Swastik Kopparty, Ronitt Rubinfeld, and Giovanni Viglietta for the collaborations on other projects. Many thanks to Eli Ben-Sasson, Henry Cohn, Nadia Heninger, Piotr Indyk, and Yohay Kaplan for many interesting and enlightening discussions, and to my friends and colleagues at MIT, especially Eric Blais, Adam Bouland, Mohammad Bavarian, Matt Coudron, Ioana Ivan, Sepideh Mahabadi, Ludwig Schmidt, Aaron Sidford, Madars Virza, Adrian Vladu, and Henry Yuen and many others for their good company and entertaining conversations, and for creating a warm and friendly research environment. Thanks to Kwan Li for being a good friend and roommate and making my life at home enjoyable. Most of all, thanks to Mom, Dad, Julia, and Lisa for their endless encouragement, love, and support.

This work was supported by NSF grants CCF-0829672, CCF-1065125, CCF-6922462, MIT CSAIL, and an NSF Graduate Research Fellowship.

Contents

Chapter 1

Introduction

In this work we investigate the closure of the class of locally-testable affine-invariant (linear) properties under some natural operations. We define these notions below and then give some motivation for our investigation.

1.1 Main terms and results

Throughout this work $\mathbb{F}_q$ will denote the finite field consisting of q elements. We consider properties of functions mapping a big field $\mathbb{F}_{q^n}$ (for growing n) to a small field $\mathbb{F}_q$. Denoting all functions mapping $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$ by $\{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$, a property is given by a family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$, which is the family of functions that satisfy the property. Throughout the discussion below, $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$. Throughout this paper, we will consider only *linear properties*, where $\mathcal{F}$ is said to be linear if for every $\alpha, \beta \in \mathbb{F}_q$ and $f, g \in \mathcal{F}$, the function $\alpha f + \beta g$ is also in $\mathcal{F}$ (where $(\alpha f + \beta g)(x) = \alpha f(x) + \beta g(x)$).

A property $\mathcal{F}$ is said to be *affine-invariant* if $\mathcal{F}$ is invariant under affine permutations of the domain as elaborated below. A map $A : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$ is said to be an *affine permutation* if there exist $\alpha, \beta \in \mathbb{F}_{q^n}$ with $\alpha \neq 0$ such that $A(x) = \alpha x + \beta$ for every $x \in \mathbb{F}_{q^n}$. (We often drop "permutation" and often simply refer to A as affine map.) For $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ and affine A, let $f \circ A : \mathbb{F}_{q^n} \to \mathbb{F}_q$ be the function $(f \circ A)(x) = f(A(x))$. $\mathcal{F}$ is said to be affine-invariant if for every affine A, $f \in \mathcal{F} \Rightarrow f \circ A \in \mathcal{F}$.

Property testers for a property $\mathcal{F}$ aim to estimate the distance between a given function $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ and a property $\mathcal{F}$. We formalize this concept below, starting with the notion of distance. For functions $f, g : \mathbb{F}_{q^n} \to \mathbb{F}_q$, the distance between f and g, denoted $\delta(f, g) = \Pr_{x \leftarrow_U \mathbb{F}_{q^n}}[f(x) \neq g(x)]$, where the notation $x \leftarrow_U \mathbb{F}_{q^n}$ denotes x chosen uniformly at random from $\mathbb{F}_{q^n}$. We define $\delta(f, \mathcal{F})$ to be $\min_{g \in \mathcal{F}}\{\delta(f, g)\}$. We say f is δ-close to $\mathcal{F}$ if $\delta(f, \mathcal{F}) \leq \delta$ and δ-far otherwise.

A property $\mathcal{F}$ is said to be (k, ϵ) *locally-testable* if there exists a probabilistic algorithm with oracle access to a function $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ that makes k queries to the oracle f and accepts with probability 1 if $f \in \mathcal{F}$ and rejects all f with probability at least $\epsilon \cdot \delta(f, \mathcal{F})$.

Our interest in this work is in an ensemble of properties $\mathcal{F}_n \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ for infinitely many n that are testable with some fixed parameters $k < \infty$ and $\epsilon > 0$ for every n. If such k and ϵ exist we will refer to these properties as simply locally testable.

Our main results show that locally testable affine-invariant properties are closed under some basic operations.

The first operation we consider is the intersection. Given $\mathcal{F}_1, \mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$, $\mathcal{F}_1 \cap \mathcal{F}_2$ is just the set of functions satisfying both the properties. In Theorem 4.3 we prove that the class of locally testable affine-invariant properties is closed under intersection, i.e., if $\mathcal{F}_1$ and $\mathcal{F}_2$ are locally testable, then so is $\mathcal{F}_1 \cap \mathcal{F}_2$. We note that this result also follows from the general study of the closure of property testing under set-theoretic operations by Chen et al. [8, Proposition 2] who show (roughly) that $\mathcal{F}_1 \cap \mathcal{F}_2$ is locally testable if $\mathcal{F}_1 \cup \mathcal{F}_2$ is contained in an error-correcting. The fact that the hypothesis holds follows from a result of Ben-Sasson and Sudan [6], but this combination of observations does not seem to have been made before.

The second operation is almost as natural in the context of linear properties. For $\mathcal{F}_1, \mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$, their sum, denoted $\mathcal{F}_1 + \mathcal{F}_2$, is the property $\{f_1 + f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}$. In Theorem 4.8 we show that if $\mathcal{F}_1$ and $\mathcal{F}_2$ are locally testable, then so is $\mathcal{F}_1 + \mathcal{F}_2$.

The final operation we consider is a unary one. Given a property $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ and positive integer ℓ, there is a unique natural affine-invariant property $\mathcal{F}' \subseteq \{\mathbb{F}_{q^{n\ell}} \to \mathbb{F}_q\}$ that extends $\mathcal{F}$. (A formal definition is given in Chapter 4.) This property $\mathcal{F}'$ is called the *ℓ-lift* of $\mathcal{F}$ and denoted $\text{Lift}_\ell(\mathcal{F})$. In Theorem 4.15 we show the lifts of locally testable properties are also locally testable.

We now describe the reason to study affine-invariant properties and their closures.

1.2 Motivation

We start by reiterating the case for the study of affine-invariant properties briefly. (This case has already been made in many of the previous works and surveys [16, 10, 11, 6, 4, 2, 17].) Affine-invariance is the natural abstraction of a very important class of properties that have proven to be of central interest in complexity theory. Namely they abstract the property of being linear, and/or low-degree, with the feature that they offer the ability to preserve the efficiency of the proofs and techniques in this area. Finally, they offer the potential for new constructions of locally testable codes (and potentially PCPs), though such possibility would need much better understanding of the testability of affine-invariant properties.

The study of what makes an affine-invariant property locally testable is still in its early stages. We are still far from getting an exact characterization of when such properties may be tested with a constant number of queries, and the work of [2] poses many questions that remain open that need to be resolved to reach such a goal. The question as to what operations preserve testability is among the basic questions one can ask to gain understanding of testability. In the case of general property testing, the seminal work of Goldreich, Goldwasser and Ron [9] explored this question, but derived mostly negative results. The work of Chen et al. [8] explored this question under the condition that the properties were "code-like" and managed to get some positive results, under technical conditions. (As mentioned earlier, these results imply that the intersection of testable affine-invariant properties is testable, though to get this implication one needs to invoke some of the structural aspects of affine-invariant properties.) The work of Ben-Sasson et al. [2] studied this question for a restricted class of testable properties (called "single-orbit characterized" properties) which we will discuss below. Our work settles the closure questions positively, unconditionally, for all locally testable affine-invariant properties and thus represents progress towards the broad goal of understanding what makes an affine-invariant property locally testable.

1.3 Single-orbit characterized properties

Most previous works on local testability have focussed on a special route to local-testability via what are termed "single-orbit-characterizations". Single-orbit characterizations go to the heart of the most commonly studied locally testable affine invariant properties. These are properties characterized by a single local "constraint" and the feature of being affine-invariant (a k-(local)-constraint looks at the value of a function at some k values and restricts the values in some way). Canonical examples include the fact that a multivariate function f is of degree d if and only if its restriction to the first coordinate axis is of degree d and the function is invariant under affine transformation of $\mathbb{F}_q^n$ (the n dimensional vector space over $\mathbb{F}_q$).

It is known that the single-orbit characterized properties (of a local constraint) are locally testable [16]. All known locally testable properties are also known to be single-orbit characterized [2]. Motivated by these considerations Ben-Sasson et al. [2] studied the closure of single-orbit characterized properties under intersection, sums, and lifts and showed that this class was closed under these operations.

They however left open the more general question of the closure of locally testable properties under these operations. To the best of our knowledge these two classes — locally testable properties, and single-orbit properties — may be identical, but even the truth of this statement (leave alone the ability to prove it) is wide open. Indeed one path to separate these classes would have been to show that the former class is not closed under one of these operations. Our work closes this possibility.

1.4 Technical contributions

The results in this work are obtained by simple combinations of known facts in the literature on testing affine-invariant properties. These facts already tell us that affine-invariant properties should be viewed via a basis of (traces of) monomials. The set of exponents of the monomials in the support of functions contained in an affine-invariant properties form the "degree set" of the property, and completely determine the property. In the reverse direction, it is known that not every degree set corresponds to an affine-invariant property, and the structure of what the degree set can look like for the property to be affine-invariant is completely understood.

Turning to local testability, among the known families of locally testable properties, their degree set is well understood. But for a generic locally testable property, it is still open as to what the degree set may look like. In the absence of such understanding it seemed this structural feature would offer little help in understanding local testability. This is where this work manages to improve the understanding.

Our main technical lemma manages to relate the performance of testers to the degree sets of the properties. Specifically it says that the "canonical local tester" of an affine-invariant property must behave nicely with respect to the monomials appearing in the degree set, and distinguish them from a small set of "excluded" monomials, which come from the complement of the degree set. The "canonical local tester" is one introduced in the work of [3] which shows that without loss of generality any linear property can be tested by picking a distribution over "low-weight linear constraints" satisfied by all functions with the property and testing that a randomly chosen one of these constraints holds. Our lemma says that every monomial in the degree set should also satisfy all these constraints, while every monomial in in the excluded set should fail to satisfy ϵ-fraction of these constraints.

While our lemma is simple to prove given the known results on testabilility of affine-invariant properties, the resulting understanding is valuable. Indeed, the closure properties follow quite easily from this lemma, since the behavior of the degree sets is well-understood under the operations in consideration.

Organization. In Chapter 2 we introduce some of the notation and background material from the study of affine-invariant properties. In Chapter 3 we state and prove the main technical result, Theorem 3.4, of this thesis relating degree sets to testability. In Chapter 4 we then prove the closure theorems using Theorem 3.4.

Chapter 2

Preliminaries

We use $[m]$ to denote the set $\{1, \ldots, m\}$. We start with some background material on constraints and (single-orbit) characterizations. We then describe the Reed-Muller property which is known to be locally testable, and to contain all locally testable affine-invariant properties.

2.1 Constraints and Characterizations

A *k-constraint* on functions mapping $\mathbb{F}_q^n$ to $\mathbb{F}_q$ is given by a pair $C = (\vec{\alpha}, \vec{\lambda})$ of k-tuples where $\vec{\alpha} = \langle \alpha_1, \ldots, \alpha_k \rangle \in \mathbb{F}_{q^n}^k$ and $\vec{\lambda} = \langle \lambda_1, \ldots, \lambda_k \rangle \in \mathbb{F}_q^k$. A function $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ is said to satisfy C if $\sum_{i=1}^k \lambda_i f(\alpha_i) = 0$. (Note that while the notion of satisfaction is intended to apply to functions mapping to $\mathbb{F}_q$, it extends also to functions mapping to $\mathbb{F}_{q^n}$ also, and we will need this extension in this thesis.)

A collection of k-constraints $C_1, \ldots, C_m$ k-characterizes a property $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ if the following holds:

$$\forall f : \mathbb{F}_{q^n} \to \mathbb{F}_q, \;\; f \in \mathcal{F} \Leftrightarrow \; \forall j \in [m], \; f \text{ satisfies } C_j.$$

For a k-constraint $C = (\vec{\alpha}, \vec{\lambda})$ on $\{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ and affine transform $A : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$, $C \circ A$ denotes the k-constraint $(A(\vec{\alpha}), \vec{\lambda})$ where $A(\vec{\alpha}) = \langle A(\alpha_1), \ldots, A(\alpha_k) \rangle$. The orbit of a constraint C is the collection of of constraints $\text{orb}(C) = \{C \circ A \mid A \text{ is an affine transform}\}$. Note that if $\mathcal{F}$ is affine-invariant and every member of $\mathcal{F}$ satisfies C, then every member of $\mathcal{F}$ satisfies every constraint in $\text{orb}(C)$.

We say that $\mathcal{F}$ is *k-single orbit characterized* if there exists a k-constraint C such that $\text{orb}(C)$ is a k-characterization of $\mathcal{F}$.

We use the following theorem showing that single-orbit characterized properties are locally testable.

Theorem 2.1 ([16]). *There exists a constant c such that for every prime power q and integers k, n, if $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ has a k-single orbit characterization, then it is $(k, 1/(ck^2))$-locally testable.*

We note that the test from [16] simply picks an affine transformation A uniformly at random and tests if f satisfies $C \circ A$, where C is the k-constraint giving the single-orbit characterization.

2.2 Reed-Muller Property

We first introduce some basic notions. For integer d, let $d_0, d_1, \ldots$ denote its expansion in base q, so that $0 \leq d_i \leq q-1$ and $d = \sum_i d_i q^i$. The q-weight of d, denote $q\text{-wt}(d)$, is the quantity $\sum_i d_i$. Recall that every function $f : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$ is uniquely expressible as a univariate polynomial of degree at most $q^n - 1$. For $f(x) = \sum_{i=0}^{q^n-1} c_i x^i$, we say that its support is the set of integers $\text{supp}(f) = \{i \mid c_i \neq 0\}$. Let $\text{Tr}(x) = x + x^q + x^{q^2} + \cdots + x^{q^{n-1}}$ denote the "trace" function, which is a linear map from $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$.

For integer d, the Reed-Muller property $\text{RM}[n, d, q] \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is the collection of all functions which are traces of polynomials supported on integers of weight at most d, i.e.,

$$\text{RM}[n, d, q] = \{\text{Tr}(p) \mid p \in \mathbb{F}_{q^n}[x] \text{ s.t. } \forall i \in \text{supp}(f), q\text{-wt}(i) \leq d\}.$$

We note that there are several alternate definitions of Reed-Muller properties, but none of these is relevant to us. The only aspects we care about are (1) the Reed-Muller property is an affine-invariant property that forms an error-correcting code for constant d, (2) the Reed-Muller property is locally testable for constant d, and (3) every affine-invariant property that admits a local constraint is contained in the Reed-Muller property. We give references below.

Proposition 2.2 (Folklore). *1. For every prime power q and positive integers n and d, the Reed-Muller property* $\mathrm{RM}[n,d,q]$ *is $\mathbb{F}_q$-linear and affine-invariant.*

2. For every prime power q and positive integer d there exists $\delta > 0$ such that for every n, the Reed-Muller property $\mathrm{RM}[n,d,q]$ *is a code of distance δ, i.e., for every pair of distinct functions $f, g \in \mathrm{RM}[n,d,q]$, $\delta(f,g) \geq \delta$.*

Theorem 2.3 ([15]). *For every prime power q and positive integer d there exists $k < \infty$ and $\epsilon > 0$ such that for every n, the Reed-Muller property* $\mathrm{RM}[n,d,q]$ *is (k,ϵ)-locally testable.*

We note that the study of testability of the Reed-Muller property was initiated by Alon et al. [1] who analyzed the case of $q = 2$. The case of prime q was proved independently by [15] and [14]. By now, improved analyses of the tests (with better k and ϵ) are also available (see [7, 13]).

Theorem 2.4 ([6]). *For every prime power q and integer k there exists an integer w such that for every n the following holds: Suppose $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is an affine-invariant linear property and C is a k-constraint satisfied by every member of $\mathcal{F}$. Then $\mathcal{F} \subseteq \mathrm{RM}[n,w,q]$.*

Chapter 3

ϵ-separators and local tests

In this section we introduce the notion of an ϵ-separating test, and prove a theorem relating the existence of a tester to the existence of such separating tests. This theorem will be employed repeatedly in Chapter 4 to get testers for various composed properties.

We start with a result of Ben-Sasson et al. [3] that shows that all testers for linear properties can be made "canonical", i.e., described by a collection of k-local "constraints" and a distribution over them. We describe their result first.

A *canonical k-test T* on functions mapping $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$ is given by a sequence of k-constraints $C_1, \ldots, C_m$ and a distribution D on $[m]$. To test a function f, the tester picks $j \leftarrow_D [m]$ and accepts if and only if f satisfies C_j.

Proposition 3.1 ([3]). *A linear family $\mathcal{F}$ is (k, ϵ)-locally testable if and only if there exists a canonical k-test T that accepts $f \in \mathcal{F}$ with probability 1, while rejecting $f : \mathbb{F}_{q^n} \to \mathbb{F}_q$ with probability at least $\epsilon \cdot \delta(f, \mathcal{F})$.*

Our notions will consider the performance of canonical tests on a certain selection of monomials (viewed as functions mapping $\mathbb{F}_{q^n}$ to $\mathbb{F}_{q^n}$).

Definition 3.2. *For sets $A \subseteq B \subseteq \{0, \ldots, q^n - 1\}$, we say that a k-canonical test $T = (C_1, \ldots, C_m; D)$ (for functions mapping $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$) ϵ-separates A from B if the following hold:*

Completeness: $\forall a \in A, \quad \Pr_{j \leftarrow_D [m]}[x^a \text{ satisfies } C_j] = 1$.

Soundness: $\forall b \in B - A, \quad \Pr_{j \leftarrow_D [m]}[x^b \text{ does not satisfy } C_j] \geq \epsilon$.

To identify sets appropriate for separation by canonical tests, we move to the structural aspects. It is by now well-known that an affine invariant family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ has an associated degree set $\mathrm{Deg}(\mathcal{F}) \subseteq \{0, \ldots, q^n - 1\}$, which uniquely specifies $\mathcal{F}$. Specifically, $\mathrm{Deg}(\mathcal{F}) = \cup_{f \in \mathcal{F}} \mathrm{supp}(f)$. The degree set of a family $\mathcal{F}$ is well-studied and the following lemma is an easy consequence of its well-known properties.

Lemma 3.3. *For every affine-invariant linear family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$, for every $d \in \mathrm{Deg}(\mathcal{F})$, and for every $\lambda \in \mathbb{F}_{q^n}$, the function $\mathrm{Tr}(\lambda \cdot x^d) \in \mathcal{F}$. Conversely, if $d \notin \mathrm{Deg}(\mathcal{F})$, then there exists $\lambda \in \mathbb{F}_{q^n}$ such that $\mathrm{Tr}(\lambda \cdot x^d) \notin \mathcal{F}$.*

For an affine-invariant family $\mathcal{F}$, let $\mathrm{wt}(\mathcal{F}) = \max_{d \in \mathrm{Deg}(\mathcal{F})} q\text{-wt}(d)$. Let $\text{RM-Deg}(\mathcal{F}) = \{d \in \{0, \ldots, q^n - 1\} \mid q\text{-wt}(d) \leq \mathrm{wt}(\mathcal{F}) + 1\}$. (The notation RM-Deg recalls the fact that the Reed-Muller family contains all degrees of q-weight bounded by w.) Our main result about testability of a family $\mathcal{F}$ is summarized below.

Theorem 3.4. *A linear affine-invariant family $\mathcal{F}$ is locally testable if and only if a canonical test separates $\mathrm{Deg}(\mathcal{F})$ from $\text{RM-Deg}(\mathcal{F})$. More precisely:*

$\Rightarrow$ *$\forall q, k, \epsilon > 0$, $\exists k' < \infty$ and $\epsilon' > 0$ such that $\forall n$ the following holds: If $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k, ϵ)-locally testable, then there exists a k'-canonical test ϵ'-separating $\mathrm{Deg}(\mathcal{F})$ from $\text{RM-Deg}(\mathcal{F})$.*

$\Leftarrow$ *$\forall q, k', \epsilon' > 0$, $\exists k < \infty$ and $\epsilon > 0$ such that $\forall n$ the following holds: If $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ has a k'-canonical test ϵ'-separating $\mathrm{Deg}(\mathcal{F})$ from $\text{RM-Deg}(\mathcal{F})$, then $\mathcal{F}$ is (k, ϵ)-locally testable.*

3.1 Proof of Theorem 3.4

We prove Theorem 3.4 in this section. We give a brief overview first. The forward direction is straightforward - any canonical local tester T for $\mathcal{F}$ gives a (k, ϵ)-canonical test separating

Deg($\mathcal{F}$) from RM-Deg($\mathcal{F}$), and the proof is almost immediate from definitions and basic properties of the trace function.

The reverse direction takes a few steps. To start with, it is not the case that the canonical test T_1 separating Deg($\mathcal{F}$) from RM-Deg($\mathcal{F}$) is itself a tester for $\mathcal{F}$ (or at least we do not know how to prove this). So we combine this test with a test for the Reed-Muller property corresponding to the degree set RM-Deg($\mathcal{F}$). Completeness of this test is immediate, but soundness takes some calculations. Roughly, if a function f is far from the Reed-Muller property, then the Reed-Muller test detects this with high probability. If f is very close to $\mathcal{F}$ but not contained in it, then also the Reed-Muller test rejects it with sufficiently high probability. The only remaining case is when f is close to the Reed-Muller family, but its closest codeword in the Reed-Muller property is a function $g \notin \mathcal{F}$. In this case, we note first that the function g is rejected by T_1 with high-probability (based on the soundness condition of canonical tests separating Deg($\mathcal{F}$) from RM-Deg($\mathcal{F}$)), and then argue that f, being close to g, is rejected with roughly the same probability. Details below.

Proof. Let $P = \mathrm{RM}[n, \mathrm{wt}(\mathcal{F}) + 1, q]$ and let $\delta_P = \delta(P)$ be the relative minimum distance of this code. (Note by Proposition 2.2 that δ_P does not depend on n.)

($\Rightarrow$) By Proposition 3.1, there exists a canonical k-test $T = (C_1, \dots, C_m; D)$ that accepts $f \in \mathcal{F}$ with probability 1, while rejecting $f \notin \mathcal{F}$ with probability at least $\epsilon \cdot \delta(f, \mathcal{F})$. We claim that T ϵ'-separates Deg($\mathcal{F}$) from RM-Deg($\mathcal{F}$) where $\epsilon' = \epsilon \cdot \delta_P$. Suppose $C_i = (\vec{\alpha_i}, \vec{\lambda_i})$ where $\vec{\alpha_i} = \langle \alpha_{i1}, \dots, \alpha_{ik} \rangle$ and $\vec{\lambda_i} = \langle \lambda_{i1}, \dots, \lambda_{ik} \rangle$.

Completeness: If $d \in \mathrm{Deg}(\mathcal{F})$, then $\mathrm{Tr}(\lambda x^d) \in \mathcal{F}$ for every $\lambda \in \mathbb{F}_{q^n}$ (from Lemma 3.3), and so

$$0 = \sum_{j=1}^{k} \lambda_{ij} \,\mathrm{Tr}(\lambda \alpha_{ij}^d) = \mathrm{Tr}\left(\lambda \sum_{j=1}^{k} \lambda_{ij} \alpha_{ij}^d\right)$$

for every $i \in [m]$ and $\lambda \in \mathbb{F}_{q^n}$, which implies that $\sum_{j=1}^{k} \lambda_{ij}\alpha_{ij}^d = 0 \,\forall i$, i.e. $\Pr_{i \leftarrow_D [m]}[x^d \text{ satisfies } C_i] = 1$.

Soundness: If $e \in \mathrm{RM\text{-}Deg}(\mathcal{F}) - \mathrm{Deg}(\mathcal{F})$, then there exists $\lambda \in \mathbb{F}_{q^n}$ such that $\mathrm{Tr}(\lambda x^e) \in P -$

$\mathcal{F}$. In particular, $\mathrm{Tr}(\lambda x^e)$ is a codeword of P and $\mathcal{F}$ is a subcode of P, so $\delta(\mathrm{Tr}(\lambda x^e), \mathcal{F}) \geq \delta_P$. If x^e satisfies C_i, then so does $\mathrm{Tr}(\lambda x^e)$ since

$$\sum_{j=1}^{k} \lambda_{ij} \,\mathrm{Tr}(\lambda \alpha_{ij}^e) = \mathrm{Tr}\left(\lambda \sum_{j=1}^{k} \lambda_{ij} \alpha_{ij}^e\right) = 0$$

and so $\Pr_{i \leftarrow_D [m]}[x^e \text{ satisfies } C_i] \leq \Pr_{i \leftarrow_D [m]}[\mathrm{Tr}(\lambda x^e) \text{ satisfies } C_i] \leq 1 - \epsilon \cdot \delta(\mathrm{Tr}(\lambda x^e), \mathcal{F}) \leq 1 - \epsilon \cdot \delta_P$.

($\Leftarrow$) We prove this direction in two steps. We first prove that there is a k_1-local test T_1 that accepts all $f \in \mathcal{F}$ while rejecting $f \in P - \mathcal{F}$ with probability at least ϵ_1. We then prove that T_1 can be combined with a tester for membership in P to get a tester for the family $\mathcal{F}$. We start with a description of the test T_1 and its analysis.

The test T_1: Let $(C_1, \ldots, C_m; D)$ be a k'-canonical test ϵ'-seperating $\mathrm{Deg}(\mathcal{F})$ from $\text{RM-Deg}(\mathcal{F})$. Our test T_1 consists of picking $i \leftarrow_D [m]$, and picking an affine transformation $A : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$ uniformly at random, and checking if f satisfies $C_i \circ A$.

To see the completeness condition is met, note that $C_i \circ A$ accepts f if and only if C_i accepts $f \circ A$. Since $f \circ A \in \mathcal{F}$, it follows that every monomial in the support of $f \circ A$ is accepted by C_i and so C_i also accepts f. We thus conclude that T_1 accepts $f \in \mathcal{F}$ with probability 1.

Now, we analyze the soundness.

Let w be the weight given by Theorem 2.4 so that every property satisfying some k'-constraint is contained in $\mathrm{RM}[n, w, q]$. Let δ_0 be the distance of $\mathrm{RM}[n, w, q]$ from Proposition 2.2. Let $\epsilon_0 = \delta_0/(c(k')^2)$ where c is the constant from Theorem 2.1. We will show below that T_1 rejects every member $f \in P - \mathcal{F}$ with probability at least $\epsilon_1 = \epsilon' \cdot \epsilon_0$. Note that all the constants are indeed independent of n, as desired.

Fix $f \in P - \mathcal{F}$. There must exist $e \in \text{RM-Deg}(\mathcal{F}) - \mathrm{Deg}(\mathcal{F})$ such that $e \in \mathrm{supp}(f)$. With probability at least ϵ' our choice of $i \leftarrow_D [m]$ will be such that x^e does not satisfy C_i. We show below that for every i such that x^e does not satisfy C_i the probability, over the choice of A, that $C_i \circ A$ rejects f is at least ϵ_0, which yields the desired soundness.

Let $\mathcal{F}'$ be the family of functions that satisfy $C_i \circ A'$ for every affine transformations A'. Then $\mathcal{F}'$ has a k'-single orbit characterization, given by $\mathrm{orb}(C_i)$. Since $f \notin \mathcal{F}'$, by Theorem 2.1, the test consisting of randomly choosing A' and accepting if and only if f satisfies $C_i \circ A'$ rejects all f with probability at least $\delta(f, \mathcal{F}')/(c(k')^2)$. Since $\{f\} \cup \mathcal{F}' \subseteq \mathrm{RM}[n, w, q]$, it follows that $\delta(f, \mathcal{F}') \geq \delta_0$. We thus conclude that $C_i \circ A$ rejects f with probability at least $\delta_0/(c(k')^2) = \epsilon_0$. Combining with the probability that i is such that C_i rejects x^e, we get that T_1 rejects every $f \in P - \mathcal{F}$ with probability at least $\epsilon' \cdot \epsilon_0$.

Tester for $\mathcal{F}$. We now turn to using T_1 to build a tester for $\mathcal{F}$. Let T_2 be a (k_2, ϵ_2)-local test for P, as guaranteed by Theorem 2.3. Our tester T for $\mathcal{F}$ works as follows: With probability $1/2$ it runs T_1 and accepts if T_1 accepts, and with probability $1/2$ it runs T_2 and accepts if T_2 accepts.

We now analyze the test. The completeness is obvious: If $f \in \mathcal{F}$, then both T_1 and T_2 accept with probability one and so T accepts with probability one. So we turn below to the soundness.

If $f \in P - \mathcal{F}$, then the probability that T rejects is at least half the probability that T_1 rejects, and so T rejects with probability at least $\epsilon_1/2$. Now consider the case where $f \notin P$. Let $\delta(f, \mathcal{F}) = \delta$ and $\delta(f, P) = \delta_1$. Note that $\delta_1 \leq \delta$. If $\delta < \frac{\delta_P}{2}$, then the nearest codeword in P to f is actually in $\mathcal{F}$, hence $\delta_1 = \delta$. In this case, T_2 rejects with probability at least $\epsilon_2 \cdot \delta$ and so T rejects with probability at least $\epsilon_2 \cdot \delta/2$. Otherwise, there is some $g \in P - \mathcal{F}$ such that $\delta(f, g) = \delta_1$. The probability that T_1 rejects f is at least $\epsilon_1 - k_1\delta_1$, since T_1 rejects g with probability at least ϵ_1 and the probability that f disagrees with g on one of the k_1 queries made by T_1 is at most $k_1\delta_1$. On the other hand, T_2 rejects f with probability at least $\epsilon_2 \cdot \delta_1$. Therefore, in this case T rejects with probability at least $(\epsilon_2 \cdot \delta_1 + \epsilon_1 - k_1\delta_1)/2 \geq \frac{\epsilon_1\epsilon_2}{2(k_1+\epsilon_2)}$. Putting this together with the case that $\delta < \frac{\delta_P}{2}$, our test T rejects $f \notin P$ with probability at least $\min\{\epsilon_2 \cdot \delta/2, \frac{\epsilon_1\epsilon_2}{2(k_1+\epsilon_2)}\}$. Putting this together with the case that $f \in P - \mathcal{F}$ and noting that $\frac{\epsilon_1\epsilon_2}{2(k_1+\epsilon_2)} \leq \epsilon_1/2$, we get a $\max\{k_1, k_2\}$-local test T that rejects $f \notin \mathcal{F}$ with probability at least $\min\{\epsilon_2 \cdot \delta/2, \frac{\epsilon_1\epsilon_2}{2(k_1+\epsilon_2)}\}$.

Putting it all together, we get a (k, ϵ)-local test T for $\mathcal{F}$ where $k = \max\{k', k_2\}$ and $\epsilon = C \cdot \delta(f, \mathcal{F})$ where both k_2 and C depend only on k'.

□

Chapter 4

Closure of locally testable properties

In this section we use our structural characterization of locally testable families (Theorem 3.4) to prove that the class of locally testable affine-invariant properties is closed under sums, intersections, and lifts. Our approach for each operation is the same. First, we examine how the degree sets of the original properties relate to the degree set of the sum, intersection, or lift. Next, we use this knowledge to construct a test which separates the degree set of the new property from its Reed-Muller degree set, using separating tests for the degree sets of the original properties. Finally, we apply Theorem 3.4 to immediately obtain local testability.

4.1 Closure under intersection

Proposition 4.1. *Let $\mathcal{F}_1, \mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ be affine-invariant properties. Then* $\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2) = \mathrm{Deg}(\mathcal{F}_1) \cap \mathrm{Deg}(\mathcal{F}_2)$.

Proof. Consider $d \in \mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$. Since $\mathrm{Tr}(\lambda x^d) \in \mathcal{F}_1 \cap \mathcal{F}_2$ for every λ, it follows that $d \in \mathrm{Deg}(\mathcal{F}_1) \cap \mathrm{Deg}(\mathcal{F}_2)$ and so $\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2) \subseteq \mathrm{Deg}(\mathcal{F}_1) \cap \mathrm{Deg}(\mathcal{F}_2)$. The reverse direction is similar. If $d \in \mathrm{Deg}(\mathcal{F}_1) \cap \mathrm{Deg}(\mathcal{F}_2)$ then $\mathrm{Tr}(\lambda x^d) \in \mathcal{F}_1 \cap \mathcal{F}_2$ for every λ and so $d \in \mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$. □

Lemma 4.2. *For $i \in \{1, 2\}$ if there exist k-canonical tests ϵ-separating* $\mathrm{Deg}(\mathcal{F}_i)$ *from* $\text{RM-Deg}(\mathcal{F}_i)$*, then there is a k-canonical test $\epsilon/2$-separating* $\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$ *from* $\text{RM-Deg}(\mathcal{F}_1 \cap$

$\mathcal{F}_2)$.

Proof. For $i \in \{1,2\}$ let $(C_1^{(i)}, \ldots, C_m^{(i)}; D^{(i)})$ be the k-canonical tests that ϵ-separate $\mathrm{Deg}(\mathcal{F}_i)$ from RM-$\mathrm{Deg}(\mathcal{F}_i)$. Then we claim that the natural test which picks $i \in \{1,2\}$ uniformly at random and then picks j according to $D^{(i)}$ is a k-canonical test that $\epsilon/2$-separates $\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$ from RM-$\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$. To verify the claim note that all tests accept x^a for $a \in \mathrm{Deg}(\mathcal{F}_1) \cap \mathrm{Deg}(\mathcal{F}_2) = \mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$. On the other hand if $a \in$ RM-$\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2) - \mathrm{Deg}(\mathcal{F}_\ell)$ then a is also in RM-$\mathrm{Deg}(\mathcal{F}_\ell) - \mathrm{Deg}(\mathcal{F}_\ell)$ and in such case with probability $1/2$ we pick $i = \ell \in \{1,2\}$ and then with further probability ϵ we pick j such that x^a does not satisfy $C_j^{(i)}$. □

The following theorem now follows immediately from Lemma 4.2 above and Theorem 3.4.

Theorem 4.3. *For all $q, k_1, k_2, \epsilon_1, \epsilon_2 > 0$, there exists $k < \infty$ and $\epsilon > 0$ such that, for every n, if $\mathcal{F}_1 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k_1, ϵ_1)-locally testable and $\mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k_2, ϵ_2)-locally testable, then $\mathcal{F}_1 \cap \mathcal{F}_2$ is (k, ϵ)-locally testable.*

Proof. Fix $q, k_1, k_2, \epsilon_1, \epsilon_2 > 0$. By Theorem 3.4, for each $i \in \{1,2\}$, there is a k'_i-canonical test T_i that ϵ'_i-separates $\mathrm{Deg}(\mathcal{F}_i)$ from RM-$\mathrm{Deg}(\mathcal{F}_i)$. By Lemma 4.2, there is a k'-canonical test $\epsilon'/2$-separating $\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$ from RM-$\mathrm{Deg}(\mathcal{F}_1 \cap \mathcal{F}_2)$ where $k' = \max\{k'_1, k'_2\}$ and $\epsilon' = \min\{\epsilon'_1, \epsilon'_2\}$. By Theorem 3.4, there exist $k, \epsilon > 0$, independent of n, such that $\mathcal{F}_1 \cap \mathcal{F}_2$ is (k, ϵ)-locally testable. □

4.2 Closure under summation

Proposition 4.4. *Let $\mathcal{F}_1, \mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ be affine-invariant properties. Then* $\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2) = \mathrm{Deg}(\mathcal{F}_1) \cup \mathrm{Deg}(\mathcal{F}_2)$.

Proof. Since $\mathcal{F}_1 \cup \mathcal{F}_2 \subseteq \mathcal{F}_1 + \mathcal{F}_2$ it follows that $\mathrm{Deg}(\mathcal{F}_1) \cup \mathrm{Deg}(\mathcal{F}_2) \subseteq \mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$. In the reverse direction, for every $f \in \mathcal{F}_1 + \mathcal{F}_2$, we have $f = f_1 + f_2$ with $f_i \in \mathcal{F}_i$ for $i \in \{1,2\}$. It follows that $\mathrm{supp}(f) \subseteq \mathrm{supp}(f_1) \cup \mathrm{sup}(f_2) \subseteq \mathrm{Deg}(\mathcal{F}_1) \cup \mathrm{Deg}(\mathcal{F}_2)$. Hence we get $\mathrm{Deg}(\mathcal{F}) = \cup_{f \in \mathcal{F}} \mathrm{supp}(f) \subseteq \mathrm{Deg}(\mathcal{F}_1) \cup \mathrm{Deg}(\mathcal{F}_2)$. □

Definition 4.5. *If $\vec{u} = \langle u_i \rangle_{i=1}^{s} \in \mathbb{F}_{q^n}^{s}$ and $\vec{v} = \langle v_j \rangle_{j=1}^{t} \in \mathbb{F}_{q^n}^{t}$, then the* tensor *of $\vec{u}$ and $\vec{v}$ is*

$$\vec{u} \otimes \vec{v} = \langle u_i v_j \rangle_{\substack{1 \le i \le s \\ 1 \le j \le t}} \in \mathbb{F}_{q^n}^{st}$$

The tensor *of two constraints $C_1 = (\vec{\alpha}, \vec{\lambda})$ and $C_2 = (\vec{\beta}, \vec{\mu})$ is*

$$C_1 \otimes C_2 := (\vec{\alpha} \otimes \vec{\beta}, \vec{\lambda} \otimes \vec{\mu}).$$

The following proposition is implicit in [2, 5].

Proposition 4.6. *x^d satisfies the constraint $C_1 \otimes C_2$ if and only if x^d satisfies at least one of the constraints C_1 or C_2.*

Proof. Let $C_1 = (\vec{\alpha}, \vec{\lambda})$ where $\vec{\alpha} = \langle \alpha_1, \ldots, \alpha_{k_1} \rangle \in \mathbb{F}_{q^n}^{k_1}$ and $\vec{\lambda} = \langle \lambda_1, \ldots, \lambda_{k_1} \rangle \in \mathbb{F}_{q^n}^{k_1}$, and similarly let $C_2 = (\vec{\beta}, \vec{\mu})$ where $\vec{\beta} = \langle \beta_1, \ldots, \beta_{k_2} \rangle \in \mathbb{F}_{q^n}^{k_2}$ and $\vec{\mu} = \langle \mu_1, \ldots, \mu_{k_2} \rangle \in \mathbb{F}_{q^n}^{k_2}$. Then x^d satisfies $C_1 \otimes C_2$ if and only if

$$0 = \sum_{i=1}^{k_1} \sum_{j=1}^{k_2} \lambda_i \mu_j (\alpha_i \beta_j)^d = \left(\sum_{i=1}^{k_1} \lambda_i \alpha_i^d \right) \left(\sum_{j=1}^{k_2} \mu_j \beta_j^d \right)$$

if and only if x^d satisfies at least one of C_1 or C_2. □

Lemma 4.7. *For $i \in \{1,2\}$ if there exist k-canonical tests ϵ-separating* $\mathrm{Deg}(\mathcal{F}_i)$ *from* RM-$\mathrm{Deg}(\mathcal{F}_i)$*, then there is a k^2-canonical test ϵ^2-separating* $\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$ *from* RM-$\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$.

Proof. For $\ell \in \{1,2\}$, let $T_\ell = (C_1^{(\ell)}, \ldots, C_m^{(\ell)}; D_i)$ be the k-canonical test ϵ-separating $\mathrm{Deg}(\mathcal{F}_i)$ from RM-$\mathrm{Deg}(\mathcal{F}_i)$. We claim that the test T, which picks $i \leftarrow_{D_1} [m]$ and $j \leftarrow_{D_2} [m]$ and accepts if and only if f satisfies $C_i^{(1)} \otimes C_j^{(2)}$, is a k^2-canonical test ϵ^2-separating $\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$ from RM-$\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$. For completeness, note that if $d \in \mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2) = \mathrm{Deg}(\mathcal{F}_1) \cup \mathrm{Deg}(\mathcal{F}_2)$, then for any $i, j \in [m]$, x^d satisfies at least one of $C_i^{(1)}$ or $C_j^{(2)}$, hence, by Proposition 4.6, x^d satisfies $C_i^{(1)} \otimes C_j^{(2)}$. For soundness, suppose $e \in$ RM-$\mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2) - \mathrm{Deg}(\mathcal{F}_1 + \mathcal{F}_2)$. The probability that x^e does not satisfy $C_i^{(1)} \otimes C_j^{(2)}$ equals the probability that x^e satisfies neither $C_i^{(1)}$ nor $C_j^{(2)}$. These events are independent and each happen with probability at least ϵ, hence the probability that neither constraint is satisfied is at least ϵ^2. □

The following theorem now follows immediately from Lemma 4.7 above and Theorem 3.4.

Theorem 4.8. *For every q, k_1, k_2 and $\epsilon_1, \epsilon_2 > 0$, there exists $k < \infty$ and $\epsilon > 0$ such that, for every n, if $\mathcal{F}_1 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k_1, ϵ_1)-locally testable and $\mathcal{F}_2 \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k_2, ϵ_2)-locally testable, then $\mathcal{F}_1 + \mathcal{F}_2$ is (k, ϵ)-locally testable.*

Proof. Fix $q, k_1, k_2, \epsilon_1, \epsilon_2 > 0$. By Theorem 3.4, for each $i \in \{1, 2\}$, there is a k_i'-canonical test T_i that ϵ_i'-separates $\mathrm{Deg}(\mathcal{F}_i)$ from RM-$\mathrm{Deg}(\mathcal{F}_i)$. By Lemma 4.7, there is a k'-canonical test ϵ'^2-separating $\mathrm{Deg}(\mathcal{F}_1+\mathcal{F}_2)$ from RM-$\mathrm{Deg}(\mathcal{F}_1+\mathcal{F}_2)$ where $k' = \max\{k_1', k_2'\}$ and $\epsilon' = \min\{\epsilon_1', \epsilon_2'\}$. By Theorem 3.4, there exist $k, \epsilon > 0$, independent of n, such that $\mathcal{F}_1 + \mathcal{F}_2$ is (k, ϵ)-locally testable. □

4.3 Lifts, and closure under lifts

Given a family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ its ℓ-lift defines a family of functions mapping $\mathbb{F}_{q^{nm}}$ to $\mathbb{F}_q$ as defined next. In viewing the definition below, we use the notation $f|_S$ to denote the restriction of f to the domain S. We also use the fact that $\mathbb{F}_{q^n} \subset \mathbb{F}_{q^{n\ell}}$.

Definition 4.9 (Lift [4]). *Given a family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ its ℓ-lift, denoted* $\mathrm{Lift}_\ell(\mathcal{F})$, *is the family*

$$\mathrm{Lift}_\ell(\mathcal{F}) = \{f : \mathbb{F}_{q^{n\ell}} \to \mathbb{F}_q \mid (f \circ A)|_{\mathbb{F}_{q^n}} \in \mathcal{F}, \forall \text{ affine } A : \mathbb{F}_{q^{n\ell}} \to \mathbb{F}_{q^{n\ell}}\}.$$

We note that while the definition above seems somewhat unnatural, it turns out (as noted in [12]) to be equivalent to the following much more natural definition.

Definition 4.10 (Lift, alternate definition). *Given a family $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ its ℓ-lift, denoted* $\mathrm{Lift}_\ell(\mathcal{F})$, *is the family*

$$\mathrm{Lift}_\ell(\mathcal{F}) = \{f : \mathbb{F}_{q^n}^\ell \to \mathbb{F}_q \mid f_L \in \mathcal{F}, \forall \text{ one dimensional affine subspaces } L\}.$$

This definition, equivalent under every linearity preserving isomorphism between $\mathbb{F}_{q^n}^{\ell}$ and $\mathbb{F}_{q^{n\ell}}$, makes the notion very natural, and as pointed out in [12] very useful. In this work, however, we work with the original definition.

To prove that local testability is closed under lifts, we will need to use a bit more of the well-known aspects of degree sets, and in particular the notion of "shadows".

Let $q = p^s$ for prime p. Let $d_0, d_1, \ldots$ be the base-p expansion of d (i.e., $0 \leq d_i < p$ and $d = \sum_i d_i p^i$). Similarly let $e_0, e_1, \ldots$ be the base-p expansion of e. We say that e is in the p-shadow of d, denoted $e \leq_p d$, if $e_i \leq d_i$ for every $i \geq 0$. The following proposition is well-known (see, for instance, [2]).

Proposition 4.11. *Let $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ be an affine-invariant linear property and let $q = p^s$ for prime p. Then $\mathrm{Deg}(\mathcal{F})$ is p-shadow-closed, i.e., if $d \in \mathrm{Deg}(\mathcal{F})$ and $e \leq_p d$ then $e \in \mathrm{Deg}(\mathcal{F})$.*

The following proposition relates the degree set of the lifted family to the degree set of a given family. The relationship uses the notion of p-shadows and a variation of the standard modular reduction, which is termed mod*, where $a \pmod{^* b}$ sends $a \in \mathbb{Z}^{\geq 0}$ to an integer in $\{0, \ldots, b-1\}$ so as to satisfy $x^a = x^{a \ (\mathrm{mod}^* \ b)} \pmod{x^b - x}$.

Proposition 4.12 ([4]). *Let $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ be an affine-invariant property. Then for every m,*

$$\mathrm{Deg}(\mathrm{Lift}_\ell(\mathcal{F})) = \{d \in \{0, \ldots, q^{n\ell} - 1\} \mid \forall e \leq_p d, e \ (\mathrm{mod}^* \ q^n - 1) \in \mathrm{Deg}(\mathcal{F})\}.$$

Proposition 4.13. *For every q, k and w, there exists a positive constant $\epsilon = \epsilon(q,k,w) > 0$ such that for every n the following is true. If $e \leq_p d$ and q-$\mathrm{wt}(d) \leq w$ and x^e does not satisfy some k-constraint C, then x^d does not satisfy ϵ fraction of the k-constraints $\{C \circ A\}_{\text{affine } A}$.*

Proof. Let w_0 be the weight from Theorem 2.4 such that families satisifying constraints of weight k are contained in $\mathrm{RM}[n, w_0, q]$. Let $w_1 = \max\{w_0, w\}$ and let $\delta_0 = \delta(\mathrm{RM}[n, w_1, q])$ be the minimum distance of $\mathrm{RM}[n, w_1, q]$. Let $\epsilon = \delta_0/(ck^2)$, where c is the constant from Theorem 2.1. We prove the lemma for this choice of ϵ.

Consider the affine-invariant family $\mathcal{F}' = \{f \mid f \text{ satisifes } C \circ A,\ \forall \text{ affine } A\}$. $\mathcal{F}'$ is a single-orbit characterized family and $e \notin \mathrm{Deg}(\mathcal{F}')$. It follows (since degree sets are shadow-closed, Proposition 4.11) that $d \notin \mathrm{Deg}(\mathcal{F}')$. Thus there exists some $\gamma \in \mathbb{F}_{q^n}$ such that $\mathrm{Tr}(\gamma x^d) \notin \mathcal{F}'$. Since $\{\mathrm{Tr}(\lambda x^d)\} \cup \mathcal{F}' \subseteq \mathrm{RM}[n, w_1, q]$ it follow that $\mathrm{Tr}(\gamma \cdot x^d)$ is δ_0-far from $\mathcal{F}'$. Applying Theorem 2.1 we get that the probability that a random affine map A would lead to a constraint $C \circ A$ that rejects $\mathrm{Tr}(\gamma \cdot x^d)$ is at least $\delta_0/(ck^2) = \epsilon$. For a choice of A such that $\mathrm{Tr}(\gamma \cdot x^d)$ does not satisfy $C \circ A$, it is also the case that x^d does not satisfy $C \circ A$, thus yielding the lemma. □

Lemma 4.14. *For every q, k and $\epsilon > 0$ there exist k' and ϵ' such that for every ℓ, n the following holds: If $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ has a k-canonical test ϵ-separating* $\mathrm{Deg}(\mathcal{F})$ *from* $\mathrm{RM\text{-}Deg}(\mathcal{F})$, *then there is a k'-canonical test ϵ'-separating* $\mathrm{Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$ *from* $\mathrm{RM\text{-}Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$.

Proof. Let w be the constant from Theorem 2.4 so that every affine-invariant family mapping $\mathbb{F}_{q^{n'}} \to \mathbb{F}_q$ is contained in $\mathrm{RM}[n', w-1, q]$. (We note that we intend to apply this to $n' = n\ell$. But w does not depend on n' and so doesn't depend on ℓ.) Let $\epsilon_1 = \epsilon(q, k, w)$ be the constant from Proposition 4.13. We prove the lemma for $k' = k$ and $\epsilon' = \epsilon \cdot \epsilon_1$.

Let $T = (C_1, \ldots, C_m; D)$ be a k-canonical test ϵ-separating $\mathrm{Deg}(\mathcal{F})$ from $\mathrm{RM\text{-}Deg}(\mathcal{F})$. For $i \in [m]$, let $C_i = (\vec{\alpha_i}, \vec{\lambda_i})$. Consider the tester T' which chooses a random affine $A : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$ and accepts f if and only if f satisfies $C \circ A$. We claim that T' is k-canonical tester that ϵ'-separates $\mathrm{Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$ from $\mathrm{RM\text{-}Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$ for some ϵ' independent of n.

For the completeness, suppose $d \in \mathrm{Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$. Let $d' = d \ (\mathrm{mod}^*\, q^n - 1)$. Then $d' \in \mathrm{Deg}(\mathcal{F})$, so $x^{d'}$ satisfies $C_i \circ A$ for every i and all affine $A : \mathbb{F}_{q^n} \to \mathbb{F}_{q^n}$. This follows since $x^{d'}$ certainly satisfies C_i, and if $A(x) = ax + b$ where $a, b \in \mathbb{F}_{q^n}$, then

$$\sum_{j=1}^{k} \lambda_{ij}(a \cdot \alpha_j + b)^{d'} = \sum_{j=1}^{k} \lambda_{ij} \left(\sum_{e' \leq_p d'} \binom{d'}{e'} a^{e'} \alpha_j^{e'} b^{d'-e'} \right) = \sum_{e' \leq_p d'} \binom{d'}{e'} a^{e'} b^{d'-e'} \left(\sum_{j=1}^{k} \lambda_{ij} \alpha_j^{e'} \right) = 0$$

where the first equality follows from Lucas' theorem and last equality holds since for every $e' \leq_p d'$, we have $e' \in \mathrm{Deg}(\mathcal{F})$ and so $x^{e'}$ satisfies C_i.

Now we turn to the soundness. Fix $e \in \mathrm{RM\text{-}Deg}(\mathrm{Lift}_\ell(\mathcal{F})) - \mathrm{Deg}(\mathrm{Lift}_\ell(\mathcal{F}))$. Note that since $\text{Lift}_\ell(\mathcal{F})$ satisfies the k-local constraint C_i, by Theorem 2.4, we have $\text{Lift}_\ell(\mathcal{F}) \subseteq \text{RM}[n\ell, w-1, q]$. Thus $\text{RM-Deg}(\text{Lift}_\ell(\mathcal{F})) \subseteq \text{RM}[n\ell, w, q]$ and so $q\text{-wt}(e) \leq w$. Since $e \notin \text{Deg}(\text{Lift}_\ell(\mathcal{F}))$, there exists some $e' \leq_p e$ such that $e' \pmod{^*} q^n - 1) \notin \text{Deg}(\mathcal{F})$, and moreover $q\text{-wt}(e' \ (\text{mod}^* \, q^n - 1)) \leq q\text{-wt}(e') \leq q\text{-wt}(e) \leq w + 1$. Since T is an ϵ-seperating set, we have that with probability at least ϵ, over the choice of $i \leftarrow_D [m]$, $x^{e'}$ does not satisfy C_i. Fix such an i. By Proposition 4.13, we have that with probability at least ϵ_1, over the choice of A, x^e does not satisfy $C_i \circ A$. Thus x^e fails to satisfy $C_i \circ A$ with probability at least $\epsilon \cdot \epsilon_1 = \epsilon'$. □

The following theorem now follows immediately from Lemma 4.14 above and Theorem 3.4.

Theorem 4.15. *For every $q, k, \epsilon > 0$, there exists $k' < \infty$ and $\epsilon' > 0$ such that, for every n, ℓ, the following holds: If $\mathcal{F} \subseteq \{\mathbb{F}_{q^n} \to \mathbb{F}_q\}$ is (k, ϵ)-locally testable, then* $\text{Lift}_\ell(\mathcal{F})$ *is (k', ϵ')-locally testable.*

Proof. Fix $q, k, \epsilon > 0$. By Theorem 3.4, there is a k_1-canonical test that ϵ_1-separates $\text{Deg}(\mathcal{F})$ from $\text{RM-Deg}(\mathcal{F})$. By Lemma 4.14, there is a k_2-canonical test ϵ_2-separating $\text{Deg}(\text{Lift}_\ell(\mathcal{F}))$ from $\text{RM-Deg}(\text{Lift}_\ell(\mathcal{F}))$. By Theorem 3.4, there exist k' and $\epsilon' > 0$, independent of n and ℓ, such that $\text{Lift}_\ell(\mathcal{F})$ is (k', ϵ')-locally testable. □

Bibliography

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

[2] Eli Ben-Sasson, Elena Grigorescu, Ghid Maatouk, Amir Shpilka, and Madhu Sudan. On sums of locally testable affine invariant properties. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:79, 2011.

[3] Eli Ben-Sasson, Prahladh Harsha, and Sofya Raskhodnikova. Some 3CNF properties are hard to test. *SIAM Journal on Computing*, 35(1):1–21, 2005. (Preliminary Version in *35th STOC*, 2003).

[4] Eli Ben-Sasson, Ghid Maatouk, Amir Shpilka, and Madhu Sudan. Symmetric LDPC codes are not necessarily locally testable. In *IEEE Conference on Computational Complexity*, pages 55–65, 2011.

[5] Eli Ben-Sasson, Noga Ron-Zewi, and Madhu Sudan. Testing sparse affine-invariant linear properties over all fields. Manuscript, 2012.

[6] Eli Ben-Sasson and Madhu Sudan. Limits on the rate of locally testable affine-invariant codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:108, 2010.

[7] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *FOCS*, pages 488–497, 2010.

[8] Victor Chen, Madhu Sudan, and Ning Xie. Property testing via set-theoretic operations. In *ICS*, pages 211–222, 2011.

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

[10] Elena Grigorescu, Tali Kaufman, and Madhu Sudan. 2-transitivity is insufficient for local testability. In *IEEE Conference on Computational Complexity*, pages 259–267, 2008.

[11] 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.

[12] Alan Guo, Swastik Kopparty, and Madhu Sudan. New affine-invariant codes from lifting. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:149, 2012.

[13] 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.

[14] 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.

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

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

[17] Madhu Sudan. Invariance in property testing. In *Property Testing*, pages 211–227, 2010.