



Spectral Norm, Economical Sieve, and Linear Invariance Testing of Boolean Functions

Swarnalipa Datta ✉ 


Indian Statistical Institute, Kolkata, India

Arijit Ghosh ✉ 


Indian Statistical Institute, Kolkata, India

Chandrima Kayal ✉ 

Université Paris Cité, CNRS, IRIF, France

Manaswi Paraashar ✉ 

University of Copenhagen, Denmark

Manmatha Roy ✉ 

Indian Statistical Institute, Kolkata, India

Abstract

Given Boolean functions $f, g : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, we say they are *linearly isomorphic* if there exists $A \in \text{GL}_n(\mathbb{F}_2)$ such that $f(x) = g(Ax)$ for all x . We study this problem in the tolerant property testing framework under the known–unknown model, where g is given explicitly and f is accessible only via oracle queries, meaning the algorithm may adaptively request the value of $f(x)$ for inputs $x \in \mathbb{F}_2^n$ of its choice. Given parameters $\epsilon \geq 0$ and $\omega > 0$, the goal is to distinguish whether there exists $A \in \text{GL}_n(\mathbb{F}_2)$ such that the normalized Hamming distance between f and $g(Ax)$ is at most ϵ , or whether for every $A \in \text{GL}_n(\mathbb{F}_2)$ the distance is at least $\epsilon + \omega$.

Our main result is a tolerant tester making $\tilde{O}((m/\omega)^4)$ queries to f , where m is an upper bound on the spectral norm of g , improving the previous $\tilde{O}((m/\omega)^{24})$ bound of Wimmer and Yoshida. We complement this with a nearly matching lower bound of $\Omega(m^2)$ for constant ω (for example, $\omega = 1/4$), improving the prior $\Omega(\log m)$ lower bound of Grigorescu, Wimmer and Xie. A key technical ingredient on the algorithmic side is a query-efficient local list corrector. For the lower bound, we give a reduction from communication complexity using a novel subclass of Maiorana–McFarland functions from symmetric-key cryptography.

2012 ACM Subject Classification Theory of computation \rightarrow Boolean function learning

Keywords and phrases Boolean Function, Isomorphism of Boolean Function, Fourier Analysis, Sublinear Algorithm, Property Testing

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.30

Related Version *Full Version:* <https://arxiv.org/abs/2308.02662>

Funding *Arijit Ghosh:* Arijit Ghosh acknowledges partial support from the Science and Engineering Research Board (SERB), Government of India, through the MATRICS grant MTR/2023/001527, and from the Department of Science and Technology (DST), Government of India, through grant TPN-104427.

Chandrima Kayal: Chandrima Kayal is supported by French PEPR integrated project EPiQ (ANR-22-PETQ-0007).

Manmatha Roy: Manmatha Roy acknowledges support from the ISEA Phase-III initiative of the Ministry of Electronics and Information Technology (MeitY), Govt of India, through Grant No L-14017/1/2022-HRD.



© Swarnalipa Datta, Arijit Ghosh, Chandrima Kayal, Manaswi Paraashar, and Manmatha Roy;

licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 30; pp. 30:1–30:21



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

A Boolean function $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ is said to be linearly isomorphic to $g : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ if there exists $A \in \text{GL}_n(\mathbb{F}_2)$ such that $f(x) = g(Ax)$ for all $x \in \mathbb{F}_2^n$. The linear isomorphism is a fundamental notion with applications in coding theory, circuit complexity, and cryptography [16, 6, 38, 18, 21, 27, 17, 20]. A notable example is given by Reed–Muller codes, which are affine-invariant and whose decoding algorithms crucially exploit this symmetry [1]. Closely related isomorphism problems also arise in cryptography, such as the *Isomorphism of Polynomials* problem [22].

From an algorithmic standpoint, this problem has been extensively studied. Determining whether two Boolean functions are linearly isomorphic is computationally intractable even under the strong restriction that the linear transformation is a permutation matrix. The problem is known to be coNP-hard (when the functions are given in DNF form) and lies in Σ_2^P , but it is not known to be in coNP. Agrawal and Thierauf further showed that, unless the polynomial hierarchy collapses to Σ_3^P , the problem is not Σ_2^P -complete [2]. The best known algorithm reduces the task to the *Hypergraph Isomorphism Problem* and runs in $2^{O(n)}$ time [29].

In this work, we study linear isomorphism of Boolean functions in the *property testing* framework, where the goal is to decide whether an object has a given property or is far from having it, using only a small number of queries. We work in the stronger *tolerant testing* model, which requires distinguishing objects that are close to the property from those that are far. We now formally define the problem of testing linear isomorphism in this setting.

► **Definition 1** ((ϵ, ω) -TOLERANT LINEAR ISOMORPHISM TESTING). *Let $\epsilon \geq 0$ and $\omega > 0$. An algorithm is given full access to a known function $g : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ and oracle access to an unknown function $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, meaning that for any input $x \in \mathbb{F}_2^n$ the algorithm may query the oracle to obtain $f(x)$. It must distinguish, with probability at least $2/3$, between the following cases:*

- **(Yes)** *There exists $A \in \text{GL}_n(\mathbb{F}_2)$ such that $\delta(f \circ A, g) \leq \epsilon$.*
- **(No)** *For all $A \in \text{GL}_n(\mathbb{F}_2)$, $\delta(f \circ A, g) \geq \epsilon + \omega$.*

Here $\delta(f, g)$ denotes the fraction of inputs on which f and g differ, that is, $\delta(f, g) := \Pr_{x \sim \mathbb{F}_2^n} [f(x) \neq g(x)]$.

We now review prior work on the isomorphism testing problem, covering both algorithmic upper bounds and lower bounds.

1.1 Related Works

Wimmer and Yoshida [37] were the first to study linear isomorphism testing in the tolerant setting. They gave an algorithm for the $(\omega, 3\omega)$ -tolerant problem with query complexity $\tilde{O}((m/\omega)^{24})$, where m is an upper bound on the spectral norm of the known function g . They also proved an adaptive lower bound of $\Omega(m)$ queries in a certain subconstant regime of ω (see Remark 4). In addition, Chakraborty et al. [14] showed an $\Omega(k)$ lower bound for testing linear isomorphism, when the known function is a k -Junta, which implies an $\Omega(\log m)$ lower bound in terms of the spectral norm of the known function. Grigorescu, Wimmer, and Xie [25] further proved an adaptive lower bound of $\Omega(n^2)$ queries when the known function g is the inner-product function on n bits. Since the spectral norm of the inner-product function is at most $2^{n/2}$, this also yields an $\Omega(\log m)$ lower bound in terms of m .

Linear isomorphism is an affine-invariant property, and testing affine-invariant properties of Boolean functions has been widely studied; see, e.g., [28, 8, 7, 26]. In particular, the *regularity framework* of Hatami and Lovett [26] yields tolerant testers for a broad class of

affine-invariant properties, including linear isomorphism to a fixed function. However, a major limitation of these general frameworks is that the resulting query complexity grows as a tower-type function of the relevant complexity parameter.

Characterizing which properties admit constant-query testers is a central question in property testing. In graph property testing, a landmark result is the characterization of all testable properties via Szemerédi’s regularity lemma [5]. In our setting, Wimmer and Yoshida [37] gave an analogous characterization for linear isomorphism testing, showing that functions with small spectral norm are exactly those for which linear isomorphism can be tested with a constant number of queries.

A closely related and more restricted problem is testing isomorphism under permutations of variables, which has also been extensively studied [11, 3, 15, 14, 12, 10]. Alon et al. [4] showed that testing permutation isomorphism for arbitrary Boolean functions requires $\Omega(n)$ queries and can be done with $O(n \log n)$ queries. They also proved that testing permutation isomorphism to a k -junta requires $\Omega(k)$ queries, and also gave a nearly matching $O(k \log k)$ -query algorithm.

1.2 Our Contribution

Prior work exhibits a large gap between the known upper and lower bounds, even for Boolean functions with small spectral norm, which form the only class admitting constant-query testers. In this work, we nearly close this gap by presenting a non-adaptive randomized tester with substantially improved query complexity for all $\epsilon \geq 0$ and $\omega > 0$.

► **Theorem 2.** *Let $\epsilon \geq 0$ and $\omega > 0$. There exists a non-adaptive randomized algorithm (see Algorithm 1) that solves the (ϵ, ω) -TOLERANT LINEAR ISOMORPHISM TESTING problem for a known function $g : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ with spectral norm at most m and an unknown function $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, using $\tilde{O}((m/\omega)^4)$ queries to f , and succeeding with probability at least $2/3$. Here, the $\tilde{O}(\cdot)$ notation hides polynomial factors in $\log m$ and $\log(1/\omega)$.*

We complement this upper bound with the following nearly matching lower bound.

► **Theorem 3.** *For any $m > 0$, there exists a Boolean function $h : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ with spectral norm at most m such that every adaptive algorithm for the $(0, 1/4)$ -TOLERANT LINEAR ISOMORPHISM TESTING problem with respect to h requires $\Omega(m^2)$ queries to the unknown function.*

► **Remark 4.** Wimmer and Yoshida [37] proved an $\Omega(m)$ lower bound for the $(0, \omega)$ -tolerant linear isomorphism testing problem. However, to the best of our understanding, their proof applies only to the special case $\omega = 1/m$. Even if their argument could be extended to general ω , our lower bound is quadratically stronger and is based on a different approach, which we believe may be of independent interest.

2 Proof ideas

2.1 Proof idea of Theorem 2

At a high level, we follow the paradigm of testing by implicit learning introduced by Diakonikolas et al. [19], and build on the approach of Gopalan et al. [24] for testing induced membership in subclasses of Boolean functions with small spectral norm, which was also used by Wimmer and Yoshida [37] for linear isomorphism testing. Our main contribution is a refinement of this approach that yields a tester with significantly improved query complexity,

30:4 Linear Invariance Testing of Boolean Functions

compared to Wimmer and Yoshida [37]. In the process, we introduce an efficient list-correction framework, which we call the *Economical Sieve*, and which we believe may be of independent interest in the analysis of Boolean functions. We now outline the overall proof strategy.

We first recall that if two Boolean functions are approximately related by a linear transformation, then their Fourier spectra are strongly correlated. Formally, for $f, g : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, if there exists an invertible matrix $A \in \text{GL}_n(\mathbb{F}_2)$ such that $\delta(f, g \circ A) \leq \epsilon$, then

$$\sum_{\beta \in \mathbb{F}_2^n} \widehat{f}(\beta) \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon,$$

and the converse also holds. Moreover, when the known function g has small spectral norm, this implies that it suffices to focus on the heavy Fourier coefficients of the unknown function f . The main challenge is to locate and estimate these coefficients using only oracle access to f , with a number of queries independent of n . At first glance, this seems difficult, since even identifying a linear function requires $\Omega(n)$ queries; see, e.g., [32, Exercise 1.27].

We construct an approximation f^* of the unknown function f such that there exists (yet unknown to the algorithm) $A \in \text{GL}_n(\mathbb{F}_2)$ such that $\|f - f^* \circ A\|_2^2 \leq \epsilon$, using $\text{poly}(m, 1/\epsilon)$ queries to f , where m is the spectral norm of the known function g . The crucial observation is that, for testing linear isomorphism, it suffices to learn f only up to an unknown nonsingular linear transformation.

The key tool enabling this is a local list-correction algorithm. Intuitively, given a threshold parameter θ , oracle access to f , and a point $x \in \mathbb{F}_2^n$, the algorithm identifies all heavy Fourier coefficients of f and returns a list $L_x = \{\chi_\beta(x) : |\widehat{f}(\beta)| \geq \theta\}$. The algorithm is not required to output the Fourier characters β themselves, but only their evaluations at the given point x . Its query complexity depends only on θ and is independent of the dimension n . For our purposes, this weaker guarantee is sufficient. By accessing L_x at many random inputs x (polynomially many in $1/\theta$), we show how one can (i) approximate the heavy Fourier coefficients of f , and (ii) recover a basis for them, thereby identifying the coefficients up to a linear transformation. This yields a function that approximates f up to linear isomorphism, which is adequate for our setting, since linear isomorphism is itself a linear-invariant property. Thus, it suffices to check over all possible linear transformations to determine whether a suitable one exists. This can be carried out offline and does not require any further queries to the unknown function f . It remains to describe the local list-correction algorithm, which we refer to as the *Economical Sieve*.

The *Economical Sieve* can be viewed as a query-efficient refinement of the implicit learning framework known as the *Implicit Sieve*, introduced by Wimmer and Yoshida [37]. Our improved query complexity is achieved through a more refined analysis of the *coset hashing* process. Specifically, we revisit coset hashing process and establish new concentration bounds (Claim 12 and 14) for the ℓ_1 - and ℓ_2 -norms of the projected Fourier spectrum. These results show that, within any coset, the norm of the Fourier projection is dominated by the contribution of heavy coefficients. Our analysis is inspired by techniques for heavy-hitter detection in the streaming literature: ℓ_2 -concentration plays a role analogous to the COUNT-MIN SKETCH, while ℓ_1 -concentration corresponds to the COUNT SKETCH. Below we provide a brief sketch of the overall approach.

- **Heavy Fourier Coefficient Detection.** We begin by projecting the Fourier spectrum of f onto cosets of a random subspace $H \subseteq \mathbb{F}_2^n$ of dimension $\log O(1/\theta^8)$. This induces a pairwise-independent hashing of the Fourier coefficients of f into $O(1/\theta^8)$ cosets of H . The key ingredient is Claim 12, which shows that the ℓ_2 -norm of the Fourier projection in each coset is dominated by its heaviest coefficient. Thus, given sufficiently accurate

ℓ_2 -estimates for the cosets, we can identify all cosets containing a heavy Fourier coefficient. Unlike prior work [37, 24], which relied on estimating the Fourier ℓ_4 -norm of coset buckets in order to detect cosets containing heavy Fourier coefficients, our method relies solely on these new concentration results. Consequently, we avoid the ℓ_4 -estimation step, which is inherently query-intensive.

- **Heavy Fourier Coefficient Evaluation.** Once the heavy cosets have been identified, we proceed to evaluate the corresponding heavy Fourier characters. This step relies on Claim 14, which shows that the ℓ_1 -norm of the projected Fourier spectrum within each coset is dominated by its largest coefficient. Consequently, for any $x \in \mathbb{F}_2^n$, the projection can be well approximated by the evaluation of the heaviest Fourier coefficient at x . However, since we use much smaller coset structures than in [37], the probability that this approximation is sufficiently accurate for a fixed x is only constant. As our algorithm requires many such evaluations, we apply an amplification step in the spirit of the Goldreich–Levin algorithm [23].

2.2 Proof idea of Theorem 3

We prove our lower bound for linear isomorphism testing via a reduction from the *Approximate Matrix Rank* problem in the public-coin randomized communication complexity model. Our reduction follows the general communication-based framework for proving lower bounds in function property testing introduced by Blais, Brody and Matulef [9]. A key ingredient in our approach is the structural relationship between the sparsity of a Boolean function’s Fourier support and the dimension of its linear span, known as the *Fourier dimension*. A fundamental result of Sanyal [34] shows that the Fourier dimension is at most the square root of the Fourier sparsity, up to constant factors.

To exploit this relationship, we construct a class of Boolean functions based on the Maiorana–McFarland construction, which is widely used in symmetric-key cryptography due to its rich algebraic structure and well-understood spectral properties. By composing these functions with linear maps of varying rank, we obtain a family of functions whose Fourier sparsity and Fourier dimension can be carefully controlled, while preserving the quadratic relationship established in [34]. Depending on the rank of the applied linear transformation, functions from this family are either linearly isomorphic to one another or far from any such isomorphism. This enables a reduction from the *Approximate Matrix Rank* problem to testing linear isomorphism of Boolean functions.

In this reduction, Alice and Bob are given matrices A and B , respectively, and must decide whether the rank of their sum $C = A + B$ is exactly r , or significantly smaller, using limited communication and shared public randomness. We encode these matrices into Boolean functions from our constructed family such that:

- if $A+B$ has full rank, the resulting joint function is a *yes*-instance of the linear isomorphism testing problem;
- if $A + B$ is far from full rank (by a constant factor), the resulting joint function is a *no*-instance.

To complete the reduction, we use a result of Sherstov and Storozhenko [36], which proves an $\Omega(r^2)$ lower bound on the randomized communication complexity of distinguishing whether a matrix has rank r or at most cr , for a constant $c < 1$. This lower bound translates directly into a nearly matching query lower bound for linear isomorphism testing in our construction.

3 Background

In this section, we recall standard notions from the analysis of Boolean functions and communication complexity that will be used throughout the paper.

For $\alpha, \beta \in \mathbb{F}_2^n$, $\langle \alpha, \beta \rangle$ denotes the \mathbb{F}_2 -inner product. For a function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$, we write

$$\mathbb{E}_{x \sim \mathbb{F}_2^n} [f(x)] := \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} f(x).$$

For a function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$, its Fourier coefficient at $\alpha \in \mathbb{F}_2^n$ is

$$\widehat{f}(\alpha) := \mathbb{E}_{x \sim \mathbb{F}_2^n} [f(x)\chi_\alpha(x)], \quad \text{where } \chi_\alpha(x) := (-1)^{\langle \alpha, x \rangle}.$$

All expectations and probabilities over \mathbb{F}_2^n are with respect to the uniform distribution. Every such function admits the Fourier expansion

$$f(x) = \sum_{\alpha \in \mathbb{F}_2^n} \widehat{f}(\alpha)\chi_\alpha(x).$$

► **Fact 5** (Plancherel and Parseval Identities, see [32, Chapter 1]). *For any functions $f, g : \mathbb{F}_2^n \rightarrow \mathbb{R}$,*

$$\mathbb{E}_{x \sim \mathbb{F}_2^n} [f(x)g(x)] = \sum_{\alpha \in \mathbb{F}_2^n} \widehat{f}(\alpha)\widehat{g}(\alpha).$$

In particular,

$$\mathbb{E}_{x \sim \mathbb{F}_2^n} [f(x)^2] = \sum_{\alpha \in \mathbb{F}_2^n} \widehat{f}(\alpha)^2,$$

which equals 1 when $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$.

The *Fourier support* of f is $\text{supp}(\widehat{f}) := \{\alpha \in \mathbb{F}_2^n : \widehat{f}(\alpha) \neq 0\}$, and its size is called the *Fourier sparsity*. We say that f is s -Fourier sparse if $|\text{supp}(\widehat{f})| \leq s$. We also use the *spectral norm* of f , defined as $\|\widehat{f}\|_1 := \sum_{\alpha} |\widehat{f}(\alpha)|$.

For Boolean functions $f, g : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, we measure distance by

$$\delta(f, g) := \Pr_{x \sim \mathbb{F}_2^n} [f(x) \neq g(x)].$$

The *Linear Isomorphism Distance* is

$$\delta_{\mathcal{L}}(f, g) := \min_{A \in \text{GL}_n(\mathbb{F}_2)} \delta(f \circ A, g),$$

where $\text{GL}_n(\mathbb{F}_2)$ denotes the group of invertible linear maps over \mathbb{F}_2 . Both δ and $\delta_{\mathcal{L}}$ satisfy the triangle inequality.

► **Fact 6** (Self-Correction, see [32, Proposition 1.31]). *Suppose $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$ is ϵ -close to the linear function χ_α . Then, for every $x \in \mathbb{F}_2^n$,*

$$\Pr_{y \sim \mathbb{F}_2^n} [\chi_\alpha(x) = f(y)f(x+y)] \geq 1 - 2\epsilon.$$

For the lower bound, we rely on a communication complexity result for the *Approximate Matrix Rank* problem. In this problem, Alice receives a matrix $A \in \mathbb{F}_2^{r \times r}$ and Bob receives $B \in \mathbb{F}_2^{r \times r}$; their goal is to distinguish whether $\text{rank}(A+B) = r$ or $\text{rank}(A+B) \leq r/4$. We use the following fundamental result of Sherstov and Storozhenko [36].

► **Fact 7** (Approximate Matrix Rank Problem [36, Theorem 1.1]). *Alice receives a matrix $A \in \mathbb{F}_2^{r \times r}$, and Bob receives a matrix $B \in \mathbb{F}_2^{r \times r}$. Their task is to determine, for $C = A + B$, whether*

$$\text{rank}(C) = r \quad \text{or} \quad \text{rank}(C) \leq \frac{r}{4},$$

while exchanging as few bits of communication as possible. The public-coin randomized communication complexity of this problem is $\Omega(r^2)$.

4 A query-efficient tester for the linear isomorphism problem

We begin with the local list correction framework, which plays a fundamental role in our algorithm.

► **Lemma 8** (Economical Sieve). *There exists an algorithm, *Economical Sieve*, that takes parameters θ and λ as input, makes $\tilde{O}(\max(\frac{1}{\theta^4}, \frac{\lambda}{\theta^2}))$ queries to the truth table of a Boolean function $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, and, with probability at least $\frac{9}{10}$, outputs:*

- A matrix $Q \in \{-1, +1\}^{\lambda \times k}$, where the (i, j) -th entry is $\chi_{\alpha_j}(x_i)$;
- A column vector $F \in \{-1, 1\}^\lambda$, where $F(i) = f(x_i)$ for each $i \in [\lambda]$,

where each $x_i \in \mathbb{F}_2^n$ is drawn independently and uniformly at random, and each $\alpha_j \in \mathbb{F}_2^n$ belongs to a set $\mathcal{S} = \{\alpha_1, \dots, \alpha_k\}$ satisfying:

- For every $\alpha \in \mathbb{F}_2^n$ such that $|\hat{f}(\alpha)| \geq \theta$, we have $\alpha \in \mathcal{S}$;
- For all $\alpha \in \mathcal{S}$, it holds that $|\hat{f}(\alpha)| \geq \theta/2$.

For now, we assume Lemma 8 and proceed to prove Theorem 2. The proof of Lemma 8, which establishes the correctness and performance guarantees of the *Economical Sieve*, will be presented later.

4.1 Proof of Theorem 2

We begin by restating Theorem 2 for the sake of reader's convenience.

► **Theorem 2.** *Let $\epsilon \geq 0$ and $\omega > 0$. There exists a non-adaptive randomized algorithm (see Algorithm 1) that solves the (ϵ, ω) -TOLERANT LINEAR ISOMORPHISM TESTING problem for a known function $g : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ with spectral norm at most m and an unknown function $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, using $\tilde{O}((m/\omega)^4)$ queries to f , and succeeding with probability at least $2/3$. Here, the $\tilde{O}(\cdot)$ notation hides polynomial factors in $\log m$ and $\log(1/\omega)$.*

Proof. We prove the theorem by presenting an explicit algorithm, *LinearIsoTester* (Algorithm 1), and analyzing its theoretical guarantees to show that it satisfies the requirements of Theorem 2. Consider Algorithm 1. We assume that, for the parameter settings

$$\theta = \frac{\omega}{10m} \quad \text{and} \quad \lambda = \frac{100}{\theta^2} \log \frac{100}{\theta^2},$$

Lemma 8 holds and the *Economical Sieve* correctly returns the pair (Q, F) . We now analyze the subsequent steps of Algorithm 1. Recall that each invocation of the *Economical Sieve* implicitly returns a set $\mathcal{S} = \{\alpha_1, \dots, \alpha_k\} \subseteq \mathbb{F}_2^n$ satisfying:

- For every $\alpha \in \mathbb{F}_2^n$ such that $|\hat{f}(\alpha)| \geq \theta$, we have $\alpha \in \mathcal{S}$;
- For all $\alpha \in \mathcal{S}$, it holds that $|\hat{f}(\alpha)| \geq \theta/2$.

30:8 Linear Invariance Testing of Boolean Functions

Having access to evaluations of $\chi_{\alpha_i}(x)$ for many inputs x , we can identify the characters in \mathcal{S} up to a linear transformation. Let \mathcal{T} be a subset of \mathcal{S} . Note that if $\sum_{\alpha_i \in \mathcal{T}} \alpha_i = \mathbf{0}^n$, then for all $x \in \mathbb{F}_2^n$,

$$\prod_{\alpha_i \in \mathcal{T}} \chi_{\alpha_i}(x) = \prod_{\alpha_i \in \mathcal{T}} (-1)^{\langle \alpha_i, x \rangle} = (-1)^{\langle \sum_{\alpha_i \in \mathcal{T}} \alpha_i, x \rangle} = 1.$$

Thus, the product of the corresponding columns of Q , denoted $\Pi_{\mathcal{T}}$, always equals $\mathbf{1}^\lambda$. On the other hand, if $\mathcal{B} \subseteq \mathcal{S}$ consists of linearly independent vectors, then the probability that $\prod_{\alpha_i \in \mathcal{B}} \chi_{\alpha_i}(x) = 1$ for all x is small.

▷ **Claim 9.** The probability that all entries of $\prod_{\alpha_i \in \mathcal{B}} \chi_{\alpha_i}(x_j)$ are equal to 1 is at most $2^{-\lambda}$. Moreover, for any fixed \mathcal{B} , the probability that this product equals 1 for any subset of \mathcal{B} is at most $1/100$.

Proof. Since the vectors α_i are linearly independent and the points x_1, \dots, x_λ are sampled uniformly at random from \mathbb{F}_2^n , we have $\Pr[\prod_{\alpha_i \in \mathcal{B}} \chi_{\alpha_i}(x_j) = 1] = 1/2$ for each independent sample x_j . Hence, the probability that this holds across all λ independent trials is at most $2^{-\lambda}$. Furthermore, the number of possible subsets of \mathcal{B} is at most $2^{|\mathcal{B}|} = 2^{O(1/\theta^2)}$. By applying the union bound, the probability that $\prod_{\alpha_i \in \mathcal{D}} \chi_{\alpha_i}(x_j) = 1$ occurs for any subset $\mathcal{D} \subseteq \mathcal{B}$ is at most $2^{-\lambda} \cdot 2^{O(1/\theta^2)} < o(1)$, provided that $\lambda = \Omega((1/\theta^2) \log(1/\theta))$. ◁

Having identified all the heavy Fourier coefficients of f (up to a linear transformation), we now estimate their corresponding Fourier magnitudes using the sample set (Q, F) . This is formalized in the following claim.

▷ **Claim 10.** Using random examples (Q, F) , one can estimate all heavy Fourier coefficients of f within $\pm\theta/10$, except with probability at most $1/25$.

Proof. For each $B_i \in \{B_1, B_2, \dots, B_k\}$, the algorithm uses the random examples (Q, F) to estimate $\widehat{f}(B_i)$ within an additive error of $\pm\theta/10$, achieving confidence $1 - \theta^2/100$. Let $\widetilde{f}(B_i)$ denote this estimate. A sample size of $\lambda = (100/\theta^2) \log(100/\theta^2)$ suffices to achieve this accuracy and confidence. Applying the union bound over all heavy Fourier coefficients, the probability that at least one estimate fails to meet the required accuracy is at most $1/25$. ◁

Finally, the algorithm constructs the real-valued function $f^*(x) = \sum_{B_i} \widetilde{f}(B_i) \cdot \chi_{B_i}(x)$. The following claim shows that this approximation suffices for testing linear isomorphism, as the property is linear-invariant and f^* preserves the spectral structure of f up to a linear transformation.

▷ **Claim 11.** Suppose f and g are ϵ -close to being linearly isomorphic. Then there exists a nonsingular transformation A such that $\sum_{\beta} \widehat{f^*}(\beta) \cdot \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon - \omega/100$. Conversely, if there exists a nonsingular transformation A satisfying $\sum_{\beta} \widehat{f^*}(\beta) \cdot \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon - \omega/100$, then f and g are $\epsilon + \omega/100$ -close to being linearly isomorphic.

Proof. Since f and g are ϵ -close to being linearly isomorphic, there exists an invertible linear transformation A such that $\Pr_{x \sim \mathbb{F}_2^n} [f(x) \neq g \circ A(x)] \leq \epsilon$. Using the fact that $f, g \in \{+1, -1\}$, we have:

$$\sum_{\beta \in \mathbb{F}_2^n} \widehat{f}(\beta) \cdot \widehat{g \circ A}(\beta) = \mathbb{E}_{x \sim \mathbb{F}_2^n} [f(x) \cdot g \circ A(x)] = 1 - 2 \cdot \Pr_{x \sim \mathbb{F}_2^n} [f(x) \neq g \circ A(x)] \geq 1 - 2\epsilon.$$

■ **Algorithm 1** LinearIsoTester.

Input: Given a function g such that $\|\widehat{g}\|_1 \leq m$, and query access to f
Output: $\delta_{\mathcal{L}}(f, g) \leq \varepsilon$ or $\delta_{\mathcal{L}}(f, g) > \varepsilon + \omega$ far from being linearly isomorphic
Initialization: $\theta = \frac{\omega}{10m}$, $\lambda = \frac{100}{\theta^2} \log \frac{100}{\theta^2}$

Step 1: $(Q, F) \leftarrow \mathbf{EconomicSieve}(\theta, \lambda)$

Step 2: Relabel the columns of Q as $\{B_1, \dots, B_k\}$, where

$B_1 = e_1^n, B_2 = e_2^n, \dots, B_r = e_r^n$, such that for all $i \in \{r+1, \dots, k\}$, the column B_i is a linear combination of $\{B_1, \dots, B_r\}$, where $(e_i^n)_{i \in [r]}$ denote the standard basis vectors of \mathbb{F}_2^n along coordinate direction i .

Step 3: Use (Q, F) to estimate each of $\widehat{f}(B_i)$ within $\pm \frac{\theta}{10}$ with confidence $1 - \frac{\theta^2}{100}$.

Let $\widetilde{f}(B_i)$ be the estimate

Step 4: Define the function f^* such that $\widehat{f^*}(\alpha) = \widetilde{f}(\alpha)$ for $\alpha \in \{B_1, \dots, B_k\}$, and 0 otherwise, where $\widetilde{f}(\alpha)$ denotes the estimated Fourier coefficient of f at point $\alpha \in \mathbb{F}_2^n$.

Step 5: Accept if and only if there exists a nonsingular linear transformation $A \in \mathbb{F}_2^{n \times n}$ such that

$$\sum_{\alpha \in \mathbb{F}_2^n} \widehat{g \circ A}(\alpha) \cdot \widehat{f^*}(\alpha) > 1 - 2\epsilon - \frac{\omega}{100}.$$

Since each Fourier coefficient of f^* approximates that of f within an additive error of at most $\theta/10$, we have:

$$\sum_{\beta \in \mathbb{F}_2^n} \widehat{f^*}(\beta) \cdot \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon - \frac{\theta}{10} \cdot \|\widehat{g}\|_1 \geq 1 - 2\epsilon - \frac{\omega}{100},$$

where the last inequality follows from $\theta = \frac{\omega}{10m}$ and $\|\widehat{g}\|_1 \leq m$.

For the other direction, note that by assumption, the algorithm finds a transformation A such that

$$\sum_{\beta \in \mathbb{F}_2^n} \widehat{f^*}(\beta) \cdot \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon - \frac{\omega}{100}.$$

Since each estimated Fourier coefficient in f^* is within $\frac{\theta}{10}$ of the true value in \widehat{f} , it follows that:

$$\sum_{\beta \in \mathbb{F}_2^n} \widehat{f}(\beta) \cdot \widehat{g \circ A}(\beta) \geq 1 - 2\epsilon - \frac{\omega}{100} - \frac{\theta}{10} \cdot \|\widehat{g}\|_1 \geq 1 - 2\epsilon - \frac{2\omega}{100}.$$

Using Parseval's identity, we have,

$$\begin{aligned} \sum_{\beta \in \mathbb{F}_2^n} \left(\widehat{f}(\beta) - \widehat{g \circ A}(\beta) \right)^2 &= \sum_{\beta} \widehat{f}^2(\beta) - 2 \sum_{\beta} \widehat{f}(\beta) \cdot \widehat{g \circ A}(\beta) + \sum_{\beta} \widehat{g \circ A}^2(\beta) \\ &= 2 - 2 \sum_{\beta} \widehat{f}(\beta) \cdot \widehat{g \circ A}(\beta) \\ &\leq 2 - 2 \cdot \left(1 - 2\epsilon - \frac{2\omega}{100} \right) \\ &= 4\epsilon + \frac{4\omega}{100}. \end{aligned}$$

30:10 Linear Invariance Testing of Boolean Functions

Further,

$$\Pr_{x \sim \mathbb{F}_2^n} [f(x) \neq g \circ A(x)] = \frac{1}{4} \cdot \mathbb{E}_{x \sim \mathbb{F}_2^n} [(f(x) - g \circ A(x))^2] \leq \frac{1}{4} \cdot \left(4\epsilon + \frac{4\omega}{100}\right) = \epsilon + \frac{\omega}{100}.$$

Thus, f and g are $\epsilon + \frac{\omega}{100}$ -close to being linearly isomorphic, as claimed. \triangleleft

Finally, the overall failure probability of Algorithm 1 is at most $1/10$ (Step 1) + $o(1)$ (Step 2) + $1/25$ (Step 3), which sums to at most $1/3$. Regarding the query complexity, note that the only queries to the unknown function f occur in Step 1 via the **Economical Sieve**. By Lemma 8, for the chosen parameters θ and λ , the total query complexity is $\tilde{O}(m^4/\omega^4)$. This completes the proof. \blacktriangleleft

4.2 Proof of Lemma 8

The primary tool we use here is the projection of the Fourier coefficients of f onto a random coset structure and analyzing the concentration of these coefficients within the cosets. Below, we briefly outline the concept of random coset decomposition and then proceed to the main proof.

Specifically, we project the Fourier spectrum of f onto a collection of cosets \mathcal{C} , obtained from a randomly permuted coset decomposition of a subspace of codimension

$$t = \left\lceil \log \left(\frac{2^{32}}{\theta^8} \right) \right\rceil.$$

To construct this decomposition, we sample t independent uniform vectors $\beta_1, \dots, \beta_t \in \mathbb{F}_2^n$ and define

$$H = \text{span}\{\beta_1, \dots, \beta_t\}^\perp,$$

the subspace of vectors orthogonal to all β_i . Each coset of H is indexed by a vector $b \in \mathbb{F}_2^t$ and defined as

$$D(b) := \{\alpha \in \mathbb{F}_2^n : \langle \alpha, \beta_i \rangle = b_i \text{ for all } i \in [t]\}.$$

To randomize the coset labels, we further sample a uniform shift $z \in \mathbb{F}_2^t$ and relabel each coset $D(b)$ as $D(b+z)$, yielding a randomly permuted coset structure. As shown by Gopalan et al. [24], this construction induces a pairwise independent hash family over \mathbb{F}_2^n . We will rely heavily on this property throughout the proof. We now proceed to give the proof of Theorem 8 by analyzing Algorithm 2. Throughout the proof, we will use the notations listed in Table 1.

Proof of Lemma 8 (Economical Sieve). In **Step 1**, we project the Fourier spectrum of f onto a random coset structure \mathcal{C} of codimension $\log \frac{2^{32}}{\theta^8}$. In **Step 2**, we estimate the Fourier weight of f of each coset of \mathcal{C} within $\pm \frac{\theta^2}{4}$ accuracy. Given the parameter settings in the algorithm, we demonstrate that at the end of **Step 3** of Algorithm 2, we successfully identify a set of cosets $\mathcal{C} \subseteq \mathcal{C}$ that collectively contain all heavy Fourier coefficients while ensuring that its size remains bounded. One may observe a resemblance of this guarantee to the celebrated Goldreich-Levin theorem [23]. This resemblance is not coincidental; in fact, the this part of this algorithm can be viewed as an implicit version of the Goldreich-Levin algorithm. We now formally state and prove the following claim.

■ **Table 1** Notations.

Notation	Description
\mathcal{C}	A random coset decomposition of \mathbb{F}_2^n
C	Individual cosets of \mathcal{C}
\mathcal{C}	Set of cosets of \mathcal{C} that contains a <i>heavy</i> Fourier coefficient of f
α_C	The Fourier coefficient in C with the highest absolute value $\alpha_C := \arg \max_{\beta \in C} \widehat{f}(\beta) $
\mathcal{W}_C	Total Fourier weight of coset C e.g. $\sum_{\beta \in C} \widehat{f}(\beta)^2$
\mathcal{W}_C^*	$\widehat{f}^2(\alpha_C)$, weight of the heaviest Fourier coefficient in C
$\mathcal{P}_C(z)$	$\sum_{\beta \in C} \widehat{f}(\beta) \chi_\beta(z)$, projection of $f(z)$ into coset C
$\mathcal{P}_C^*(z)$	$\widehat{f}(\alpha_C) \chi_{\alpha_C}(z)$, projection of $f(z)$ onto the heaviest Fourier coefficient of C

▷ **Claim 12.** Let \mathcal{C} be a randomly permuted coset structure of codimension $\log \frac{2^{32}}{\theta^8}$. Then, except with probability at most $\frac{1}{20}$, after **Step 3** the algorithm outputs a set of cosets $\mathcal{C} \subseteq \mathcal{C}$ such that:

- (i) Every Fourier coefficient α with $|\widehat{f}(\alpha)| \geq \frac{\theta^2}{64}$ is mapped to a distinct coset. Consequently for any $C \in \mathcal{C}$

$$\max_{\beta \in C \setminus \{\alpha_C\}} |\widehat{f}(\beta)| < \frac{\theta^2}{64} \text{ where } \alpha_C := \arg \max_{\beta \in C} |\widehat{f}(\beta)|$$

- (ii) For any $C \in \mathcal{C}$, if $|\widehat{f}(\alpha_C)| \geq \theta$, then $C \in \mathcal{C}$.
 (iii) For every $C \in \mathcal{C}$, $|\widehat{f}(\alpha_C)| \geq \frac{\theta}{2}$.
 (iv) For every $C \in \mathcal{C}$, its Fourier weight $\mathcal{W}_C = \sum_{\beta \in C} \widehat{f}(\beta)^2$ is highly concentrated around weight of its dominant fourier coefficient

$$\sum_{\beta \in C \setminus \{\alpha_C\}} \widehat{f}(\beta)^2 < \frac{\theta^4}{2^{12}}$$

Proof. We begin by defining the following events.

Event E_1 : All Fourier coefficients α with $|\widehat{f}(\alpha)| \geq \frac{\theta^2}{64}$ are mapped to distinct cosets of \mathcal{C} .

Event E_2 : For every coset C whose dominant Fourier coefficient α_C satisfies $|\widehat{f}(\alpha_C)| < \frac{\theta^2}{64}$,

$$\mathcal{W}_C < \frac{\theta^2}{16}.$$

Event E_3 : For every coset C whose dominant coefficient satisfies $|\widehat{f}(\alpha_C)| \geq \frac{\theta^2}{64}$,

$$\mathcal{W}_C \leq \widehat{f}(\alpha_C)^2 + \frac{\theta^4}{2^{12}}.$$

Event E_4 : For every coset $C \in \mathcal{C}$, the estimate $\widetilde{\mathcal{W}}_C$ produced in Step 3 satisfies

$$|\widetilde{\mathcal{W}}_C - \mathcal{W}_C| \leq \theta^2/4.$$

Condition on the event $E := E_1 \cap E_2 \cap E_3 \cap E_4$. we show that the conclusion of the claim follows deterministically.

- (i) **Distinctness of heavy coefficients.** Since $\theta > \frac{\theta^2}{64}$, event E_1 immediately implies that all Fourier coefficients of magnitude at least θ are mapped to distinct cosets, establishing the first item of the claim.

30:12 Linear Invariance Testing of Boolean Functions

(ii) All Cosets with heavy coefficients are selected. Let C be a coset whose dominant coefficient α_C satisfies $|\widehat{f}(\alpha_C)| \geq \theta$. Then $\mathcal{W}_C \geq \widehat{f}(\alpha_C)^2 \geq \theta^2$. By event E_4 , $\widetilde{\mathcal{W}}_C \geq \mathcal{W}_C - \frac{\theta^2}{4} \geq \frac{3\theta^2}{4}$. Hence, C is selected into \mathcal{C} .

(iii) No coset with small dominant coefficient is selected. Let C be a coset whose dominant coefficient satisfies $|\widehat{f}(\alpha_C)| \leq \theta/2$.

Case 1: $|\widehat{f}(\alpha_C)| \leq \frac{\theta^2}{64}$.

By event E_2 , $\mathcal{W}_C < \frac{\theta^2}{16}$. Using event E_4 , $\widetilde{\mathcal{W}}_C \leq \mathcal{W}_C + \frac{\theta^2}{4} < \frac{3\theta^2}{4}$, and hence C is not selected.

Case 2: $\frac{\theta^2}{64} \leq |\widehat{f}(\alpha_C)| \leq \theta/2$.

By event E_3 , $\mathcal{W}_C \leq \widehat{f}(\alpha_C)^2 + \frac{\theta^4}{2^{12}} < \frac{5\theta^2}{16}$. Again using event E_4 , $\widetilde{\mathcal{W}}_C \leq \mathcal{W}_C + \frac{\theta^2}{4} < \frac{3\theta^2}{4}$, and C is not selected.

(iv) Fourier weight of the cosets is highly concentrated around the weight of the heaviest part. This follows directly from Event E_3 .

It remains to show that

$$\Pr_{H,b}[\neg E] = \Pr_{H,b}[\neg E_1 \cup \neg E_2 \cup \neg E_3 \cup \neg E_4] < \frac{1}{10},$$

When the Fourier spectrum of f is projected onto a randomly permuted coset structure by choosing a subspace $H \leq \mathbb{F}_2^n$ uniformly at random of codimension $t = \log \left\lceil \frac{2^{12} \cdot 100}{\theta^8} \right\rceil$, and an independent uniformly random shift $b \in \mathbb{F}_2^t$, and considering the restriction of f to the cosets $b + H$. We analyze the events E_1, E_2, E_3, E_4 individually and bound the probability that each of them fails.

Event E_1 (Isolation of heavy coefficients). First we show that since the coset structure is induced by a uniformly random subspace of codimension t , any fixed vector lies in a given coset with probability 2^{-t} . To see why, fix $\alpha \in \mathbb{F}_2^n$ and $b \in \mathbb{F}_2^t$. By definition, the event $\alpha \in D(b + z)$ is equivalent to

$$\forall i \in [t], \quad \langle \alpha, \beta_i \rangle = b_i + z_i.$$

Each z_i is an independent uniformly random bit, and hence

$$\Pr[\alpha \in D(b + z)] = \prod_{i=1}^t \Pr[z_i = \langle \alpha, \beta_i \rangle - b_i] = 2^{-t}.$$

Now consider two distinct vectors $\alpha, \alpha' \in \mathbb{F}_2^n$. They lie in the same coset if and only if

$$\langle \alpha - \alpha', \beta_i \rangle = 0 \quad \forall i \in [t].$$

Since $\alpha - \alpha' \neq 0$, each constraint holds independently with probability $1/2$, and thus the collision probability is 2^{-t} .

Now, fix two distinct Fourier coefficients $\alpha, \beta \in S$. By the above, the probability that they collide into the same coset is 2^{-t} . Applying the union bound over all $\binom{|S|}{2}$ pairs, the total collision probability is at most

$$\binom{|S|}{2} 2^{-t} \leq |S|^2 2^{-t} \leq \left(\frac{2^{12}}{\theta^4} \right)^2 \cdot \frac{\theta^8}{2^{32}} < \frac{1}{100},$$

where we used $t = \lceil \log(2^{32}/\theta^8) \rceil$. Therefore, with probability at least 0.99, all coefficients in S are isolated into distinct cosets.

Event E_2 (Fourier Weight Concentration for Heavy cosets). Fix a coset C containing a unique coefficient $\alpha_C \in S$. Write

$$\mathcal{W}_C = \widehat{f}(\alpha_C)^2 + \sum_{\beta \neq \alpha_C} \widehat{f}(\beta)^2 I_\beta,$$

where I_β is the indicator that β hashes to C .

For each $\beta \neq \alpha_C$, $\mathbb{E}_{H,b}[I_\beta] = 2^{-t}$. Therefore,

$$\mathbb{E}_{H,b}[\mathcal{W}_C - \widehat{f}(\alpha_C)^2] = \sum_{\beta \neq \alpha_C} \widehat{f}(\beta)^2 \mathbb{E}_{H,b}[I_\beta] = 2^{-t} \sum_{\beta \neq \alpha_C} \widehat{f}(\beta)^2 \leq 2^{-t} \leq \theta^8 / 2^{32}.$$

Applying Markov's inequality,

$$\Pr_{H,b}[\mathcal{W}_C - \widehat{f}(\alpha_C)^2 \geq \frac{\theta^4}{2^{12}}] \leq \frac{\theta^8 / 2^{32}}{\theta^4 / 2^{12}} = \theta^4 / 2^{20}.$$

Since there are at most $|S| \leq 2^{12} / \theta^4$ such cosets, a union bound gives overall failure probability at most $1/100$.

Event E_3 (Bounding Fourier Weights of Light cosets). Let C be a coset such that $|\widehat{f}(\beta)| < \theta^2 / 64$ for all $\beta \in C$. Then

$$\mathbb{E}_{H,b}[\mathcal{W}_C] = \sum_{\beta} \widehat{f}(\beta)^2 \mathbb{E}_{H,b}[I_\beta] = 2^{-t} \sum_{\beta} \widehat{f}(\beta)^2 \leq 2^{-t} \leq \theta^8 / s^{32}.$$

Using pairwise independence,

$$\text{Var}_{H,b}(\mathcal{W}_C) = \sum_{\beta} \widehat{f}(\beta)^4 \text{Var}_{H,b}(I_\beta) \leq \sum_{\beta} \widehat{f}(\beta)^4 \mathbb{E}_{H,b}[I_\beta] = 2^{-t} \sum_{\beta} \widehat{f}(\beta)^4.$$

Since $|\widehat{f}(\beta)| < \theta^2 / 64$ implies $\widehat{f}(\beta)^4 \leq \frac{\theta^4}{2^{12}} \widehat{f}(\beta)^2$,

$$\sum_{\beta} \widehat{f}(\beta)^4 \leq \frac{\theta^4}{2^{12}} \sum_{\beta} \widehat{f}(\beta)^2 = \frac{\theta^4}{2^{12}}$$

Thus $\text{Var}_{H,b}(\mathcal{W}_C) \leq 2^{-t} \frac{\theta^4}{2^{12}} \leq \frac{\theta^{12}}{2^{44}}$. Applying Chebyshev's inequality,

$$\Pr_{H,b}[\mathcal{W}_C - \mathbb{E}_{H,b}[\mathcal{W}_C] \geq \frac{\theta^2}{16}] \leq \frac{(\theta^{12} / 2^{44})}{(\theta^4 / 16^2)} = \theta^8 / 2^{36}.$$

There are at most $2^t \leq 2^{32} / \theta^8$ cosets, so by a union bound the total failure probability is at most $1/16$.

Event E_4 (Fourier weight estimation are correct upto $\pm \theta^2 / 4$). We assume that it happens with probability at least $1 - o(1)$. See Claim 15 further details.

Combining the failure probabilities, we get $\Pr_{H,b}[\neg E_1 \cup \neg E_2 \cup \neg E_3 \cup \neg E_4] < \frac{1}{10}$. \triangleleft

Having implicitly identified all the heavy cosets associated with heavy Fourier coefficients, the next step is to evaluate $\chi_{\alpha_c}(x)$ at a uniformly sampled point $x \in \mathbb{F}_2^n$, for each coset $C \in \mathcal{C}$. Importantly, while each survived coset C is known to contain a unique dominant Fourier character α_c , the identity of α_c itself remains unknown. From a coding-theoretic viewpoint, this task can be interpreted as a local list-correction problem for the first-order

30:14 Linear Invariance Testing of Boolean Functions

■ Algorithm 2 Economical Sieve.

Input: Threshold: θ , Number of samples : λ
Output: Coset Samples of $f \circ A$, with respect to subspace $\mathcal{S}(\theta)$
Parameters: $\gamma = \log 100\lambda$,

Step 1: Let \mathcal{C} be a randomly permuted coset structure of a randomly chosen subspace of codimension $\log \frac{2^{32}}{\theta^8}$;

Step 2: **foreach** $C \in \mathcal{C}$ **do**
 | Estimate weight of C within $\pm\theta^2/4$ accuracy;
end

Step 3: Discard any coset with estimated weight $\leq \frac{3}{4}\theta^2$; Let \mathcal{C} be the set of surviving cosets;

Step 4: **for** $i \in \{1, 2, 3, \dots, \lambda\}$ **do**
 | Sample x_i uniformly at random from \mathbb{F}_2^n and set $F[x_i] \leftarrow f(x_i)$;
 | Sample $\{y_1, y_2, \dots, y_\gamma\}$ each uniformly at random from \mathbb{F}_2^n ;
 | **foreach** y_j **do**
 | **foreach** $C \in \mathcal{C}$ **do**
 | Estimate $P_C f(y_j)$ and $P_C f(x_i + y_j)$, each within $\pm\frac{1}{8}\theta$ accuracy;
 | **end**
 | **end**
 | Set $Q[x_i][C] \leftarrow \text{median}_j \{\text{sign}(\mathcal{P}_C[f(y_j)]) \cdot \text{sign}(\mathcal{P}_C[f(x_i + y_j)])\}$;
end

Step 5: Return (Q, F)

Reed–Muller code. Given oracle access to a corrupted codeword (represented here by the Boolean function f), the goal is to recover the value of all nearby codewords at a queried position without explicitly decoding them. To this end, for any coset C , define

$$\mathcal{P}_C(z) := \sum_{\beta \in C} \widehat{f}(\beta) \chi_\beta(z).$$

The following claim shows that a noisy variant of the standard self-correction procedure succeeds simultaneously for all survived cosets.

▷ **Claim 13.** Let \mathcal{C} be the set of survived cosets. Then for any fixed $x \in \mathbb{F}_2^n$,

$$\Pr_{y \sim \mathbb{F}_2^n} \left[\forall C \in \mathcal{C}, \text{sign}(\mathcal{P}_C(x + y)) \cdot \text{sign}(\mathcal{P}_C(y)) = \chi_{\alpha_C}(x) \right] \geq \frac{7}{8}.$$

Proof. Before proceeding the proof of Claim 13, we show that for every coset C , for uniformly sampled z , $\mathcal{P}_C(z) = \sum_{\beta \in C} \widehat{f}(\beta) \chi_\beta(z)$, the projection of $f(z)$ onto a coset C is highly concentrated around $\mathcal{P}_C^*(z) = \widehat{f}(\alpha_C) \chi_{\alpha_C}(z)$. More formally, we establish the following concentration result.

▷ **Claim 14.** Suppose $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, and let $C \in \mathcal{C}$ be a coset with dominant Fourier coefficient α_C such that

$$\sum_{\beta \in C \setminus \{\alpha_C\}} \widehat{f}(\beta)^2 < \frac{\theta^4}{2^{12}}, \quad |\widehat{f}(\alpha_C)| > \frac{\theta}{2}, \quad \text{and} \quad \max_{\beta \in C \setminus \{\alpha_C\}} |\widehat{f}(\beta)| < \frac{\theta^2}{64}.$$

Then

$$\Pr_{z \sim \mathbb{F}_2^n} \left[|\mathcal{P}_C(z) - \mathcal{P}_C^*(z)| \geq \frac{\theta^2}{64} + \tau \right] \leq \frac{\theta^4}{2^{12}\tau^2}.$$

Proof. To establish this, we first recall that Fourier characters act as pairwise independent hash functions $\chi_\alpha : \mathbb{F}_2^n \rightarrow \{-1, +1\}$. Indeed, when z is uniformly sampled from \mathbb{F}_2^n , each bit z_i is independent and uniformly distributed in $\{0, 1\}$. For any nonzero $\alpha \in \mathbb{F}_2^n$, we have:

$$\mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_\alpha(z)] = \mathbb{E}_{z \sim \mathbb{F}_2^n} [(-1)^{\langle \alpha, z \rangle}] = \frac{1}{2}(1) + \frac{1}{2}(-1) = 0.$$

Moreover, for distinct $\alpha_1, \alpha_2 \in \mathbb{F}_2^n$, we have:

$$\mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_{\alpha_1}(z) \cdot \chi_{\alpha_2}(z)] = \mathbb{E}_{z \sim \mathbb{F}_2^n} [(-1)^{\langle \alpha_1 + \alpha_2, z \rangle}] = 0 = \mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_{\alpha_1}(z)] \cdot \mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_{\alpha_2}(z)],$$

establishing pairwise independence. Next, we show that the expectation of $(\mathcal{P}_C(z) - \mathcal{P}_C^*(z))$ is small, analyzing two distinct cases separately depending on whether $\mathbf{0} \in C'$ or not, where $C' = \{C - \alpha_c\}$.

Case I: When C does not contain $\mathbf{0}$ or $\mathbf{0}$ is the leader of the coset C .

$$\mathbb{E}_{z \sim \mathbb{F}_2^n} [\mathcal{P}_C(z) - \mathcal{P}_C^*(z)] = \mathbb{E}_{z \sim \mathbb{F}_2^n} \left[\sum_{\beta \in C'} \hat{f}(\beta) \chi_\beta(z) \right] = \sum_{\beta \in C'} \hat{f}(\beta) \cdot \mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_\beta(z)] = \sum_{\beta \in C'} \hat{f}(\beta) \cdot 0 = 0.$$

Case II: When C contains $\mathbf{0}$ and $\mathbf{0}$ is not the leader of the coset C .

$$\mathbb{E}_{z \sim \mathbb{F}_2^n} [\mathcal{P}_C(z) - \mathcal{P}_C^*(z)] = \mathbb{E}_{z \sim \mathbb{F}_2^n} \left[\sum_{\beta} \hat{f}(\beta) \chi_\beta(z) \right] + \hat{f}(0) = \sum_{\beta} \hat{f}(\beta) \cdot \mathbb{E}_{z \sim \mathbb{F}_2^n} [\chi_\beta(z)] + \hat{f}(0) \leq \hat{f}(0) = \frac{\theta^2}{64}.$$

Next, we show that the variance of $(\mathcal{P}_C(z) - \mathcal{P}_C^*(z))$ is also small:

$$\text{Var}_{z \sim \mathbb{F}_2^n} \left[\sum_{\beta} \hat{f}(\beta) \chi_\beta(z) \right] = \sum_{\beta \in C'} \text{Var}_{z \sim \mathbb{F}_2^n} [\hat{f}(\beta) \chi_\beta(z)] \leq \sum_{\beta \in C'} \mathbb{E}_{z \sim \mathbb{F}_2^n} \left[\left(\hat{f}(\beta) \chi_\beta(z) \right)^2 \right] = \sum_{\beta \in C'} \hat{f}(\beta)^2 = \frac{\theta^4}{2^{12}}$$

Applying Chebyshev's inequality,

$$\Pr_{z \sim \mathbb{F}_2^n} \left[|\mathcal{P}_C(z) - \mathcal{P}_C^*(z)| > \frac{\theta^2}{64} + \tau \right] \leq \frac{\text{Var}_{z \sim \mathbb{F}_2^n} [\mathcal{P}_C(z) - \mathcal{P}_C^*(z)]}{\tau^2} \leq \frac{\theta^4}{2^{12}\tau^2}$$

This completes the proof. ◁

For our algorithm, we set $\tau = \frac{\theta}{8}$. Substituting this value yields

$$\Pr_{z \sim \mathbb{F}_2^n} \left[|\mathcal{P}_C(z) - \mathcal{P}_C^*(z)| > \frac{\theta}{4} \right] \leq \frac{\theta^2}{64}.$$

Next, observe that since $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$, the number of Fourier coefficients of magnitude at least $\theta/2$ is at most $4/\theta^2$. As each survived coset contains a unique dominant Fourier coefficient of magnitude at least $\theta/2$, it follows that $|C| \leq \frac{4}{\theta^2}$. Applying a union bound over all survived cosets, we obtain

$$\Pr_{z \sim \mathbb{F}_2^n} \left[\forall C \in \mathcal{C}, |\mathcal{P}_C(z) - \mathcal{P}_C^*(z)| \geq \frac{\theta}{4} \right] \leq \frac{1}{16}.$$

30:16 Linear Invariance Testing of Boolean Functions

Combining this event with the fact that $|\widehat{f}(\alpha_C)| \geq \frac{\theta}{2}$ for every surviving coset C , and using that for any $x \in \mathbb{F}_2^n$ the sum $x + y$ is uniformly distributed over \mathbb{F}_2^n when y is chosen uniformly at random from \mathbb{F}_2^n , we conclude that

$$\Pr_{z \sim \mathbb{F}_2^n} \left[\forall C \in \mathcal{C}, \text{sign}(\mathcal{P}_C(x + y)) \cdot \text{sign}(\mathcal{P}_C(y)) = \chi_{\alpha_c}(x) \right] \geq \frac{7}{8}.$$

However, the algorithm does not have access to the exact value of $\mathcal{P}_C(z)$; it can only estimate it with reasonable accuracy. In our setting, if the estimated value $\widetilde{\mathcal{P}}_C(z)$ lies within $\pm \frac{\theta}{8}$ of the true value, the preceding argument still holds.

Moreover, to construct the pair (Q, F) under the given parameter settings of λ and α , it is necessary to evaluate $\chi_{\alpha_c}(x)$ for a sufficiently large number of inputs x , for all cosets $C \in \mathcal{C}$. This requires enhancing the reliability of the self-correction procedure. To achieve this, we apply the standard median trick for error reduction: for each fixed input x , we perform multiple independent trials using different random choices of y , and report the median of the obtained outcomes. By a direct application of the Chernoff bound, taking $O(\log(1/\theta))$ independent samples suffices to amplify the success probability to at least $1 - 1/\text{poly}(1/\theta)$, which is sufficient to apply a union bound over all surviving cosets and all inputs x . \triangleleft

\triangleright **Claim 15.** In Algorithm 2, the total number of queries made by the algorithm is bounded by $\widetilde{O}(\max(\lambda/\theta^2, 1/\theta^4))$.

Proof. The unknown function f is queried primarily in the following two steps of Algorithm 2:

Step 2: In this step, the algorithm estimates the weight of each of the $O(\frac{1}{\theta^8})$ cosets with accuracy $\pm \frac{\theta^2}{4}$ and failure probability at most $O(\theta^8)$. Then, the total number of queries required to estimate the weights of all cosets with the specified accuracy and confidence is $\widetilde{O}(\frac{1}{\theta^4})$. Note that for any $x \in \mathbb{F}_2^n$, the weight of a coset $r + H$ can be estimated via

$$\mathcal{W}_{r+H} = \mathbb{E}_{x \sim \mathbb{F}_2^n, z \sim H^\perp} [\chi_r(z) f(x) f(x + z)].$$

Moreover, a single batch of samples can be reused simultaneously to estimate the weight of all cosets.

Step 4: In this step, the algorithm samples λ many x 's, and for each, runs $\widetilde{O}(\log(1/\theta^2))$ independent trials. In each trial, for each coset, the projections

$$\mathcal{P}_{r+H} f(x) = \mathbb{E}_{y \sim H^\perp} [\chi_r(y) f(x + y)]$$

are estimated within $\pm \frac{\theta}{8}$ accuracy with failure probability at most $O(\theta^2)$. For any fixed x , the same batch of samples can simultaneously estimate the projections for all cosets. Then, the total number of queries in this step is also bounded by $\widetilde{O}(\lambda/\theta^2)$.

So, the overall query complexity of Algorithm 2 is $\widetilde{O}(\max(\lambda/\theta^2, 1/\theta^4))$. \triangleleft

This completes the analysis of Economical Sieve (proof of Lemma 8). \blacktriangleleft

5 A lower bound for testing linear isomorphism

We begin by defining a class of Boolean functions known as the *Maiorana–McFarland* functions. These functions have a long history of applications in theoretical computer science, particularly in proving lower bounds. Notable examples include circuit lower

bounds [33, 13] and studies on the structural properties of Boolean functions relevant to complexity theory [31, 34]. Beyond complexity theory, they play a fundamental role in symmetric-key cryptography, especially in the design of stream ciphers, where they serve as building blocks for achieving good confusion (captured by the Fourier or Walsh spectrum) and diffusion (captured by the autocorrelation spectrum); see [35] for further details. We now formally define Maiorana–McFarland functions.

► **Definition 16.** *Given positive integers n and r with $r \leq n$, the family of Maiorana–McFarland functions (originally introduced in [30]), denoted $MM_{r,n}$, consists of n -variable Boolean functions $f : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ of the form*

$$g(x, y) = (-1)^{\langle x, \varphi(y) \rangle}, \quad (x, y) \in \mathbb{F}_2^r \times \mathbb{F}_2^{n-r},$$

where $\varphi : \mathbb{F}_2^{n-r} \rightarrow \mathbb{F}_2^r$ is an arbitrary mapping.

A key property of Maiorana–McFarland functions that we exploit is that, when composed with suitable linear transformations, their Fourier sparsity is governed by the rank of the underlying transformation. We state this property formally below.

▷ **Claim 17.** Let $n = r + \log r$, and let $\varphi : \mathbb{F}_2^{n-r} \rightarrow \mathbb{F}_2^r$ be a mapping whose image has cardinality r , with the image set linearly independent in \mathbb{F}_2^r and L is a linear transformation in $\mathbb{F}_2^{r \times r}$. Let

$$g_L(x, y) = (-1)^{\langle Lx, \varphi(y) \rangle},$$

Then the Fourier sparsity of g_L is at most $\text{rank}(L) \times r$.

Proof. For $(u, v) \in \mathbb{F}_2^r \times \mathbb{F}_2^{n-r}$, the Fourier coefficient of g_L is

$$\widehat{g_L}(u, v) = \mathbb{E}_{x \sim \mathbb{F}_2^r, y \sim \mathbb{F}_2^{n-r}} \left[(-1)^{\langle Lx, \varphi(y) \rangle + \langle u, x \rangle + \langle v, y \rangle} \right].$$

Reordering the expectation yields

$$\widehat{g_L}(u, v) = \mathbb{E}_{y \sim \mathbb{F}_2^{n-r}} \left[(-1)^{\langle v, y \rangle} \cdot \mathbb{E}_{x \sim \mathbb{F}_2^r} (-1)^{\langle L^T \varphi(y) + u, x \rangle} \right].$$

Now

$$\mathbb{E}_{x \sim \mathbb{F}_2^r} \left[(-1)^{\langle L^T \varphi(y) + u, x \rangle} \right] = \begin{cases} 1 & \text{if } u = L^T \varphi(y), \\ 0 & \text{otherwise.} \end{cases}$$

Hence

$$\widehat{g_L}(u, v) = \frac{1}{2^{n-r}} \sum_{\substack{y \in \mathbb{F}_2^{n-r} \\ L^T \varphi(y) = u}} (-1)^{\langle v, y \rangle}.$$

Therefore (u, v) can be in the Fourier support only if $u \in \text{Im}(L^T \circ \varphi)$. Since φ has r distinct outputs that are linearly independent in \mathbb{F}_2^r , the set $\{L^T \varphi(y) : y \in \mathbb{F}_2^{n-r}\}$ spans a subspace of dimension at most $\text{rank}(L)$. Thus there are at most $\text{rank}(L)$ distinct u values that can occur. For each such u , there are $2^{n-r} = r$ possible choices of v . Hence the total number of possible nonzero Fourier coefficients is at most $\text{rank}(L) \cdot r$, establishing the claimed sparsity bound. ◁

Proof of Theorem 3

In this section, we prove Theorem 3, establishing a lower bound that is quadratically stronger than the previously known result from [24].

► **Theorem 3.** *For any $m > 0$, there exists a Boolean function $h : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ with spectral norm at most m such that every adaptive algorithm for the $(0, 1/4)$ -TOLERANT LINEAR ISOMORPHISM TESTING problem with respect to h requires $\Omega(m^2)$ queries to the unknown function.*

Proof. We establish the lower bound via a reduction from the Approximate Matrix Rank problem in randomized communication complexity (Theorem 7). Alice receives $A \in \mathbb{F}_2^{r \times r}$, Bob receives $B \in \mathbb{F}_2^{r \times r}$, and it is promised that $\text{rank}(C)$ for $C = A + B$ is either r or $r/4$. Their goal is to determine $\text{rank}(C)$ with minimal communication and shared randomness.

Let $g = g_I$ denote the reference function defined in Claim 17, where I is the $r \times r$ identity matrix. Alice constructs a function f_A from A , Bob constructs f_B from B , and together they define $f = f_C = f_{A+B}$. By Claim 17, if $\text{rank}(C) = r$, then $\text{supp}(\hat{f}) = r^2$ and if $\text{rank}(C) = r/4$, then $\text{supp}(\hat{f}) \leq r^2/4$. Moreover, if $\text{rank}(C) = r$, then f is linearly isomorphic to g . On the other hand, the following Claim shows that when $\text{rank}(C) = r/4$, the function f is far from every such isomorphism.

▷ **Claim 18.** If $\text{rank}(C) = r$, then f is $\frac{1}{4}$ -far from any Boolean function with Fourier sparsity at most $r^2/4$.

Proof. Let $h : \mathbb{F}_2^n \rightarrow \{-1, +1\}$ be a $\frac{r^2}{4}$ -Fourier sparse function. Then,

$$\begin{aligned} \Pr_{x \sim \mathbb{F}_2^n} [h(x) \neq f(x)] &= \Pr_{x \sim \mathbb{F}_2^n} [h(x) \neq f(x)] & (1) \\ &= \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x \sim \mathbb{F}_2^n} [h(x)f(x)] \\ &= \frac{1}{2} + \frac{1}{2} \sum_{\alpha \in \mathbb{F}_2^n} \hat{h}(\alpha) \hat{f}(\alpha) \\ &= \frac{1}{2} + \frac{1}{2} \sum_{\alpha \in \text{supp}(h)} \hat{h}(\alpha) \hat{f}(\alpha). & (2) \end{aligned}$$

Recall that for Boolean function $f : \mathbb{F}_2^n \rightarrow \{\pm 1\}$, $\text{supp}(\hat{f}) := \{\alpha \in \mathbb{F}_2^n : \hat{f}(\alpha) \neq 0\}$. Now applying Cauchy-Schwarz inequality, we get

$$\left| \sum_{\alpha \in \text{supp}(h)} \hat{h}(\alpha) \hat{f}(\alpha) \right| \leq \sqrt{\sum_{\alpha \in \text{supp}(h)} \hat{h}^2(\alpha) \cdot \sum_{\alpha \in \text{supp}(h)} \hat{f}^2(\alpha)} = \sqrt{\sum_{\alpha \in \text{supp}(h)} \hat{f}^2(\alpha)} \quad (3)$$

Note that h is a $\frac{r^2}{4}$ -Fourier sparse Boolean function, that is, $|\text{supp}(\hat{h})| \leq \frac{r^2}{4}$. Observe that, by the construction of function f (see Corollary 17), the absolute values of any two non-zero Fourier coefficients are equal and the Fourier support $\text{supp}(\hat{f}) = r^2$. Using the fact that $\sum_{\alpha \in \mathbb{F}_2^n} \hat{f}(\alpha)^2 = 1$ (Parseval's identity), we get

$$\sqrt{\sum_{\alpha \in \text{supp}(h)} \hat{f}^2(\alpha)} \leq \frac{1}{2} \quad (4)$$

Finally, plugging the bound from Equation (4) into Equation (1), we conclude that

$$\Pr_{x \sim \mathbb{F}_2^n} [h(x) \neq f(x)] \geq \frac{1}{2} + \frac{1}{2} \cdot \left(-\frac{1}{2}\right) = \frac{1}{4}. \quad \triangleleft$$

Thus, distinguishing between the two rank cases reduces to distinguishing whether f is isomorphic to g or $\frac{1}{4}$ -far from every isomorphism of g . Suppose there exists a tester \mathbb{T} for linear isomorphism with respect to a function of spectral norm m , with query complexity $q(m, \frac{1}{4})$. In our case, the reference function is g with $m = O(r)$. To simulate a query $(x, y) \in \mathbb{F}_2^n$, Alice computes $f_A(x, y)$, Bob computes $f_B(x, y)$, and they exchange one bit each. Since

$$f_C(x, y) = (-1)^{\langle (A+B)x, \varphi(y) \rangle} = f_A(x, y) \cdot f_B(x, y),$$

they can compute $f(x, y)$ with two bits of communication per query. Hence, simulating \mathbb{T} requires at most $2q(m, \frac{1}{4})$ bits. However, by Fact 7, solving the promised rank problem requires $\Omega(r^2)$ bits. Therefore, $q(m, \frac{1}{4}) = \Omega(r^2)$. Since $m = O(r)$ for g , we conclude that testing linear isomorphism with respect to g requires $\Omega(m^2)$ queries. \blacktriangleleft

6 Conclusion

A Boolean function $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$ is said to be affinely isomorphic to another function $g : \mathbb{F}_2^n \rightarrow \{-1, 1\}$ if there exist an invertible matrix $A \in \text{GL}_n(\mathbb{F}_2)$ and a vector $b \in \mathbb{F}_2^n$ such that for all $x \in \mathbb{F}_2^n$,

$$g(Ax + b) = f(x).$$

By adapting the techniques from the proofs of Theorem 2 and Theorem 3, we can establish analogous results for testing affine isomorphism between a known function and an unknown function given oracle access to its truth table.

References

- 1 Emmanuel Abbe, Amir Shpilka, and Min Ye. Reed-Muller Codes: Theory and Algorithms. *IEEE Transactions on Information Theory*, 67(6):3251–3277, 2021. doi:10.1109/TIT.2020.3004749.
- 2 Manindra Agrawal and Thomas Thierauf. The Formula Isomorphism Problem. *SIAM Journal on Computing*, 30(3):990–1009, 2000. doi:10.1137/S0097539798343647.
- 3 Noga Alon and Eric Blais. Testing Boolean Function Isomorphism. In *Proceedings of the 10th International Workshop on Randomization and Computation (RANDOM)*, pages 394–405, 2010. doi:10.1007/978-3-642-15369-3_30.
- 4 Noga Alon, Eric Blais, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Nearly Tight Bounds for Testing Function Isomorphism. *SIAM Journal on Computing*, 42(2):459–493, 2013. doi:10.1137/110832677.
- 5 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 (STOC)*, pages 251–260, 2006. doi:10.1145/1132516.1132555.
- 6 Alberto Bemporad. Efficient Conversion of Mixed Logical Dynamical Systems Into an Equivalent Piecewise Affine Form. *IEEE Transactions on Automatic Control*, 49(5):832–838, 2004. doi:10.1109/TAC.2004.828315.

- 7 Arnab Bhattacharyya, Eldar Fischer, Hamed Hatami, Pooya Hatami, and Shachar Lovett. Every Locally Characterized Affine-Invariant Property is Testable. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 429–436, 2013. doi:10.1145/2488608.2488662.
- 8 Arnab Bhattacharyya, Elena Grigorescu, and Asaf Shapira. A Unified Framework for Testing Linear-Invariant Properties. *Random Structures & Algorithms*, 46(2):232–260, 2015. A preliminary version of this work appeared as an extended abstract in the proceedings of FOCS 2010. doi:10.1002/RSA.20507.
- 9 Eric Blais, Joshua Brody, and Kevin Matulef. Property Testing Lower Bounds via Communication Complexity. *Computational Complexity*, 21(2):311–358, 2012. A preliminary version of this work appeared as an extended abstract in the proceedings of CCC 2011. doi:10.1007/S00037-012-0040-X.
- 10 Eric Blais, Clément L. Canonne, Talya Eden, Amit Levi, and Dana Ron. Tolerant Junta Testing and the Connection to Submodular Optimization and Function Isomorphism. *ACM Transactions on Computation Theory*, 11(4):24:1–24:33, 2019. A preliminary version of this work appeared as an extended abstract in the proceedings of SODA 2019. doi:10.1145/3337789.
- 11 Eric Blais and Ryan O’Donnell. Lower Bounds for Testing Function Isomorphism. In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity (CCC)*, pages 235–246, 2010. doi:10.1109/CCC.2010.30.
- 12 Eric Blais, Amit Weinstein, and Yuichi Yoshida. Partially symmetric functions are efficiently isomorphism testable. *SIAM Journal on Computing*, 44(2):411–432, 2015. A preliminary version of this work appeared as an extended abstract in the proceedings of FOCS 2012. doi:10.1137/140971877.
- 13 N. Blum. A Boolean function requiring $3n$ network size. *Theoretical Computer Science*, 28:337–345, 1984. doi:10.1016/0304-3975(83)90029-4.
- 14 Sourav Chakraborty, Eldar Fischer, David García-Soriano, and Arie Matsliah. Junto-Symmetric Functions, Hypergraph Isomorphism and Crunching. In *Proceedings of the 27th Conference on Computational Complexity (CCC)*, pages 148–158, 2012. doi:10.1109/CCC.2012.28.
- 15 Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Nearly Tight Bounds for Testing Function Isomorphism. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1683–1702, 2011. doi:10.1137/1.9781611973082.130.
- 16 Valentina Ciriani. Synthesis of SPP three-level logic networks using affine spaces. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 22(10):1310–1323, 2003. doi:10.1109/TCAD.2003.818121.
- 17 Lingguo Cui and Yuanda Cao. A new S-box structure named Affine-Power-Affine. *International Journal of Innovative Computing, Information and Control*, 3(3):751–759, 2007.
- 18 Stanisa Dautovic and Ladislav A. Novak. A comment on "Boolean functions classification via fixed polarity Reed-Muller form". *IEEE Transactions on Computers*, 55(8):1067–1069, 2006. doi:10.1109/TC.2006.114.
- 19 Ilias Diakonikolas, Homin K Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A Servedio, and Andrew Wan. Testing for Concise Representations. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 549–558, 2007. doi:10.1109/FOCS.2007.32.
- 20 Orr Dunkelman, Nathan Keller, and Adi Shamir. Minimalism in Cryptography: The Even-Mansour Scheme Revisited. In *Proceedings of the 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 336–354, 2012. doi:10.1007/978-3-642-29011-4_21.
- 21 Iwan M. Duursma, Carlos Rentería-Márquez, and Horacio Tapia-Recillas. Reed-Muller Codes on Complete Intersections. *Applicable Algebra in Engineering, Communication and Computing*, 11(6):455–462, 2001. doi:10.1007/S002000000047.

- 22 Jean-Charles Faugère and Ludovic Perret. Polynomial Equivalence Problems: Algorithmic and Theoretical Aspects. In *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 30–47, 2006. doi:10.1007/11761679_3.
- 23 Oded Goldreich and Leonid A. Levin. A Hard-Core Predicate for all One-Way Functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 25–32, 1989. doi:10.1145/73007.73010.
- 24 Parikshit Gopalan, Ryan O’Donnell, Rocco A Servedio, Amir Shpilka, and Karl Wimmer. Testing fourier dimensionality and sparsity. *SIAM Journal on Computing*, 40(4):1075–1100, 2011. doi:10.1137/100785429.
- 25 Elena Grigorescu, Karl Wimmer, and Ning Xie. Tight Lower Bounds for Testing Linear Isomorphism. In *Proceedings of the 17th International Workshop on Randomization and Computation (RANDOM)*, pages 559–574, 2013. doi:10.1007/978-3-642-40328-6_39.
- 26 Hamed Hatami and Shachar Lovett. Estimating the Distance from Testable Affine-Invariant Properties. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 237–242, 2013. doi:10.1109/FOCS.2013.33.
- 27 Xiang-Dong Hou. Classification of cosets of the Reed-Muller code $R(m - 3, m)$. *Discrete Mathematics*, 128(1):203–224, 1994. doi:10.1016/0012-365X(94)90113-9.
- 28 Tali Kaufman and Madhu Sudan. Algebraic Property Testing: The Role of Invariance. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 403–412, 2008. doi:10.1145/1374376.1374434.
- 29 Eugene M. Luks. Hypergraph Isomorphism and Structural Equivalence of Boolean Functions. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC)*, pages 652–658, 1999. doi:10.1145/301250.301427.
- 30 R. L. McFarland. A Family of Noncyclic Difference Sets. *Journal of Combinatorial Theory, Series A*, 15:1–10, 1973.
- 31 Noam Nisan and Mario Szegedy. On the degree of boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994. A preliminary version of this work appeared as an extended abstract in the proceedings of STOC 1992. doi:10.1007/BF01263419.
- 32 Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- 33 Wolfgang J. Paul. A $2.5N$ -Lower Bound on the Combinational Complexity of Boolean Functions. *SIAM Journal on Computing*, 6(3):427–443, 1977. A preliminary version of this work appeared as an extended abstract in the proceedings of the STOC 1975. doi:10.1137/0206030.
- 34 Swagato Sanyal. Fourier Sparsity and Dimension. *Theory of Computing*, 15(11):1–13, 2019. doi:10.4086/TOC.2019.V015A011.
- 35 Palash Sarkar and Subhamoy Maitra. Construction of nonlinear boolean functions with important cryptographic properties. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 485–506. Springer, 2000. doi:10.1007/3-540-45539-6_35.
- 36 Alexander A. Sherstov and Andrey A. Storozhenko. The Communication Complexity of Approximating Matrix Rank. In *Proceedings of the IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 433–462, 2024. doi:10.1109/FOCS61266.2024.00035.
- 37 Karl Wimmer and Yuichi Yoshida. Testing Linear-Invariant Function Isomorphism. In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 840–850, 2013. doi:10.1007/978-3-642-39206-1_71.
- 38 Boyan Yordanov, Jana Tumova, Ivana Cerna, Jiří Barnat, and Calin Belta. Temporal Logic Control of Discrete-Time Piecewise Affine Systems. *IEEE Transactions on Automatic Control*, 57(6):1491–1504, 2012. doi:10.1109/TAC.2011.2178328.