Electronic Colloquium on Computational Complexity, Report No. 139 (2019)

Direct Sum Testing: The General Case

Irit Dinur[1] and Konstantin Golubev[2]

[1]The Weizmann Institute of Science, Israel
[2]ETH Zurich, Switzerland

October 8, 2019

Abstract

A function $f : [n_1] \times \cdots \times [n_d] \to \mathbb{F}_2$ is a direct sum if it is of the form $f(a_1, \ldots, a_d) = f_1(a_1) \oplus \ldots \oplus f_d(a_d)$, for some d functions $f_i : [n_i] \to \mathbb{F}_2$ for all $i = 1, \ldots, d$, and where $n_1, \ldots, n_d \in \mathbb{N}$. We present a 4-query test which distinguishes between direct sums and functions that are far from them. The test relies on the BLR linearity test (Blum, Luby, Rubinfeld, 1993) and on an agreement test which slightly generalizes the direct product test.

In multiplicative ± 1 notation, our result reads as follows. A d-dimensional tensor with ± 1 entries is called a tensor product if it is a tensor product of d vectors with ± 1 entries, or equivalently, if it is of rank 1. The presented tests can be read as tests for distinguishing between tensor products and tensors that are far from being tensor products.

We also present a different test, which queries the function at most $(d + 2)$ times, but is easier to analyze.

1 Introduction

Let us first fix some notations and definitions. By $[n]$ we mean the set $\{0, 1, 2, \ldots, n\}$. For d positive integers $n_1, \ldots, n_d$, we denote $[\overline{n}; d] = [n_1] \times \cdots \times [n_d]$. For two functions $F, G : X \to Y$, we denote by $\mathrm{dist}(F, G)$ the relative Hamming distance between them, namely $\mathrm{dist}(F, G) = \Pr_{x \in X}[F(x) \neq G(x)]$. We say that $F : X \to Y$ is ε-close to have some Property, if there exists a function $G : X \to Y$ such that g has the Property and $\mathrm{dist}(F, G) \leq \varepsilon$.

Given d functions $f_i : [n_i] \to \mathbb{F}_2$, $i = 1, \ldots, d$, where $n_1, \ldots, n_d \in \mathbb{N}$, their direct sum is the function $f : [\overline{n}; d] \to \mathbb{F}_2$ given by $f(a_1, \ldots, a_d) = f_1(a_1) \oplus f_2(a_2) \oplus \ldots \oplus f_d(a_d)$, where $\oplus$ stands for addition is in the field $\mathbb{F}_2$. We denote $f = f_1 \oplus \cdots \oplus f_d$. We study the testability question: given a function $f : [\overline{n}; d] \to \mathbb{F}_2$ test if it is a direct sum, namely if it belongs to the set

$$\mathit{DirectSum}_{[\overline{n};d]} = \{f_1 \oplus \cdots \oplus f_d \mid f_i : [n_i] \to \mathbb{F}_2,\ i = 1, \ldots, d\}.$$

Direct sum is a natural construction that is often used in complexity for hardness amplification [Y82, IJK06, IJKW08, STV01, T03]. It is related to the direct product construction: a function $f : [\overline{n}; d] \to \mathbb{F}_2^d$ is the direct product of $f_1, \ldots, f_d$ as above if $f(a_1, \ldots, a_d) = (f_1(a_1), \ldots, f_d(a_d))$ for all $(a_1, \ldots, a_d) \in [\overline{n}; d]$. The testability of direct products has received attention [GS97, DR06, DG08, IKW12, DS14] as abstraction of certain PCP tests. It was not surprising to find [DDG+17] that there is a connection between testing direct products to testing direct sum. However, somewhat unsatisfyingly this connection was confined to testing a certain type of *symmetric* direct sum. A symmetric direct sum is a function $f : [n]^d \to \mathbb{F}_2$ that is a direct product with all components equal; namely such that there is a single $g : [n] \to \mathbb{F}_2$ such that

$$f(a_1, \ldots, a_d) = g(a_1) \oplus g(a_2) \oplus \cdots \oplus g(a_d).$$

In [DDG+17], a 3-query test was presented for testing if a given f is a symmetric direct sum, and the analysis carried out relying on the direct product test. It was left as an open question to devise and analyze a test for the property of being a (not necessarily symmetric) direct sum.

We design and analyze a four-query test which we call the "square in a cube" test, and show that it is a strong absolute local test for being a direct sum. That is, the number of queries is an absolute constant (namely, 4), and the distance from a function to the subspace of direct sums is bounded by

ISSN 1433-8092

some absolute constant (independent of n and d) times the probability of the failure of the test on this function. We also describe a simpler $(d+1)$-query test, whose easy analysis we defer to section 3.

In order to define the test, we need to introduce the following notation. Given two strings $a, b \in [\overline{n}; d]$ and a set $S \subseteq [d]$, denote by $a_S b$ the string in $[\overline{n}; d]$ whose i-th coordinate equals a_i if $i \in S$ and b_i otherwise.

Test 1 Square in a Cube test. Given a query access to a function $f : [\overline{n}; d] \to \mathbb{F}_2$:

1. Choose $a, b \in [\overline{n}; d]$ uniformly at random.
2. Choose two subsets $S, T \subset [d]$ uniformly at random, and let $U = S \triangle T$ be their symmetric difference.
3. Accept iff
$$f(a) \oplus f(a_S b) \oplus f(a_T b) \oplus f(a_U b) = 0.$$

We prove the following theorem for Test 1.

Theorem 1.1 (Main). *There exists an absolute constant $c > 0$ s.t. for all $d \in \mathbb{N}$ and $n_1, \dots, n_d \in \mathbb{N}$, given $f : [\overline{n}; d] \to \mathbb{F}_2$,*

$$\operatorname{dist}(f, DirectSum_{[\overline{n};d]}) \leq c \cdot \Pr_{a,b,S,T}[f(a) \oplus f(a_S b) \oplus f(a_T b) \oplus f(a_{S \triangle T} b) \neq 0]$$

where a, b are chosen independently and uniformly from the domain of f, and S, T are random subsets of $[d]$.

Our proof, similarly to [DDG+17], relies on a combination of the BLR linearity testing theorem [BLR93] and a direct product test, similar to the one analyzed in [DS14]. These two components were also used in the proof of [DDG+17] for the symmetric case, but here we use the components differently. The trick is to find the right combination. We first observe that once we fix a, b, the test is confined to a set of at most 2^d points in the domain, and can be viewed as performing a BLR (affinity rather than linearity) test on this piece of the domain. From the BLR theorem, we deduce an affine linear function on this piece. The next step is to combine the different affine linear functions, one from each piece, into one global direct sum, and this is done by reducing to direct product.

Testing if a tensor has rank 1. An equivalent way to formulate our question is as a test for whether a d-dimensional tensor with ± 1 entries has rank 1. Indeed moving to multiplicative notation and writing $h_i = (-1)^{f_i}$ and $h = (-1)^f$, we are asking whether there are $h_1, \dots, h_d$ such that

$$h = h_1 \otimes \cdots \otimes h_d.$$

Denoting
$$TensorProduct_{[\overline{n};d]} = \{h_1 \otimes \cdots \otimes h_d \mid h_i : [n_i] \to \{-1, 1\},, i = 1, \dots, d\}$$
we have

Corollary 1.2. There exists an absolute constant $c > 0$ s.t. for all $d \in \mathbb{N}$ and $n_1, \dots, n_d \in \mathbb{N}$, for every $h : [\overline{n}; d] \to \{-1, 1\}$,

$$\operatorname{dist}(h, TensorProduct_{[\overline{n};d]}) \leq c \cdot \Pr_{a,b,S,T}[h(a) \cdot h(a_S b) \cdot h(a_T b) \cdot h(a_{S \triangle T} b) \neq 1].$$

Structure of the Paper. In Sections 2 and 3 we present two different approaches for testing whether a d-dimensional binary tensor is a tensor product. In Section 5 we discuss possible directions for future research. In Section 4, we explain how to derive the specific direct product test that we need from the agreement testing theorem of [DD19]. This is used in the course of the proof in Section 2. The numbering is section-wise. Finally, in Section 5 we discuss possible directions for future research.

2 Square in a Cube Test

In this section we present the Square in a Cube Test. Then we introduce the required background: the BLR test for a function being Affine in Subsection 2.1, the direct product test in Subsection 2.2. Finally, in Subsection 2.3 we prove the main result on the test.

We start by introducing some notation.

Given two vectors $a = (a_1, \dots, a_d)$, $b = (b_1, \dots, b_d) \in [\overline{n}; d]$, define

- $\Delta(a, b) = \{i : a_i \neq b_i\} \subseteq [d]$;
- the induced subcube $C_{a,b}$ is the binary cube $\mathbb{F}_2^{\Delta(a,b)}$;

- the projection map $\rho_{a,b} : C_{a,b} \to [\overline{n}; d]$ defined for $x \in C_{a,b}$ as

$$\rho_{a,b}(x)_i = \begin{cases} a_i = b_i, & i \notin \Delta(a,b); \\ b_i, & i \in \Delta(a,b) \text{ and } x_i = 1; \\ a_i, & i \in \Delta(a,b) \text{ and } x_i = 0; \end{cases}$$

The following test is the same as Test 1 in Introduction.

Test 2 Square in a Cube test. Given a query access to a function $f : [\overline{n}; d] \to \mathbb{F}_2$:

1. Choose $a, b \in [\overline{n}; d]$ uniformly at random.
2. Choose $x, y \in C_{a,b}$ uniformly at random.
3. Query f at $\rho_{a,b}(0), \rho_{a,b}(x), \rho_{a,b}(y)$ and $\rho_{a,b}(x \oplus y)$.
4. Accept iff $f(\rho_{a,b}(0)) \oplus f(\rho_{a,b}(x)) \oplus f(\rho_{a,b}(y)) \oplus f(\rho_{a,b}(x \oplus y)) = 0$.

Theorem 2.1. *Suppose a function $f : [\overline{n}; d]^d \to \mathbb{F}_2$ passes Test 2 with probability $1 - \varepsilon$ for some $\varepsilon > 0$, then f is $O(\varepsilon)$-close to a tensor product.*

2.1 The BLR affinity test

The Blum-Luby-Rubinfeld linearity test was introduced in [BLR93], where its remarkable properties were proven. Later a simpler proof via Fourier analysis was presented, e.g. see [BCH+95]. Below we give a variation of this test for affine functions, see [O'D14, Chapter 1].

Definition 2.2. A function $g : \mathbb{F}_2^d \to \mathbb{F}_2$ is called affine, if there exists a set $S \subseteq [d]$ and a constant $c \in \mathbb{F}_2$ such that for every vector $x \in \mathbb{F}_2^d$

$$g(x) = c \oplus \bigoplus_{i \in S} x_i.$$

Note that (see [O'D14, Exercise 1.26]) a function g is affine iff for any two vectors $x, y \in \mathbb{F}_2^d$ it satisfies

$$g(0) \oplus g(x) \oplus g(y) \oplus g(x \oplus y) = 0. \tag{1}$$

The BLR test implies that if a function $g : \mathbb{F}_2^d \to \mathbb{F}_2$ satisfies (1) with high probability, then it is close to an affine function.

Test 3 The BLR affinity test. Given a query access to a function $f : \mathbb{F}_2^d \to \mathbb{F}_2$:

1. Choose $x \sim \mathbb{F}_2^d$ and $y \sim \mathbb{F}_2^d$ independently and uniformly at random.
2. Query g at $0, x, y$ and $x \oplus y$.
3. Accept if $g(0) \oplus g(x) \oplus g(y) \oplus g(x \oplus y) = 0$.

Theorem 2.3 ([BLR93])**.** *Suppose $g : \mathbb{F}_2^d \to \mathbb{F}_2$ passes the affinity test with probability $1 - \varepsilon$ for some $\varepsilon > 0$. Then g is ε-close to being affine.*

2.2 Generalized Direct Product Test

Definition 2.4. For $k, M, N_1, \ldots, N_k \in \mathbb{N}$, and k functions $g_1, \ldots, g_k : [N_i] \to [M]$, their direct product is the function $g : \prod_i [N_i] \to [M]^k$ denoted $g = g_1 \times \cdots \times g_k$ and defined as $g((x_1, \ldots, x_k)) = (g_1(x_1), \ldots, g_k(x_k))$. A function $g : \prod_i [N_i] \to [M]^k$, is called a direct product if there exist k functions $g_1, \ldots, g_k : [N_i] \to [M]$ such that $g = g_1 \times \cdots \times g_k$ for all $(x_1, \ldots, x_k) \in \prod_i [N_i]$.

Dinur and Steurer [DS14] presented a 2-query test, very similar to Test 4 below, that, with constant probability, distinguishes between direct products and functions that are far from direct product. The proof in [DS14] works for the special case of $N_1 = \cdots = N_k$ and can easily be modified to work for the more general situation. Nevertheless, for completeness, we will rely on a newer and more general agreement theorem of [DD19] that directly implies what we need.

Theorem 2.5 (Generalized direct product testing theorem)**.** *Let $k, M, N_1, \ldots, N_k \in \mathbb{N}$ be positive integers, and let $\varepsilon > 0$. Let $g : \prod_i [N_i] \to [M]^k$ be a function that passes Test 4 with parameter $\alpha = 0.75$ with probability at least $1 - \varepsilon$. Then there exist functions $h_i : [N_i] \to [M]$ such that*

$$\Pr_x [g(x) = (h_1(x), h_2(x), \ldots, h_k(x))] \geq 1 - O(\varepsilon).$$

We will show in Section 4 how to derive the above theorem from the agreement theorem of [DD19].

Test 4 Two-query test $\mathcal{T}(\alpha)$. Given a query access to a function $g : \prod_{i=1}^k [N_i] \to [M]^k$:

- Choose $x \in \prod_{i=1}^k [N_i]$ uniformly.
- For each i, with probability α set $y_i = x_i$ and add i to A, and otherwise choose $y_i \in [N_i]$ uniformly.
- Query g at x and y.
- Accept iff $g(x)_A = g(y)_A$.

2.3 Proof of Theorem 2.1

For a positive integer D, we denote by $\mu_{2/3}(\mathbb{F}_2^D)$ the distribution on $\mathbb{F}_2^D$, where each coordinate, independently, is equal to 0 with probability 1/3 and to 1 with probability 2/3.

We use the following proposition in the course of the proof.

Proposition 2.6. Let $S \subseteq [D]$ be a set and $\chi_S : \mathbb{F}_2^D \to \mathbb{F}_2$ be the corresponding linear function, i.e., $\chi_S(x) = \bigoplus_{i \in S} x_i$. Suppose

$$\Pr_{x \sim \mu_{2/3}(\mathbb{F}_2^D)} (\chi_S(x) = 0) > \frac{2}{3},$$

then $S = \emptyset$.

Proof. Consider $(-1)^{\chi_S}$. Then

$$\Pr_{x \sim \mu_{2/3}(\mathbb{F}_2^D)} (\chi_S(x) = 0) = \Pr_{x \sim \mu_{2/3}(\mathbb{F}_2^D)} \left((-1)^{\chi_S(x)} = 1 \right).$$

Also the following holds

$$\frac{1}{3} < \left| 2 \Pr_{x \sim \mu_{2/3}(\mathbb{F}_2^D)} \left((-1)^{\chi_S(x)} = 1 \right) - 1 \right| = \left| \mathbb{E}_{x \sim \mu_{2/3}(\mathbb{F}_2^D)} (-1)^{\chi_S(x)} \right| =$$

$$\left| \prod_{i \in [D]} \mathbb{E}_{x_i \sim \mu_{2/3}(\mathbb{F}_2)} (-1)^{x_i} \right| = \left| \left(-\frac{1}{3} \right)^{|S|} \right| = \left(\frac{1}{3} \right)^{|S|},$$

and the statement follows. □

Proof. (of Theorem 2.1.) Assume Test 2 rejects a function $f : [\overline{n}; d] \to \mathbb{F}_2$ with probability less than ε, i.e.,

$$\Pr_{\substack{a,b \sim [\overline{n}; d] \\ x,y \sim C_{a,b}}} (f_{a,b}(0) \oplus f_{a,b}(x) \oplus f_{a,b}(y) \oplus f_{a,b}(x \oplus y) = 0) > 1 - \varepsilon,$$

where all distributions are uniform, and $f_{a,b}$ is a shorthand for $f \circ \rho_{a,b}$. Then there exists $a \in [\overline{n}; d]$ such that

$$\Pr_{\substack{b \sim [\overline{n}; d] \\ x,y \sim C_{a,b}}} (f_{a,b}(0) \oplus f_{a,b}(x) \oplus f_{a,b}(y) \oplus f_{a,b}(x \oplus y) = 0) > 1 - \varepsilon.$$

Note that the operations re-indexing the domain $[\overline{n}; d]$[1], as well as *flipping* a function, i.e., adding the constant one function to it element-wise, preserve the distance between functions. Hence, w.l.o.g. we can assume for convenience that $a = (0, \dots, 0)$ and that $f(a) = 0$.

We write C_b for $C_{a,b}$ and f_b for $f_{a,b}$. Then for every $b \in [\overline{n}; d]$,

$$\Pr_{x,y \sim C_b} (f_b(0) \oplus f_b(x) \oplus f_b(y) \oplus f_b(x \oplus y) = 0) = 1 - \varepsilon_b.$$

The BLR theorem (Theorem 2.3) implies that for each $b \in [\overline{n}; d]$ there exists a subset $S(b) \subseteq \Delta(a, b)$, such that

$$\Pr_{x \sim C_b} \left(f_b(x) = \chi_{S(b)}(x) \right) = 1 - \varepsilon_b.$$

Remark 2.7. By the BLR theorem, there should be the "greater or equal to" sign instead of the equality. We assume equality for convenience.

[1]By this we mean selecting permutations π_i on $[n_i]$ for $i = 1, \dots, d$, and setting $f^{\pi_1, \dots, \pi_d}(x_1, \dots, x_d) = f(\pi_1(x_1), \dots, \pi_d(x_d))$

Let $F : [\overline{n}; d] \to \mathbb{F}_2^d$ be a function defined as follows. For each $b \in [\overline{n}; d]$, the set $S(b) \subseteq \Delta(a, b)$ can be viewed as a subset of $[d]$, since $\Delta(a, b) \subseteq [d]$. Then $F(b)$ is defined as the element of $\mathbb{F}_2^d$ corresponding to the set $S(b)$.

We now show that F passes Test 4 with high probability and hence is close to a direct product.

Let $b \in [\overline{n}; d]$ be chosen uniformly at random, and let $b' \in [\overline{n}; d]$ be chosen with respect to the following distribution $D(b)$. For each $i \in [d]$,

$$b'_i = \begin{cases} b_i, & \text{w.p. } 3/4; \\ \text{chosen uniformly at random from } [n] \setminus \{b_i\}, & \text{w.p. } 1/4. \end{cases}$$

Note that the distribution on pairs (b, b'), where b is chosen uniformly from $[\overline{n}; d]$ and b' w.r.t. $D(b)$, is equivalent to the following: for each $i \in [d]$,

$$\begin{cases} b_i = b'_i \text{ chosen uniformly from } [n], & \text{w.p. } 3/4; \\ b_i \neq b'_i \text{ both chosen uniformly from } [n] & \text{w.p. } 1/4. \end{cases} \tag{2}$$

In particular, it is symmetric in the sense that choosing $b' \sim [\overline{n}; d]$ uniformly at random first, and then $b \sim D(b')$, leads to the same distribution on pairs (b, b') as the one described above.

For such a pair (b, b') define distribution $\mathcal{D}_{b,b'}$ on $[\overline{n}; d]$ as follows. For a vector $x \sim \mathcal{D}_{b,b'}$,

$$x_i = \begin{cases} 0, & \text{if } i \in \Delta(b, b'); \\ 0, \quad \text{w.p. } 1/3; & \text{if } i \notin \Delta(b, b'). \\ b_i = b'_i \quad \text{w.p. } 2/3. & \end{cases}$$

Note that the distribution $\mathcal{D}_{b,b'}$ is supported on a binary cube of dimension $d - |\Delta(b, b')|$ inside $[\overline{n}; d]$. Denote

$$\varepsilon_{b,b'} = \Pr_{x \sim \mathcal{D}_{b,b'}} \left(f(x) \neq \chi_{F(b)(x)} \right).$$

We claim that the following holds

$$\varepsilon_b = \Pr_{x \sim C_b} \left(f(x) \neq \chi_{F(b)(x)} \right) = \mathop{\mathbb{E}}_{b' \sim D(b)} \varepsilon_{b,b'}. \tag{3}$$

To see (3) note that since b is chosen uniformly, b' is chosen w.r.t. $D(b)$, and $x \sim \mathcal{D}_{b,b'}$, the resulting distribution for x is

$$x_i = \begin{cases} 0, & \text{w.p. } 1/2; \\ b_i & \text{w.p. } 1/2, \end{cases}$$

which is exactly the uniform distribution on C_b.

We now show that

$$\Pr_{\substack{b \sim [\overline{n}; d] \\ b' \sim D(b)}} \left(\varepsilon_{b,b'} + \varepsilon_{b',b} > \frac{1}{3} \right) < 6\varepsilon \tag{4}$$

First note that it follows from the definitions that

$$\mathop{\mathbb{E}}_{b \sim [\overline{n}; d]} \mathop{\mathbb{E}}_{b' \sim D(b)} \varepsilon_{b,b'} = \mathop{\mathbb{E}}_{b \sim [\overline{n}; d]} \varepsilon_b = \varepsilon.$$

And by the symmetry of the distribution on pairs (b, b'),

$$\mathop{\mathbb{E}}_{b \sim [\overline{n}; d]} \mathop{\mathbb{E}}_{b' \sim D(b)} \varepsilon_{b',b} = \mathop{\mathbb{E}}_{b' \sim D(b)} \mathop{\mathbb{E}}_{b \sim [\overline{n}; d]} \varepsilon_{b',b} = \varepsilon.$$

Combined together, the previous two equations imply that

$$\mathop{\mathbb{E}}_{b \sim [\overline{n}; d]} \mathop{\mathbb{E}}_{b' \sim D(b)} \left(\varepsilon_{b,b'} + \varepsilon_{b',b} \right) = 2\varepsilon,$$

and by the Markov inequality, Inequality 4 follows. By the definition of $\varepsilon_{b,b'}$,

$$\Pr_{x \sim \mathcal{D}_{b,b'}} \left(\chi_{F(b)}(x) = \chi_{F(b')}(x) \right) > 1 - \left(\varepsilon_{b,b'} + \varepsilon_{b',b} \right).$$

which is equivalent to

$$\Pr_{x \sim \mathcal{D}_{b,b'}} \left(\chi_{F(b) \Delta F(b')}(x) = 1 \right) > 1 - \left(\varepsilon_{b,b'} + \varepsilon_{b',b} \right).$$

Proposition 2.6 implies that if $1 - (\varepsilon_{b,b'} + \varepsilon_{b',b}) > \frac{2}{3}$, then

$$F(b)_{C_b \cap C_{b'}} = F(b')_{C_b \cap C_{b'}}.$$

By Theorem 2.5, the function $F : [\overline{n}; d] \to \mathbb{F}_2^d$ is close to a direct product, i.e., there exist d functions $F_1, \ldots, F_d : [n] \to \mathbb{F}_2$ such that

$$\Pr_{b \sim [\overline{n};d]} (F(b) = (F_1(b_1), \ldots, F_d(b_d))) \geq 1 - O(\varepsilon).$$

Therefore,

$$\Pr_{b \sim [\overline{n};d]} \left(f(b) = \bigoplus_{i=1}^{d} F_i(b_i) \right) \geq 1 - O(\varepsilon).$$

□

3 The Shapka Test

In this section we present a different test for whether a tensor is a tensor product. It queries the tensor at $(d+2)$ places at most, but the proof is simpler than for the previous test.

In [KL14], Kaufman and Lubotzky showed an interesting connection between the theory of high-dimensional expanders and property testing. Namely, they showed that $\mathbb{F}_2$-coboundary expansion of a 2-dimensional complete simplicial complex implies testability of whether a symmetric $\mathbb{F}_2$-matrix is a tensor square of a vector. The following test is inspired by their work and in a way generalizes it. However, since the description below does not employ neither terminology nor machinery of high-dimensional expanders, we refer to [KL14] for the connection between this theory and property testing.

Given two strings $a, b \in [\overline{n}; d]$, for $i \in [d]$ denote by $a_b^i \in [\overline{n}; d]$ the vector which coincides with a in every coordinate except for the i-th one, where it coincides with b, i.e.,

$$(a_b^i)_j = \begin{cases} a_j, & \text{if } j \neq i; \\ b_i, & \text{if } j = i. \end{cases}$$

For a string $a \in [\overline{n}; d]$, and a number $x \in [n_i]$, we write a_x^i for the string which is equal to a in every coordinate except for the i-th one, where it is equal to x, i.e.,

$$a_x^i = (a_1, \ldots, a_{i-1}, x, a_{i+1}, \ldots, a_d).$$

Test 5 The Shapka Test. Given a query access to a function $f : [\overline{n}; d] \to \mathbb{F}_2$:

1. Choose $a, b \in [\overline{n}; d]$ uniformly at random.
2. Define the query set $Q_{a,b} \subseteq [\overline{n}; d]$ to consist of a, a_b^j for all $j \in [d]$, and also b if d is even.
3. Query f at the elements of $Q_{a,b}$.
4. Accept iff $\bigoplus_{q \in Q_{a,b}} f(q) = 0$.

Remark 3.1. Shapka is the Russian word for a winter hat (derived from Old French *chape* for a *cap*). The name *the Shapka test* comes from the fact that the set $Q_{a,b}$ consists of the two top layers of the induced binary cube $C_{a,b}$ (and also the bottom layer if d is even).

Theorem 3.2. *Suppose a function $f : [\overline{n}; d] \to \mathbb{F}_2$ passes Test 8 with probability $1 - \varepsilon$ for some $\varepsilon > 0$, then f is ε-close to a tensor product.*

Proof. Let δ be the relative Hamming distance from f to the subspace of direct sums, i.e., for every direct sum $g : [\overline{n}; d] \to \mathbb{F}_2$ it holds that

$$Pr_{x \sim [\overline{n};d]} (f(x) \neq g(x)) \geq \delta.$$

For a vector $a \in [\overline{n}; d]$, let us define the local view of f from a, that is d functions $f_1^a, \ldots, f_d^a$, where $f_i^d : [n_i] \to \mathbb{F}_2$, $i = 1, \ldots, d$, that are defined as follows. For $1 \leq i \leq d-1$, and $x \in [n_i]$,

$$f_i^a(x) = f(a_x^i).$$

For $i = d$, the definition of $f_d^a : [n_d] \to \mathbb{F}_2$ depends on the parity of d and goes as follows

$$\begin{cases} f_d^a(x) = f(a_x^d), & \text{if } d \text{ is odd}, \\ f_d^a(x) = f(a_x^d) \oplus f(a), & \text{if } d \text{ is even}. \end{cases}$$

Given a collection of d functions, $g_i : [n_i] \to \mathbb{F}_2$, $i = 1, \ldots, d$, recall that their direct sum is the function $g_1 \oplus \cdots \oplus g_d$ such that for a vector $x \in [\overline{n}; d]$ the following holds

$$g_1 \oplus \cdots \oplus g_d = \bigoplus_{i \in [d]} g_i(x_i).$$

The following holds for any $[\overline{n}; d]$,

$$(f - f_1^a \oplus \cdots \oplus f_d^a))\,(b_1, \ldots, b_d) = \bigoplus_{q \in Q_{a,b}} f(q). \tag{5}$$

As $f_1^a \oplus \cdots \oplus f_d^a$ is a direct sum, it is at least δ-far from f, and hence for any $a \in [\overline{n}; d]$,

$$\Pr_{b \sim [\overline{n}; d]} ((f - f_1^a \oplus \cdots \oplus f_d^a)\,(b) = 1) \geq \delta. \tag{6}$$

Assume now that f fails Test 8 with probability ε, i.e.,

$$\varepsilon = \Pr_{a,b \sim [\overline{n}; d]} \left(\bigoplus_{q \in Q_{a,b}} f(q) = 1 \right).$$

Combining this equality with (5) and (6), we get the following

$$\varepsilon = \mathop{\mathbb{E}}_{a \sim [\overline{n}; d]} \Pr_{b \sim [\overline{n}; d]} \left(\left(f - f_a^1 \oplus \cdots \oplus f_a^d \right) (b_1, \ldots, b_d) = 1 \right) \geq \left(\mathop{\mathbb{E}}_{a \sim [\overline{n}; d]} \delta \right) = \delta,$$

which completes the proof. □

4 Generalized direct product test

In this section we prove Theorem 2.5, restated directly below, by relying on known agreement test results.

Theorem 2.5 (restated) *Let $k, M, N_1, \ldots, N_k \in \mathbb{N}$ be positive integers, and let $\varepsilon > 0$. Let $g : \prod_i [N_i] \to [M]^k$ be a function that passes Test 4 with parameter $\alpha = 0.75$ with probability at least $1 - \varepsilon$. Then there exist functions $h_i : [N_i] \to [M]$ such that*

$$\Pr_x [g(x) = (h_1(x), h_2(x), \ldots, h_k(x))] \geq 1 - O(\varepsilon).$$

This theorem was proven "in spirit" in [DS14] although formally that proof is written only for the case of $N_1 = N_2 = \cdots = N_k$. Instead of reworking the details we will rely on a newer work that generalizes the [DS14] paper to a broader context of agreement testing.

First, let us move from the distribution of Test 4 to a related distribution. It turns out that if g passes one of these two-query tests with good probability then we can draw conclusions regarding its success in related tests.

Test 6 Two-query test with fixed intersection size $\mathcal{T}(t)$. Given $g : \prod_i [N_i] \to [M]^k$

- Choose $x \in \prod_i [N_i]$ uniformly.
- Choose a subset $T \subset [k]$ of size t uniformly.
- Choose $y \in \prod_i [N_i]$ uniformly conditioned on $y|_T = x|_T$.
- Accept iff $g(x)|_T = g(y)|_T$.

Claim 4.1. Suppose g passes Test 4 with $\alpha = 0.75$ with probability $1 - \varepsilon$ then it passes Test 6 with parameter $k/10 < t < k/4$ probability $1 - O(\varepsilon)$.

We prove this claim later in Section 4.1. Theorem 2.5 will follow by invoking a theorem from [DD19] about agreement testing. In agreement testing the input is a collection of local functions each defined on its own small domain. The agreement test checks that whenever the small domains overlap the functions agree with each other. An agreement theorem deduces a single global function (on a domain that contains all the smaller ones) from the given local pairwise agreements. To see who are the small domains in our context let us construct the following set system.

- Vertices: Let $V_1, \ldots, V_k$ be k disjoint sets of vertices, $|V_i| = N_i$ and we identify V_i with $[N_i]$.
- Subsets: We have a subset for every choice of one element from each V_i,

$$\mathcal{S} = \{\{v_1, \ldots, v_k\} \;:\; \forall i = 1, \ldots, k,\; v_i \in V_i\}.$$

  There is a straightforward bijection between $\mathcal{S}$ and the domain of g, namely $\prod_i [N_i]$.
- Local functions: For a set $S = \{v_1, \ldots, v_k\} \in \mathcal{S}$ we have a local function $f_S : S \to [M]$ defined by

$$f_S(v_i) = g(\bar{v}_1, \ldots, \bar{v}_k)_i$$

  where $\bar{v}_i \in [N_i]$ is associated with v_i in the identification of V_i and $[N_i]$.

A direct product function $g : \prod_i [N_i] \to [M]^k$ can thus be represented as a collection $\{f_S\}$ of local functions. The direct product test, Test 6, can be rephrased as Test 7 below. Given $g : \prod_i [N_i] \to [M]^k$ we view it as a family of local functions $\{f_S\}$ and would like to invoke the following agreement test theorem,

Test 7 Two-query test $\mathcal{T}(t)$. Given a family of local functions $\{f_S \in [M]^S \;:\; S \in \mathcal{S}\}$

- Choose a set $S_1 \in \mathcal{S}$ uniformly.
- Choose a subset $T \subset S_1$ of size t uniformly.
- Choose $S_2 \in \mathcal{S}$ uniformly conditioned on $S_2 \supset T$.
- Accept iff $f_{S_1}|_T = f_{S_2}|_T$.

Theorem 4.2 ([DD19, Theorem 4.4])**.** *Suppose $\mathcal{S}$ is a collection of subsets that are top faces of a λ-one-sided k-partite $\frac{1}{k^3}$-high dimensional expander. Then given $\{f_S\}$ for which Test 7 succeeds with probability $1-\varepsilon$, and assuming $t < k/4$, there exists a function $h : V_1 \sqcup \cdots \sqcup V_k \to [M]$ such that*

$$Pr_{S\in\mathcal{S}}[f_S = h|_S] \geq 1 - O(\varepsilon).$$

We will show in Section 4.2 that we are justified to apply this theorem because our collection of subsets, also known as the "complete multi-partite complex", is a λ-one-sided-HDX for any $\lambda \geq 0$. Assuming this is the case, we can now take $h_i = h|_{V_i}$ and get the desired conclusion of Theorem 2.5,

$$\Pr_x[g(x) = (h_1(x_1), \ldots, h_k(x_k))] = \Pr_S[f_S = h|_S] \geq 1 - O(\varepsilon).$$

4.1 Moving between different variants of agreement tests

Claim 4.1 follows immediately from the following lemma, (one needs to apply the first item 3 times to get from $\alpha = 0.75$ to α^2 then α^4 and then $\alpha^8 < 0.25$ and then item 2 once).

Lemma 4.3. *Let $g : \prod_i [N_i] \to [M]^k$ be a function that passes Test 4 with parameter α with probability at least $1-\varepsilon$. Then,*

- *g passes Test 4 with parameter α^2 with probability at least $1 - 2\varepsilon$.*
- *There exists a number t, $\alpha k - \sqrt{k} \leq t \leq \alpha k + \sqrt{k}$, such that g passes Test 6 with parameter t with probability at least $1 - O(\varepsilon)$*

Proof. We first prove the first item. Choosing two queries x, y according to the test distribution in Test 4 and then another pair x, y' conditioned on the first query being x, we get a pair y, y' whose distribution is exactly as if the were chosen from Test 4 with parameter α^2. Suppose A was the set of indices in which y_i was chosen to equal X_i, and suppose A' was that set for the pair x, y'. Setting $B = A \cap A'$ it remains to notice that the event that $g(y)|_B \neq g(y')|_B$ is contained in at least one of the events $g(y)|_A \neq g(x)|_A$ or $g(y')|_{A'} \neq g(x)|_{A'}$, so its probability is at most 2ε.

For the second item, observe that with probability $p > 0.1$ the size of the set A defined by the test is some t such that $\alpha k - \sqrt{k} \leq t \leq \alpha k + \sqrt{k}$ (this follows from Hoefding's tail bound). There must be some t in this range for which the failure probability of the test is at most $2\varepsilon/p$. Otherwise, even if the test succeeds with probability 1 when t is outside this range, we would still not be able to reach a sucess probability of $1 - \varepsilon$ since

$$\Pr[fail] \geq p \cdot 2\varepsilon/p > \varepsilon$$

□

4.2 The complete multi-partite complex

The collection of subsets defined in the beginning of this section gives rise to the so-called complete multi-partite simplicial complex, by downwards closing that set system.

We wish to show that it satisfies the requirements of Theorem 4.2. For this we briefly recall the relevant definitions. For a more comprehensive introduction to this topic we refer the reader to [DD19] and the references therein.

- Simplicial Complex: A simplicial complex is a hypergraph that is closed downward with respect to containment. It is $(d-1)$-dimensional if the largest hyperedge has size d. We refer to $X(\ell)$ as the hyperedges (also called faces) of size $\ell+1$. $X(0)$ are the vertices. It is d-partite if the vertices are partitioned into d parts, and each hyperedge in $X(d-1)$ has one vertex from each part.
- Link: Given a i-face σ, the link of σ is the collection of faces that are disjoint from σ and whose union belongs to X,
$$X_\sigma = \{\tau \in X \;:\; \tau \cap \sigma = \phi \text{ and } \tau \cup \sigma \in X\}\,.$$
This is a simplicial complex whose dimension is $dim(X) - |\sigma| - 1$.
- Distribution: Given any probability distribution on the top faces $X(d-1)$, it propagates to a distribution on the edges by selecting a top face and then a pair of vertices in it uniformly. This gives a weighted graph that is called the **1-skeleton** of the complex.
- HDX: A $(d-1)$-dimensional simplicial complex is a λ-one-sided HDX if for every face $\sigma \in X(t)$, $t \le d-3$, the 1-skeleton of the link X_σ is a λ-one-sided expander graph, meaning that the random walk Markov chain on this weighted graph has all non-trivial normalized eigenvalues at most λ.

The complete d-partite complex has parameters $n_1, \ldots, n_d$ and has a vertex set V_i of size n_i. It is defined by the following distribution over d-hyperedges: For each i choose $x_i \in V_i$ uniformly. This gives a probability distribution on faces $\{x_1, \ldots, x_d\}$ in $X(d-1)$. The 1-skeleton of this complex is a graph whose vertices are $V_1 \sqcup \cdots \sqcup V_d$ and whose weighted edges are obtained by selecting a random hyperedge in $X(d)$ and then a random pair of vertices inside it. The link of a face in this complex is itself a complete partite complex, with fewer parts. To show that this complex is a λ-one-sided HDX it remains to prove the following lemma,

Lemma 4.4. *Let G be the 1-skeleton of a complete d-partite complex with parameters $n_1, \ldots, n_d$. Then the normalized adjacency matrix of G has one eigenvalue of 1, eigenvalue of 0 with multiplicity $\sum_i n_i - d$, and the remaining $(d-1)$ eigenvalues have value $-1/(d-1)$.*

In particular, except for one eigenvalue of 1, all of G's remaining eigenvalues are non-positive.

Proof. Let, as before, V_i denote the part of vertices of size n_i. The distribution on edges induced by the uniform distribution on the maximal faces is as follows. For an edge (v_i, v_j), where $v_i \in V_i$, $v_j \in V_j$ and $i \neq j$, its probability is equal to

$$p(v_i, v_j) = p_{i,j} = \frac{1}{\binom{d}{2} n_i n_j}.$$

Hence the transition probability of moving from the vertex v_i to the vertex v_j is equal to

$$\frac{p_{i,j}}{\sum_{j=1,\, j\neq i}^{d} n_j p_{i,j}} = \frac{p_{i,j}}{2/(dn_i)} = \frac{1}{(d-1)n_j},$$

The transition matrix is of the following form

$$A = \frac{1}{d-1}\begin{bmatrix} 0 & \frac{1}{n_2}J_{n_1\times n_2} & \frac{1}{n_3}J_{n_1\times n_3} & \cdots & \frac{1}{n_d}J_{n_1\times n_d} \\ \frac{1}{n_1}J_{n_2\times n_1} & 0 & \frac{1}{n_3}J_{n_2\times n_3} & \cdots & \frac{1}{n_d}J_{n_2\times n_d} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \frac{1}{n_1}J_{n_d\times n_1} & \frac{1}{n_2}J_{n_d\times n_2} & \frac{1}{n_3}J_{n_d\times n_3} & \cdots & 0 \end{bmatrix},$$

where $J_{n_i \times n_j}$ stands for the all-one matrix of size $n_i \times n_j$. In order to show that A has a single positive eigenvalue, we use the approach developed in [EH80]. First, note that the multiplicity of 0 is $n-d$, where $n = \sum_{i=1}^{d} n_i$, because the matrix A is of rank $n-d$. Next, note that if f is an eigenfunction with eigenvalue $\lambda \neq 0$, then

1. it is constant on V_i for each $i = 1, \ldots, d$;
2. and
$$\lambda \alpha_i = \frac{1}{d-1} \sum_{j=1,\, j\neq i}^{d} \alpha_j,$$
where α_i is the value of f on V_i.

For $v \in V_i$,

$$\lambda f(v) = \frac{1}{d-1} \sum_{j=1, j\neq i}^{d} \left(\frac{1}{n_j} \sum_{u \in V_j} f(u) \right).$$

The expression on r.h.s. is the same for every $v \in V_i$, and $\lambda \neq 0$, which completes the proof of (1). To show (2), it is enough to substitute $f(u) = \alpha_j$ for $u \in V_j$ in the equality above.

It follows from the above that the non-zero eigenvalues of A are exactly the eigenvalues of the matrix

$$\frac{1}{d-1} \left(J_{d\times d} - I_{d\times d} \right),$$

which has eigenvalue 1 with multiplicity 1, and $-\frac{1}{d-1}$ with multiplicity $(d-1)$. □

5 Further Directions

Below we present possible directions for future research.

1. Can the original function $f : [\overline{n}; d] \to \mathbb{F}_2$ be reconstructed by a voting scheme using the Shapka Test 8?
2. It is plausible that the Square in the Cube test 2 can be analyzed by the Fourier transform approach similarly to the analysis of the BLR test.
3. Another test in the spirit of the paper is the following. We conjecture that this test is also good,

Test 8 Given a query access to a function $f : [\overline{n}; d] \to \mathbb{F}_2$:

(a) Choose $a, b \in [\overline{n}; d]$ uniformly at random.
(b) Choose $x \in C_{a,b}$ uniformly at random.
(c) Query f at $\rho_{a,b}(0), \rho_{a,b}(x), \rho_{a,b}(1)$ and $\rho_{a,b}(x \oplus 1)$.
(d) Accept iff $f(\rho_{a,b}(0)) \oplus f(\rho_{a,b}(x)) \oplus f(\rho_{a,b}(1)) \oplus f(\rho_{a,b}(x \oplus 1)) = 0$.

i.e., if a function passes the test with high probability then it is close to a tensor product.

Acknowledgements

The authors would like to thank Oded Goldreich for pointing out a gap in the proof in a previous version of this manuscript.

The first author is supported by ERC-CoG grant number 772839. A substantial part of the work was done while the second author held a joint postdoctoral position at The Weizmann Institute and Bar-Ilan University funded by the ERC grant number 336283. Currently, the second author is supported by the SNF grant number 200020_169106. The second author would also like to thank the Swiss Mathematical Society for travel funding related to this paper.

References

[BCH+95] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. Linearity testing in characteristic two. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 432–441, 1995.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of computer and system sciences*, 47(3):549–595, 1993.

[DDG+17] Roee David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. Direct sum testing. *SIAM J. Comput.*, 46(4):1336–1369, 2017.

[DG08] Irit Dinur and Elazar Goldenberg. Locally testing direct products in the low error range. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, 2008.

[DR06] Irit Dinur and Omer Reingold. Assignment testers: Towards combinatorial proofs of the PCP theorem. *SIAM Journal on Computing*, 36(4):975–1024, 2006. Special issue on Randomness and Computation.

[DD19] Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In *60th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2019.

[DS14] Irit Dinur and David Steurer. Direct product testing. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 188–196, 2014.

[EH80] Friedrich Esser and Frank Harary. On the spectrum of a complete multipartite graph. In *European Journal of Combinatorics*, 1(3), 211–218, 1980.

[GS97] Oded Goldreich and Shmuel Safra. A combinatorial consistency lemma with application to proving the PCP theorem. In *RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science*. LNCS, 1997.

[IJK06] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Approximately listdecoding direct product codes and uniform hardness amplification. In *Proc. 47th IEEE Symp. on Foundations of Computer Science*, 187–196, 2006.

[IJKW08] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. In *Proc. 40th ACM Symp. on Theory of Computing*, 39(4), 1637–1665, 2008.

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query PCPs. *SIAM J. Comput.*, 41(6):1722–1768, 2012.

[KL14] Tali Kaufman and Alexander Lubotzky. High dimensional expanders and property testing. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 501–506, New York, NY, USA, 2014. ACM.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions*.

[STV01] Madhu Sudan, Luca Trevisan and Salil Vadhan. Pseudorandom Generators without the XOR Lemma. In *Journal of Computer and System Sciences*, 62(2), pages 236–266, 2001.

[T03] Luca Trevisan. List-decoding using the XOR lemma. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003.* Proceedings., pages 126–135, 2003.

[Y82] A. C. Yao. Theory and application of trapdoor functions. 23rd Annual Symposium on Foundations of Computer Science (sfcs 1982), Chicago, IL, USA, 1982, pp. 80-91.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il