Direct Product Testing

Irit Dinur*   David Steurer†

December 23, 2013

Abstract

A *direct product* is a function of the form $g(x_1, \ldots, x_k) = (g_1(x_1), \ldots, g_k(x_k))$. We show that the direct product property is locally testable with 2 queries, that is, a canonical two-query test distinguishes between direct products and functions that are from direct products with constant probability.

This local testing question comes up naturally in the context of PCPs, where direct products play a prominent role for gap amplification. We consider the natural two query test for a given function $f : [N]^k \to [M]^k$

> **Two query direct product test**: Choose x, y that agree on a random set A of t coordinates and accept if $f(x)_A = f(y)_A$.

We provide a comprehensive analysis of this test for all parameters $N, M, k, t \leqslant O(k)$ and success probability $\delta > 0$.

Our main result is that if a given function $f : [N]^k \to [M]^k$ passes the test with probability $\delta \geqslant 1 - \varepsilon$ then there is a direct product function g such that $\mathbb{P}[f(x) = g(x)] \geqslant 1 - O(\varepsilon)$. This is the first result relating success in the (or any) test to the fraction of the domain on which f is equal to a direct product function. This test has been analyzed in previous works for the case $t \ll k \ll N$, and results show closeness of f to a direct product under a less natural measure of "approximate agreement".

In the small soundness "list-decoding" regime, we show that if the test above passes with probability $\delta \geqslant \exp(-k)$, then the function agrees with a direct product function on local parts of the domain. This extends the previous range of parameters of $\delta \geqslant \exp(-\sqrt[3]{k})$ to the entire meaningful range of $\delta > \exp(-k)$.

Keywords: parallel repetition, direct product testing.

*Department of Computer Science and Applied Mathematics, Weizmann Institute. Part of this work was done at Microsoft Research New England and Radcliffe Institute for Advanced Study.

†Computer Science Department, Cornell University. Part of this work was done at Microsoft Research New England.

1 Introduction

The direct product of functions $g_1, \dots, g_k : [N] \to [M]$ is defined as the function $g = g_1 \times \cdots \times g_k : [N]^k \to [M]^k$ given by $g(x_1, \dots, x_k) = g_1(x_1), \dots, g_k(x_k)$. Product operations are studied extensively in various domains, including products of *functions* as above, products of *graphs*, and products of (one-round two-player) *games*—known as parallel repetition. In this work we resolve the following question about the structure of direct products:

> "Does there exist a two-query test that distinguishes between direct products and functions that are far from direct products with constant probability?"

This question comes up naturally in the context of PCPs where products play an important role for gap amplification [Raz98, Din07]. In the analysis of PCPs, one is given information about local consistency and wants to deduce a global direct product behavior. Thus, the direct product (DP) testing question lies at the heart of many questions regarding products, for CSPs, for two-player games, and possibly also for graphs. Its relation to PCP constructions is analogous to that of the low degree testing question. Indeed, past results analyzing direct product tests have had direct applications for constructions of PCPs [DR06, IKW12, DM11].

Our main result is an affirmative answer to the question above. We consider the following canonical two-query test:

> $\mathcal{T}(t)$ **- Two-query test with intersection** t
>
> Choose a random set $A \subset [k]$ of size t. Choose $x, y \in [N]^k$ uniformly at random, conditioned on $x_A = y_A$. Accept iff $f(x)_A = f(y)_A$.

We give a comprehensive analysis of this test for all parameters $N, M, k, t \leqslant O(k)$,

Theorem 1.1. *Let k, N, M be positive integers, let $t \leqslant k/2$, and let $\varepsilon > 0$. Let $f : [N]^k \to [M]^k$ be given such that*

$$\mathop{\mathbb{P}}_{A,x,y}[f(x)_A = f(y)_A] \geqslant 1 - \varepsilon$$

where x, y, A are chosen from the test distribution $\mathcal{T}(t)$. Then, there exists a direct product function $g = g_1 \times \cdots \times g_k$ such that $\mathbb{E}_x[dist(f(x), g(x))] = O(\varepsilon k/t)$, where $dist(\cdot, \cdot)$ denotes the Hamming distance. In particular, when $t = k/2$ this implies

$$\mathop{\mathbb{P}}_{x}[f(x) = g(x)] \geqslant 1 - O(\varepsilon)$$

Moreover, if f is rotation-invariant[1]*, then g is rotation invariant, i.e. $g_1 = \cdots = g_k$.*

This is the first testing result to show that f is equal to a direct product function on most of its domain. Previously, results on DP testing showed closeness of f to a direct product under a less natural measure of closeness, namely, where f and g agree on "most of their coordinates" [GS97, DR06, DG08, IKW12]. In addition, previous results only handled the case $t \ll k \ll N$ whereas we allow any N, M, k and any $t \leqslant O(k)$. We remark that the approximate notion of agreement that is present in all of the previous works is inherent for $t \ll k$.[2]

Our proof of Theorem 1.1 begins by first establishing direct product structure for the function on local pieces, and then merging them all.

[1] A function is rotation-invariant if $\sigma f(x) = f(\sigma x)$ for all x and permutations σ, see Section 2.2.

[2] Consider a function f obtained by taking $g_1 \times \cdots \times g_k$ and for each x changing one coordinate of $f(x)$ independently at random. This function still passes $\mathcal{T}(t)$ with probability at least $(1 - t/k)^2 = 1 - O(t/k)$.

The result on local direct product structure may be of independent interest and can be viewed as a DP testing result in the so-called list decoding regime. We show (see Lemma 1.2 below) that if a function passes the test with any non-negligible probability δ (and here δ can be as small as $\exp(-k)$) then it must have some local direct product structure. This extends the previous range of $\delta \geqslant \exp(-\sqrt[3]{k})$ for which this result was known [IKW12].

Our analysis is related to a recent work of the authors [DS13] on parallel repetition (and in fact, that work started from working on the direct product question). Similarly to [DS13], we make technical gains by comparing how the given function f behaves on a family of related two-query tests,

> $\mathcal{T}(a,b)$ **- Two-query test with parameters** a **and** b **:**
>
> Choose disjoint subsets $A, B \subset [k]$ of sizes a and b and let $C = [k] - (A \cup B)$. Choose u, v, w, w' uniformly at random such that $u \in [N]^A$, $v \in [N]^B$ and $w, w' \in [N]^C$. Accept if $f(uvw)_A = f(uvw')_A$.

Clearly, this test generalizes the previous one, with $a = t, b = 0$. The extra b parameter provides a way to compare between the acceptance probability of the test on nearby values of a, b. We denote by $win(a,b)$ the success probability of the test with parameters a and b. Whereas the difference between $win(t,0)$ and $win(t+1,0)$ is hard to control, we show through a hybrid argument that there are values a, b for which $win(a,b) \approx win(a+1,b) \approx win(a,b+1)$. This approximate equality is quite powerful. Intuitively, it means that conditioned on already finding agreement on some a coordinates, the value of the function on one extra coordinate is pretty much determined already (because whether or not we check consistency there will hardly affect the acceptance probability). This is the key to our local direct product lemma.

In this lemma we partition the event that $\mathcal{T}(a,b)$ accepts into smaller events '$\mathcal{T}(a,b)$ accepts on s' where s is a tuple $s = (A, B, u, v, \alpha)$ and the event '$\mathcal{T}(a,b)$ accepts on s' occurs iff $\mathcal{T}$ chooses A, B, u, v and $f(uvw)_A = \alpha = f(uvw')$. For each s we define (by plurality) a direct product function $\beta_s : [N]^C \to [M]^C$ where $C = [k] - (A \cup B)$ and where $\beta_s = (\beta_{s,i})_{i \in C}$ for $\beta_{s,i} : [N] \to [M]$. We prove that the event 'accept on s' coincides, up to an η margin of error, with the event that f agrees with β_s. Equivalently, we assign each s a probability proportional to $\mathbb{P}[\mathcal{T}(a,b)$ accepts on $s]$ and show that in expectation over s, the probability of f being equal to a direct product, conditioned on the event 'accept on s', is close to 1.

Lemma 1.2 (Local Structure)**.** *Let $k, N, M, t \leqslant k/2$ be positive integers, let $\delta > 0$, and set $\eta = 1 - \delta^{2/t}$. Let $f : [N]^k \to [M]^k$ be given such that*

$$\mathbb{P}[\mathcal{T}(t) \text{ accepts } f] \geqslant \delta = (1-\eta)^{t/2}.$$

Then, there are values $t/2 \leqslant a \leqslant t$ and $b \leqslant t/2$ such that $\delta \leqslant win(a,b)$ and, a direct product function $\beta_s : [N]^C \to [M]^C$ for each $s = ABuv\alpha$ such that,

$$\mathop{\mathbb{E}}_{s \sim \mu} \mathop{\mathbb{E}}_{i \in C} \mathop{\mathbb{P}}_{z|s}[f(z)_i = \beta_{s,i}(z_i)]^2 \geqslant 1 - O(\eta),$$

where s is a tuple $s = ABuv\alpha$ distributed according to μ given by $\mu(s) = \mathbb{P}[\mathcal{T}(a,b)$ accepts on $s] / \mathbb{P}[\mathcal{T}(a,b)$ accepts$]$. Moreover, if f is rotation-invariant then β_s is rotation-invariant for each s.

The step in our analysis in which we deduce a direct product structure based on $win(a,b) \approx win(a,b+1) \approx win(a+1,b)$ is reminiscent of a part of the analysis of Feige and Kilian, in their work on parallel repetition [FK94] (which was later followed in [DG08] for obtaining a direct product testing result). However, the quantities studied in [FK94] are more like $win(a, 1)$ and $win(a, 0)$ which are very hard to relate, and in particular, the relationship proven in [FK94] has an additive and not multiplicative error. This limits the [FK94, DG08] analysis to values of δ that are polynomial in k and not exponential.

We remark that there is one interesting aspect in which the result of [IKW12] is more powerful than our lemma. [IKW12] show that conditioned on the test accepting, the probability of f agreeing with a direct product is $1 - \text{poly}(\delta)$, whereas in our lemma it is only $1 - O(\eta) = 1 - O(\delta^{2/t})$. The stronger conclusion allows them to lift their DP testing result into a PCP construction, and also to extend the local structure to a global structure by adding another query to the test (getting a "Z-test").

Finally, we mention that a well-studied variant of the direct product test is one that randomly permutes the entries of x before querying f, and then permutes back the answer. This twist is essentially equivalent to forcing the given f to be itself permutation invariant, in that $f(\sigma x) = \sigma f(x)$. Our results immediately extend to this test, and we describe this in Section 2.2.

Related Work

Goldreich and Safra pioneered the study of direct product tests as a combinatorial alternative to low degree tests, potentially leading to constructions of PCPs [GS97]. They proved an analog of Theorem 1.1 for a derandomized test, and for $t = \sqrt{k}$. The two query test above[3] was analyzed in [DR06] for $t = k^{2/3}$, where it was used for a PCP construction. Further work was done in [DG08, IKW12] on the list decoding regime, but always with intersection size t that is some polynomial fraction of k. As already mentioned, for such intersection parameters the conclusion $f(x) = g(x)$ cannot be expected for any non-negligible fraction of the domain and must be weakened to agreement of $f(x)$ and $g(x)$ on 'most of their coordinates'.

A very recent application of our Theorem 1.1 is in the study of direct sum testing [DDG+13]. The k-fold direct *sum* of a (Boolean, i.e. $M = 2$) function $g : [N] \to \{0, 1\}$ is a function $g^{(\oplus k)}$ that maps to each k-tuple $x_1, \ldots, x_k$ a single bit $\sum_i g(x_i) \mod 2$. This has the extra feature of maintaining the alphabet size, a desirable feature for PCP constructions, that is usually achieved by an additional alphabet-reduction step. It was shown in [DDG+13], that a certain natural three query test distinguishes between direct sums and functions that are far from direct sums. The proof is by reduction to the direct product question, and then invoking our main theorem. The high sensitivity of the XOR function to even one bit-flip requires the stronger notion of exact agreement rather than an approximate agreement.

We remark that the work of [IKW12] also contained a very elegant derandomization for the DP test. It would be very interesting to see if that construction can be strengthened to give a tester as in our Theorem 1.1, but derandomized. This could potentially lead to a PCP construction that requires no alphabet reduction step.

Organization

We begin with some preliminaries and a quick discussion of the rotation-invariant version of the direct product test. We prove Lemma 1.2 in Section 3 and then Theorem 1.1 in Section 4. In the appendix we add some treatment of the direct product test for sets instead of tuples (Section A), and of the relationship to parallel repetition (Section B).

[3] To be precise, [DR06] analyzed a smooth version of this test, where the intersection size is t in expectation.

2 Preliminaries

2.1 Notation

We represent a mapping of $[N]^k$ to $[M]^k$ by a function $f : [N]^k \times [M]^k \to \mathbb{R}_{\geqslant 0}$ such that $z \mapsto \alpha$ then $f(z;\alpha) = 1$ and $f(z,\beta) = 0$ for all $\beta \neq \alpha$. More generally, we consider functions $f : [N]^k \times [M]^k \to \mathbb{R}_{\geqslant 0}$ such that for each $z \in [N]^k$, $\sum_{\alpha \in [M]^k} f(z;\alpha) = 1$. These represent randomized mappings. For any $z \in [N]^k$ we let $f(z)$ be the random variable that equals α with probability $f(z;\alpha)$.

For a string $z \in [N]^k$ and a subset $A \subset [k]$ we denote $z_A \in [N]^A$ the substring obtained by restricting z to the coordinates of A. Given $u \in [N]^A, v \in [N]^B, w \in [N]^C$ we denote $z = uvw$ the string in $[N]^k$ obtained by concatenating u, v, w appropriately, so that $z_A = u, z_B = v, z_C = w$. For a partition $[k] = A \cup B \cup C$ denote

$$f(u, v, \star; \alpha, \star, \star) = \mathop{\mathbb{E}}_{w} \sum_{\beta,\gamma} f(u, v, w; \alpha, \beta, \gamma)$$

where the implicit dependence on A, B, C is in that $(u, \alpha) \in ([N]\times[M])^A$, $(v, \beta) \in ([N]\times[M])^B$ and $(w, \gamma) \in ([N]\times[M])^C$. We interpret this quantity as the probability of selecting a random w and having $f(u, v, w)$ output an answer in $[M]^k$ that is equal to α on the coordinates of A.

Denote

$$win(a, b) = \mathop{\mathbb{E}}_{A,B} \mathop{\mathbb{E}}_{u,v} \sum_{\alpha} f(u, v, \star; \alpha, \star\star)^2$$

where the expectation is over all disjoint sets A, B with $|A| = a$ and $|B| = b$.

Proposition 2.1. *The quantity $win(a, b)$ is the acceptance probability of $\mathcal{T}(a, b)$.*

Proof. The proof follows by definition: for any u, v, the quantity $\sum_\alpha f(u, v, \star; \alpha, \star, \star)^2$ is the probability of selecting w, w' at random and then selecting β, β' and then accepting if there is some α such that $\beta_A = \alpha = \beta'_A$. □

2.2 Rotation Invariance

The direct product test we described in the introduction has a well-studied "rotation-invariant" version, which is more relevant when testing that $f = g \times \cdots \times g$, i.e. that f is the k-fold direct product of some $g : [N] \to [M]$ with itself.

For a permutation $\sigma : [k] \to [k]$ and a string $x \in X^k$ let σx be the string $x_{\sigma(1)}, \ldots, x_{\sigma(k)}$.

> $\mathcal{T}_{rot}(t)$ **- rotation-invariant two-query test with intersection** t
>
> Choose $A \subset [k]$ of size t. Choose $x, y \in [N]^k$ uniformly at random, conditioned on $x_A = y_A$. Choose a random permutation $\sigma : [k] \to [k]$. Accept iff $f(x)_A = (\sigma^{-1} f(\sigma y))_A$.

For any function $f : [N]^k \times [M]^k \to \mathbb{R}_{\geqslant 0}$, the "symmetrized" version of f is defined as $f_{sym} : [N]^k \times [M]^k \to \mathbb{R}_{\geqslant 0}$ given by

$$f_{sym}(x;\alpha) = \mathop{\mathbb{E}}_{\sigma} f(\sigma x; \sigma\alpha)$$

It is not hard to see that f_{sym} is "rotation invariant" : for every σ and $x;\alpha$ $f_{sym}(x;\alpha) = f_{sym}(\sigma x; \sigma\alpha)$. Moreover, the success of f_{sym} in the original DP test is equal to the success of f in the rotation invariant test. Thus, we can deduce for example from Theorem 1.1 that f_{sym} is close to a direct product that is rotation-invariant, and this implies the closeness of f to the same direct product.

3 Local Structure

In this section we prove Lemma 1.2. So assume f is given such that $win(t, 0) \geqslant \delta$, and let $1 - \eta = \delta^{2/t}$.

3.1 A Potential Function Argument

We first find values a, b for which $win(a, b+1) \approx win(a, b) \approx win(a+1, b)$. We prove,

Lemma 3.1. *There are values $t/2 \leqslant a \leqslant t$ and $b \leqslant t/2$, such that $\delta \leqslant win(a, b)$ and such that*

$$(1-\eta) \leqslant \frac{win(a,b)}{win(a,b+1)} \leqslant 1 \qquad \text{and} \qquad (1-\eta) \leqslant \frac{win(a+1,b)}{win(a,b+1)} \leqslant 1$$

First, we show a generic monotonicity property of the value $win(a, b)$. This property is very useful, and is the main reason to consider $\mathcal{T}(a, b)$ instead of $\mathcal{T}(t)$, as there is no direct way to compare $\mathcal{T}(t)$ and $\mathcal{T}(t+1)$.

Proposition 3.2 (Monotonicity)**.** *Suppose $a + b < k$. Then*

$$win(a, b+1) \geqslant win(a+1, b)$$

and

$$win(a, b+1) \geqslant win(a, b).$$

Proof. For a tuple $\tau = (A, B, i, u, v, \alpha)$ where $|A| = a, |B| = b$ and $A \cap B = \phi$ and $u \in [N]^A, v \in [N]^B$ and $\alpha \in [M]^A$, we denote

$$\forall x \in [N], \beta \in [M] \qquad h_\tau(x, \beta) = f(ux, v, \star; \alpha\beta, \star, \star). \tag{3.1}$$

and define

$$T(\tau) = \mathop{\mathbb{E}}_x h_\tau(x, \star)^2, \quad I(\tau) = \mathop{\mathbb{E}}_x \sum_\beta h_\tau(x, \beta)^2, \quad K(\tau) = (\mathop{\mathbb{E}}_x h_\tau(x, \star))^2. \tag{3.2}$$

Let $h_\tau(x, \star) = \sum_\beta h_\tau(x, \beta)$. Now

$$win(a, b+1) = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha \mathop{\mathbb{E}}_x h_\tau(x, \star)^2 \qquad = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha T(\tau) \tag{3.3}$$

$$win(a+1, b) = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha \mathop{\mathbb{E}}_x \sum_\beta h_\tau(x, \beta)^2 \qquad = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha I(\tau) \tag{3.4}$$

$$win(a, b) = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha (\mathop{\mathbb{E}}_x h_\tau(x, \star))^2 \qquad = \mathop{\mathbb{E}}_{A,B,\{i\}} \mathop{\mathbb{E}}_{uv} \sum_\alpha K(\tau) \tag{3.5}$$

where the expectations are over all uniformly chosen disjoint sets A, B of size a, b uniformly chosen $i \notin A \cup B$, and $u \in [N]^A$, $v \in [N]^B$, $\alpha \in [M]^A$. We will show that for each $\tau = (A, B, i, u, v, \alpha)$

$$T(\tau) \geqslant I(\tau) \qquad \text{and} \qquad T(\tau) \geqslant K(\tau).$$

Indeed, for any non-negative function $h : [N] \times [M] \to \mathbb{R}_{\geqslant 0}$, (here we also denote $h(x, \star) = \sum_\beta h(x, \beta)$):

$$\mathop{\mathbb{E}}_x h(x, \star)^2 \geqslant \mathop{\mathbb{E}}_x \sum_\beta h(x, \beta)^2$$

which means that $T(\tau) \geqslant I(\tau)$ so $win(a, b+1) \geqslant win(a+1, b)$. Also, due to Cauchy-Schwarz,

$$\mathop{\mathbb{E}}_x h(x, \star)^2 \geqslant (\mathop{\mathbb{E}}_x h(x, \star))^2$$

which means that $T(\tau) \geqslant K(\tau)$ so $win(a, b+1) \geqslant win(a, b)$. $\square$

Using this monotonicity we can now easily prove Lemma 3.1,

Proof. The two upper bounds are immediate from Proposition 3.2 for all values of a, b. For the lower bounds, we run a hybrid argument. Fix (a_0, b_0) to equal $(t, 0)$. Inductively set $(a_{i+1}, b_{i+1}) = (a_i - 1, b_i + 1)$ or $(a_{i+1}, b_{i+1}) = (a_i, b_i + 1)$ whichever maximizes $win(a_{i+1}, b_{i+1})$. By Proposition 3.2 the sequence $win(a_i, b_i)$ is non-decreasing with i (and in particular, each $win(a_i, b_i)$ is at least δ). However, the multiplicative increase from step i to $i + 1$ cannot exceed

$$\frac{win(a_{i+1}, b_{i+1})}{win(a_i, b_i)} > (1 - \eta)^{-1}$$

for more than $t/2$ steps because $win(a_i, b_i) \leqslant 1$. □

3.2 Local Direct Product

With the values a, b in hand, we can complete the proof of Lemma 1.2. For a tuple $\tau = (A, B, i, u, v, \alpha)$ recall the definition of h_τ from (3.1) and define

$$\forall x \in [N] \qquad \beta_\tau(x) = \operatorname{argmax}_\beta h_\tau(x, \beta)\,. \tag{3.6}$$

Also recall the definition of $I(\tau), T(\tau), K(\tau)$ from (3.2)

$$K(\tau) = (\mathop{\mathbb{E}}_x h_\tau(x, \star))^2, \quad I(\tau) = \mathop{\mathbb{E}}_x \sum_\beta h_\tau(x, \beta)^2, \quad T(\tau) = \mathop{\mathbb{E}}_x h_\tau(x, \star)^2$$

and we add

$$D(\tau) = (\mathop{\mathbb{E}}_x h_\tau(x, \beta(x)))^2\,.$$

Define a distribution ν over tuples by $\nu(\tau) = \frac{T(\tau)}{\sum_\tau T(\tau)}$. Lemma 1.2 will follow directly from

Proposition 3.3. *If* $\mathbb{E}_{\tau\sim\nu} \frac{I(\tau)}{T(\tau)} \geqslant 1 - \eta$ *and* $\mathbb{E}_{\tau\sim\nu} \frac{K(\tau)}{T(\tau)} \geqslant 1 - \eta$ *then* $\mathbb{E}_{\tau\sim\nu} \frac{D(\tau)}{T(\tau)} \geqslant 1 - 8\eta$.

Proof of Lemma 1.2 from Proposition 3.3. Observe that

$$\mathop{\mathbb{E}}_{\tau\sim\nu} \frac{I(\tau)}{T(\tau)} = \sum_\tau \nu(\tau)\frac{I(\tau)}{T(\tau)} = \sum_\tau \frac{T(\tau)}{\sum_{\tau'} T(\tau')} \cdot \frac{I(\tau)}{T(\tau)} = \frac{\sum_\tau I(\tau)}{\sum_{\tau'} T(\tau')} = \frac{win(a+1, b)}{win(a, b+1)}$$

and similarly

$$\mathop{\mathbb{E}}_{\tau\sim\nu} \frac{K(\tau)}{T(\tau)} = \frac{\sum_\tau K(\tau)}{\sum_\tau T(\tau)} = \frac{win(a, b)}{win(a, b+1)}\,.$$

We already know from Lemma 3.1 that the two ratios above are at least $1 - \eta$. So the assumption of the proposition holds. As for the conclusion, the proposition asserts that

$$\mathop{\mathbb{E}}_{\tau\sim\nu} \frac{D(\tau)}{T(\tau)} = \frac{\sum_\tau D(\tau)}{\sum_\tau T(\tau)} \geqslant 1 - 8\eta$$

which implies

$$(1 - 8\eta)^{-1} \sum_\tau D(\tau) \geqslant \sum_\tau T(\tau) \geqslant \sum_\tau K(\tau)$$

where the second inequality follows from Proposition 3.2. Equivalently, by similar arguments we get $\mathbb{E}_{\tau\sim\mu} \frac{D(\tau)}{K(\tau)} \geqslant 1 - 8\eta$ for $\mu(\tau) = \frac{K(\tau)}{\sum_{\tau'} K(\tau')}$. One can check that $\frac{D(s,i)}{K(s,i)} = \mathbb{P}_{z|s}[f(z)_i = \beta_{s,i}(z_i)]$, and it remains to verify that the definition of μ here coincides with the definition in Lemma 1.2:

$$\frac{\mathbb{P}[\mathcal{T}(a, b) \text{ accepts on } s]}{\mathbb{P}[\mathcal{T}(a, b) \text{ accepts}]} = \frac{\mathbb{P}[ABiuv] \cdot K(ABiuv\alpha)}{\mathbb{E}_{ABiuv} K(ABiuv\alpha)} = \frac{K(\tau)}{\sum_{\tau'} K(\tau')}$$ □

Proof. (of Proposition 3.3) Let $1-\eta(\tau)=\frac{1}{2}\left(\frac{I(\tau)}{T(\tau)}+\frac{K(\tau)}{T(\tau)}\right)$, so that $\mathbb{E}_\tau\,\eta(\tau)\leqslant\eta$. Observe that $\frac{I(\tau)}{T(\tau)}\geqslant 1-2\eta(\tau)$ and $\frac{K(\tau)}{T(\tau)}\geqslant 1-2\eta(\tau)$. We now rely on the following lemma, plugging in $h=h_\tau$, to deduce $\frac{D(\tau)}{T(\tau)}\geqslant(1-2\eta(\tau))(1-\frac{2\eta(\tau)}{1-2\eta(\tau)})^2\geqslant 1-8\eta(\tau)$. Taking expectation over τ (according to ν) completes the proof. □

Lemma 3.4. *Let $h:[N]\times[M]\to\mathbb{R}_{\geqslant 0}$ and $\eta\geqslant 0$ be such that*

$$(1-\eta)\mathop{\mathbb{E}}_x h(x,\star)^2\leqslant\mathop{\mathbb{E}}_x\sum_\beta h(x,\beta)^2 \tag{3.7}$$

and

$$(1-\eta)\mathop{\mathbb{E}}_x h(x,\star)^2\leqslant(\mathop{\mathbb{E}}_x h(x,\star))^2 \tag{3.8}$$

Then, letting $\beta(x)=\operatorname{argmax}_\beta h(x,\beta)$,

$$(\mathop{\mathbb{E}}_x h(x,\beta(x)))^2\geqslant(1-\eta)(1-\frac{\eta}{1-\eta})^2\mathop{\mathbb{E}}_x h(x,\star)^2\,.$$

Proof. Define $p(x)=h(x,\beta(x))/h(x,\star)$. From (3.7),

$$\sum_x p(x)h(x,\star)^2=\sum_x\max_\beta h(x,\beta)\sum_\beta h(x,\beta)\geqslant\sum_{x,\beta}h(x,\beta)^2\geqslant(1-\eta)\sum_x h(x,\star)^2$$

and rearranging,

$$\sum_x h(x,\star)^2(1-p(x))\leqslant\eta\sum_x h(x,\star)^2 \tag{3.9}$$

So according to the distribution given by $\mathbb{P}(x)=\frac{h(x,\star)^2}{\sum_x h(x,\star)^2}$ the expectation of $1-p$ is small. Next, let $\bar{h}=\mathbb{E}_x h(x,\star)$ and write

$$\mathop{\mathbb{E}}_x(h(x,\star)-\bar{h})^2=\mathop{\mathbb{E}}_x[h(x,\star)^2-2\bar{h}h(x,\star)+\bar{h}^2]=\mathop{\mathbb{E}}_x h(x,\star)^2-\bar{h}^2\leqslant\eta\mathop{\mathbb{E}}_x h(x,\star)^2 \tag{3.10}$$

where the inequality at the end is from (3.8). Next,

$$\bar{h}\sum_x h(x,\star)(1-p(x))=\sum_x(\bar{h}-h(x,\star))h(x,\star)(1-p(x))+\sum_x h(x,\star)^2(1-p(x)) \tag{3.11}$$

The second term is immediately bounded by $\eta\sum_x\ h(x,\star)^2$ by (3.9). We bound the first term using Cauchy Schwarz to get

$$(\sum_x(\bar{h}-h(x,\star))^2)^{1/2}(\sum_x h(x,\star)^2(1-p(x))^2)^{1/2}\leqslant\eta\sum_x h(x,\star)^2$$

where we have used (3.10) and again (3.9) (and using $1-p(x)\leqslant 1$).

Plugging $\sum_x h(x,\star)^2\leqslant\frac{\bar{h}}{1-\eta}\sum_x h(x,\star)$ (from (3.8)) into (3.11) and dividing by $\bar{h}$ gives

$$\mathop{\mathbb{E}}_x[h(x,\star)-h(x,\beta(x))]\leqslant\frac{\eta}{1-\eta}\mathop{\mathbb{E}}_x h(x,\star)$$

Finally, by rearranging and squaring we get

$$(\mathop{\mathbb{E}}_x h(x,\beta(x)))^2\geqslant(1-\frac{\eta}{1-\eta})^2(\mathop{\mathbb{E}}_x h(x,\star))^2\geqslant(1-\eta)(1-\frac{\eta}{1-\eta})^2\mathop{\mathbb{E}}_x h(x,\star)^2\,.$$

where the second inequality is from (3.8). □

Finally, we note that if f is rotation-invariant then for every $s=ABuv\alpha$ it is easy to see that for every $i,i'\notin A\cup B$, $\beta_{s,i}=\beta_{s,i'}$ because for each x, $\beta_{s,i}(x)$ is the most popular β among all $z|s$ with $z_i=x$. But permuting each z so that i and i' are swapped (this does not change $z|s$) and implies that this equals $\beta_{s,i'}(x)$.

4 Global Direct Product

In this section we prove Theorem 1.1. We now assume that the function f passes the test with probability $\delta \geqslant 1-\varepsilon$ for some small $\varepsilon > 0$, i.e. $win(t,0) > 1-\varepsilon$. We first show in Lemma 4.1 that the high winning probability implies that for nearly all of $ABuv$ there is really only one α responsible for almost all of the winning probability, and so we get a version of Lemma 1.2 that does not involve μ. We then proceed to prove Theorem 1.1 from Lemma 4.1 by averaging arguments: we show that the local functions on each $ABuv$ must agree with each other essentially by arguing that a random element that sits in both $ABuv$ and $A'B'u'v'$ must agree with both functions. This gives rise to one global function.

Let $t/2 \leqslant a \leqslant t$ and $b \leqslant t/2$ be the values guaranteed by Lemma 1.2, and let $\eta = 1-\delta^{2/t}$ and observe that $\eta = \Theta(\varepsilon/t)$. Throughout this section we refer to $ABuv\alpha$ implicitly assuming that A, B are disjoint of sizes a, b respectively, and that $u \in [N]^A, v \in [N]^B, \alpha \in [M]^A$. We first prove

Lemma 4.1. *For each A, B, u, v there is a direct product function $\beta_{ABuv} : [N]^C \to [M]^C$ (where $C = [k] \setminus (A \cup B)$), such that*

$$\mathop{\mathbb{E}}_{A,B,u,v} \mathop{\mathbb{P}}_{D,w} [f(uvw)_D = \beta_{ABuv}(w_D)] \geqslant 1 - O(\varepsilon)$$

where $D \subset C$ is a random subset of size t.

The difference between this statement and the conclusion of Lemma 1.2 is in that here the average is over all uniform A, B, u, v and there is no dependence on μ, and in that the local direct product functions are independent of α.

Proof. From Lemma 1.2 we have

$$\mathop{\mathbb{E}}_{s=ABuv\alpha\sim\mu} \mathop{\mathbb{E}}_{i\notin A\cup B} \mathop{\mathbb{P}}_{z|s}[f(z)_i \neq \beta_\tau(z_i)] \leqslant O(\eta)\,.$$

For any $D \subset C$, $|D| = t$, we now do a simple union bound writing $\sum_{i\in D} \mathbb{P}_{z|s}[f(z)_i \neq \beta_{s,i}(z_i)] \geqslant \mathbb{P}_{z|s}[f(z)_i \neq \beta_{s,i}(z_i)$ for some $i \in D]$, and get

$$\sum_s \mu(s) \mathop{\mathbb{P}}_{D;\, z|s}[f(z)_D \neq \beta_s(z_D)] \leqslant O(t\eta) \leqslant O(\varepsilon)$$

where $\beta_s(z_D)$ is short for $(\beta_{s,i}(z_i))_{i\in D}$. Taking complements we get

$$\mathop{\mathbb{E}}_{s\sim\mu} \mathop{\mathbb{P}}_{D;\, z|s}[f(z)_D = \beta_s(z_D)] \geqslant 1 - O(\varepsilon)\,. \tag{4.1}$$

Plugging in $\mu(s) = \mathbb{P}[\mathcal{T}(a,b,c+1) \text{ accepts on } s] / \mathbb{P}[\mathcal{T}(a,b,c+1) \text{ accepts}]$ into this, and multiplying by $\mathbb{P}[\mathcal{T}(a,b,c+1) \text{ accepts}] \geqslant 1-\varepsilon$ we get (keeping $s = ABuv\alpha$ and letting 'win with α' be the event that $\mathcal{T}(a,b,c+1)$ accepts with $f(uvw)_A = \alpha = f(uvw')_A$)

$$\mathop{\mathbb{E}}_{ABuv} \sum_\alpha \mathbb{P}[\text{win with } \alpha \,|\, ABuv] \cdot \mathop{\mathbb{P}}_{D;\, z|s}[f(z)_D = \beta_s(z_D)] \geqslant 1 - O(\varepsilon)\,. \tag{4.2}$$

We now claim that $\mathbb{P}[\mathcal{T}(a,b,c+1) \text{ accepts}] \geqslant 1-\varepsilon$ implies

$$\sum_{ABuv} \sum_{\alpha\neq\alpha_{max}} \mathbb{P}[\mathcal{T}(a,b,c+1) \text{ accepts on } s = ABuv\alpha] \leqslant 2\varepsilon \tag{4.3}$$

where $\alpha_{max} = \mathrm{argmax}_\alpha \mathbb{P}[\text{win with } \alpha \,|\, ABuv]$. Indeed, write

$$\begin{aligned} 1 - \varepsilon_{ABuv} &= \sum_\alpha \mathbb{P}[\text{win with } \alpha \,|\, ABuv] = \sum_\alpha f(u,v,\star;\alpha,\star,\star)^2 \\ &\leqslant \sum_\alpha f(u,v,\star;\alpha,\star,\star) \cdot \max_\alpha f(u,v,\star;\alpha,\star,\star) \leqslant f(u,v,\star;\alpha_{max},\star,\star) \end{aligned}$$

which means that

$$1 - 2\varepsilon = \mathop{\mathbb{E}}_{ABuv}[1 - 2\varepsilon_{ABuv}] \leqslant \mathop{\mathbb{E}}_{ABuv}(1-\varepsilon_{ABuv})^2 \leqslant \mathop{\mathbb{E}}_{ABuv} \mathbb{P}[\text{win with } \alpha_{max} \,|\, ABuv]$$

implying (4.3). So we can neglect in (4.2) all α except for α_{max}, while only losing another at most 2ε, so, denoting $\beta_{ABuv} = \beta_{ABuv\alpha_{max}}$ we have

$$\begin{aligned} 1 - O(\varepsilon) &\leqslant \mathop{\mathbb{E}}_{ABuv} \mathbb{P}[\text{win with } \alpha_{max} \,|\, ABuv] \cdot \mathop{\mathbb{P}}_{D;\, z|ABuv\alpha_{max}}[f(z)_D = \beta_{ABuv}(z_D)] \\ &\leqslant \mathop{\mathbb{E}}_{ABuv} \mathop{\mathbb{P}}_{D;z|ABuv}[f(z)_A = \alpha_{max}] \cdot \mathop{\mathbb{P}}_{z|ABuv\alpha_{max}}[f(z)_D = \beta_{ABuv}(z_D)] \\ &= \mathop{\mathbb{E}}_{ABuv} \mathop{\mathbb{P}}_{D,w}[f(uvw)_A = \alpha_{max} \text{ and } f(uvw)_D = \beta_{ABuv}(w_D)] \\ &\leqslant \mathop{\mathbb{E}}_{ABuv} \mathop{\mathbb{P}}_{D,w}[f(uvw)_D = \beta_s(w_D)] \end{aligned}$$

The inequality in the first line follows from removing the non-maximal α in (4.2) and using the bound in (4.3). The inequality in the second line follows because $\mathbb{P}[\text{win with } \alpha \,|\, ABuv] = \mathbb{P}_{z|ABuv}[f(z)_A = \alpha]^2 \leqslant \mathbb{P}_{z|ABuv}[f(z)_A = \alpha]$. In the last two lines we rewrote $z|ABuv$ by choosing a random $w \in [N]^{[k]-(A\cup B)}$ and letting $z = uvw$. □

Now that the local structure is more uniformly distributed we can finally merge the piecewise local structure into a global structure. This is done by first using a union bound to show that for a random $ABuv$, the local direct product function β_{ABuv} agrees with f whp on *almost all* z, including those z that are not consistent with u, v. This already shows that β_{ABuv} has global agreement with f. Next, since β_{ABuv} is only defined on the coordinates $C = [k] - (A \cup B)$, we need to find another global function that takes care of the remaining coordinates. This is done essentially by repeating this argument once more with a disjoint pair of A', B'.

Proof of Theorem 1.1. Choose A, B at random disjoint of size a, b and let C be their complement. Choose $D \subset C$ at random such that $|D| = t$.

Choose u, v at random on A, B respectively, and complete them to $z = (u, v, w_1, w_2)$ where $w_2 \in [N]^D$ is defined on the coordinates in D. Choose u', v', w_1' independently from u, v, w_1 and set $z' = (u', v', w_1', w_2)$. The probability that $f(z)_D = f(z')_D$ is at least $1 - \varepsilon$ by the assumption $win(a, b, c+1) \geqslant 1 - \varepsilon$. We also know by the previous lemma, that the probability that $f(z)_D = \beta_{ABuv}(z_D)$ is at least $1 - O(\varepsilon)$. Together both events happen with probability $1 - O(\varepsilon) - \varepsilon$, and then we have

$$f(z')_D = f(z)_D = \beta_{uvAB}(z_D) = \beta_{uvAB}(z'_D)$$

where we have used the direct-product-ness of β for the last equality. We conclude that for a random choice of A, B, u, v and a uniformly random z (independent of A, B, u, v) and a random subset $D \subset [k] - (A \cup B)$, $\mathbb{P}[f(z)_D = \beta_{ABuv}(z_D)] \geqslant 1 - O(\varepsilon)$. A final union bound allows us to choose at random both A, B, u, v *and* A', B', u', v' such that $A' \cup B'$ is disjoint from $A \cup B$ and such that both

$$\mathop{\mathbb{P}}_{D\subset[k]-(A\cup B)}[f(z)_D = \beta_{ABuv}(z_D)] \geqslant 1 - O(\varepsilon) \quad \textit{and} \quad \mathop{\mathbb{P}}_{D\subset[k]-(A'\cup B')}[f(z)_D = \beta_{A'B'u'v'}(z_D)] \geqslant 1 - O(\varepsilon)\,.$$

where the probabilities are also over z and the choices of $ABuv, A'B'u'v'$. Thus, there must be a choice of $ABuv, A'B'u'v'$ such that conditioning on this choice doesn't lower the probability, and we can choose $\beta_i = \beta_{ABuv,i}$ for all $i \notin A \cup B$ as well as $\beta_i = \beta_{A'B'u'v',i}$ for $i \in A \cup B$. Finally,

$$\mathop{\mathbb{P}}_{z \in [N]^k, |D|=t} [f(z)_D = \beta(z_D)] \geqslant 1 - O(\varepsilon)$$

follows by splitting D into $D_1 = D \cap (A \cup B)$ and $D_2 = D \setminus (A \cup B)$. The above implies that with probability $1 - O(\varepsilon)$ both $f(z)_{D_1} = \beta(z_{D_1})$ and $f(z)_{D_2} = \beta(z_{D_2})$.

For each z, $\mathbb{P}_D[f(z)_D \neq \beta(z_D)] \geqslant \Omega(dist(f(z), \beta(z)) \cdot t/k$, and taking expectation over z and multiplying by k/t we get $\mathbb{E}_z[dist(f(z), \beta(z))] \leqslant O(\varepsilon t/k)$. For $t = k/2$ this implies that the probability of $f(x) \neq g(x)$ is $O(\varepsilon)$.

Finally, we note that if f is rotation-invariant, then by Lemma 1.2 each β_s is rotation invariant, and this carries over to each β_{uvAB}. Our final β is obtained from gluing two such together, but we can infer from their overlap that *all* coordinates of β are identical. We omit the details. □

We remark that this theorem is tight (up to the constants in the $O(\cdot)$ notation). Clearly, the conclusion must allow for ε fraction of total junk in f. Moreover, consider a function f obtained by taking $g_1 \times \cdots \times g_k$ and for each x changing $r = \varepsilon k/t$ coordinates of $f(x)$ independently at random. This function still passes $\mathcal{T}(t)$ with probability at least $(1 - t/k)^{2r} = 1 - O(\varepsilon)$.

Acknowledgements

We are thankful to Igor Shinkar and Elazar Goldenberg for helpful comments.

References

[DDG+13] Roee David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. Direct sum testing. Submitted, 2013. 4

[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. 2, 3, 4

[Din07] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3), 2007. 2

[DM11] Irit Dinur and Or Meir. Derandomized parallel repetition via structured PCPs. *Computational Complexity*, 20(2):207–327, 2011. 2

[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. 2, 4

[DS13] Irit Dinur and David Steurer. Analytical approach to parallel repetition. *CoRR*, abs/1305.1979, 2013. 3, 13

[FK94] U. Feige and J. Kilian. Two prover protocols–low error at affordable rates. In *Proc. 26th ACM Symp. on Theory of Computing*, pages 172–183, 1994. 3, 4

[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. 2, 4

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query PCPs. *SIAM J. Comput.*, 41(6):1722–1768, 2012. 2, 3, 4

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998. 2

A From tuples to sets

This section is joint work with Igor Shinkar.

The direct product construction takes a function $g : [N] \to [M]$ into a function $f : [N]^k \to [M]^k$, by assigning to each k-tuple in $[N]^k$ a value in $[M]^k$ $f(z) = g(z_1), \dots, g(z_k)$. A very similar construction is to assign to each k-element subset of N a value in $[M]^k$. In other words, the set-direct product of a function $g : [N] \to [M]$ is the function $f : \binom{[N]}{k} \to [M]^k$ defined by $f(S) = g|_S$.

In this setting the same two-query test can be analysed: choose two subsets $Z, Z' \subset [N]$ of size k, such that $|Z \cap Z'| = t$, and check that $f(Z)_{Z\cap Z'} = f(Z')_{Z\cap Z'}$.

Intuitively, any DP testing theorem that holds for the tuple-direct-product should also hold for the set-direct-product.

The case $k \ll \sqrt{N}$. In this case a random k-tuple has whp k distinct elements, so the theorems proven about the tuple version can be used for analyzing the set version. One simply defines a function $f' : [N]^k \to [M]^k$ from f by setting $f'(z) = f(set(z))$ where $set(z) = \{z_1, \dots, z_k\}$. We can ignore the tuples that have repeated coordinates (for which $|set(z)| < k$) because they occur with negligible probability when $k \ll \sqrt{N}$. So the success probability of the newly defined f' is almost identical to the success probability of f. Similarly, the conclusion for f' carries over to f. We remark that the fact that f' came from a function on sets means that it is rotation invariant and this implies that the direct product components are identical.

The case $k \leqslant N/2$. For larger values of k, we can no longer assume that for a random $z \in [N]^k$, $|set(z)| = k$ with high probability. And in fact $|set(z)|$ is not concentrated on one single value. Nevertheless, it is concentrated enough on values "around" k to enable the proofs to go through. Let $f : \binom{[N]}{k} \to [M]^k$ and let $win_{set}(t)$ denote the success probability of the test with intersection parameter t.

Claim A.1. Let $t = k/2 \leqslant N/4$. If $win_{set}(t) > 1 - \varepsilon$ then $win_{set}(s) > 1 - 2\varepsilon$ for every $s < t$.

Proof. We express the distribution of $\mathcal{T}_{set}(s)$ as two steps in the distribution $\mathcal{T}_{set}(t)$. Choose a random pair of k-subsets X, Y that intersect on exactly s elements. Choose a third subset $Z = (X \cap Y) \cup X_1 \cup Y_1 \cup R$ where $X_1 \subset X \setminus Y$ is a random subset of size $t - s$ and symmetrically $Y_1 \subset Y \setminus X$ is a random subset of size $t - s$, and R is a random subset of $[N] \setminus (X \cup Y)$ of size $k - 2t + s$ (more generally, need $3k - 2t \leqslant N$ for this to work). It is easy to check that the distribution of X, Z is as in the set-DP-test(t) and similarly the distribution of Y, Z. Therefore, with probability at least $1 - 2\varepsilon$ both pairs agree on their intersection in which case so does the pair X, Y. $\square$

From this we can prove the analogous Theorem 1.1 for sets following these steps.

1. Given f, define a (randomized) function $f' : [N]^k \to [M]^k$ by setting $f'(z)$ as follows. Choose a random k-element set $Z \supseteq set(z)$, let $\alpha = f(Z)$ be viewed as a function from the k elements of Z to $[M]$. Set $f'(z)$ to be $(\alpha(z_i))_{i=1}^k$.

2. Show that if f passes the test with probability $1 - \varepsilon$ then f' passes the test with probability $1 - 2\varepsilon$. (Opening up the random choices of the test on f' one gets a convex combination of tests on f with intersection parameter s for various s).

3. Deduce that f' is close to a global direct product and therefore so is f.

B Relation to Parallel Repetition of Games

We describe in this section a specific parallel repeated game is equivalent to the direct product test, in that the *squared collision value* (defined below) of the game with strategy f is exactly equal to the success probability of the test $\mathcal{T}(t)$ (or $\mathcal{T}(a, b)$) on f.

We assume familiarity with the setting of [DS13]. Briefly, a one round two player projection game G is given by a bipartite graph with constraints on the edges, such that for each edge (u, v) Bob's answer to v determines at most one possible answer for Alice to u. A game is viewed as a linear map from a strategy $f : V \times \Sigma \to \mathbb{R}_{\geqslant 0}$ for Bob to a strategy $Gf : U \times \Sigma \to \mathbb{R}_{\geqslant 0}$ for Alice,

$$Gf(u, \alpha) = \mathop{\mathbb{E}}_{v|u} \sum_{\beta : \alpha \hookleftarrow \beta} f(v, \beta)$$

The direct product of two games G and H is the game $G \otimes H$ described by the tensor of the operators.

The *squared collision value* of a projection game G is described by choosing a random Alice question u, and then two independent neighbors of it v, v', and then accepting if the answers of Bob on v and v' project to the same answer for u.

All of the games below will have right vertex set $[N]$, and right alphabet $[M]$. This means that they all map, as linear operators a function in $\mathbb{R}^{[N]\times[M]}$ to some other space.

I: The identity game described by the graph $([N], [N], E = \cup_{x\in[N]}(x, x))$, alphabet $[M]$, and whose edges carry equality constraints. The corresponding operator is the identity $I : \mathbb{R}^{[N]\times[M]} \to \mathbb{R}^{[N]\times[M]}$.

T: The sum-over-answers game described by the graph $([N], [N], E = \cup_{x\in[N]}(x, x))$, right alphabet $[M]$, left alphabet $\{0\}$, and whose edges carry a free constraint: $\pi_e(a) = 0$ for every a and e. The corresponding operator is defined by

$$Tf(v, 0) = \sum_b f(v, b)\,.$$

K: The averaging game described by the graph $(U = \{0\}, [N], E = \cup_{x\in[N]}(0, x))$, right alphabet $[M]$, left alphabet $\{0\}$, and whose edges carry a free constraint: $\pi_e(a) =$ for every a and e. The corresponding operator is defined by

$$Kf(0, 0) = \mathop{\mathbb{E}}_v \sum_b f(v, b)\,.$$

The direct product test $\mathcal{T}(t)$ is given by first choosing a random subset $A \subset [k]$ of size t and then letting $G = I^{\otimes A} \otimes K^{\otimes [k]-A}$ be the game. Thus, we have

$$win(t) = \mathop{\mathbb{E}}_{|A|=t} \|(I^{\otimes A} \otimes K^{\otimes [k]-A}) f\|^2 .$$

and similarly

$$win(a,b) = \mathop{\mathbb{E}}_{|A|=a,|B|=b,A\cap B=\phi} \|(I^{\otimes A} \otimes T^{\otimes B} \otimes K^{\otimes [k]-A-B}) f\|^2 .$$

Our analysis can be seen as taking the original game G that has t copies of I and $k-t$ copies of K and replacing at each step one more coordinate with a copy of T. At each step the value can only go up, but there is a point where it doesn't change much. This is exactly captured in the proof of Lemma 3.1 with quantities $K(\tau)$, $T(\tau)$ and $I(\tau)$ (see (3.2)) that represent the change on some coordinate i between games K, T or I.