A Near-Optimal Polynomial Distance Lemma over Boolean Slices

Prashanth Amireddy
Harvard University, Cambridge, MA, USA

Amik Raj Behera
University of Copenhagen, Denmark

Srikanth Srinivasan
University of Copenhagen, Denmark

Madhu Sudan
Harvard University, Cambridge, MA, USA

Abstract

The celebrated Ore-DeMillo-Lipton-Schwartz-Zippel (ODLSZ) lemma asserts that n-variate non-zero polynomial functions of degree d over a field $\mathbb{F}$, are non-zero over any "grid" (points of the form S^n for finite subset $S \subseteq \mathbb{F}$) with probability at least $\max\{|S|^{-d/(|S|-1)}, 1 - d/|S|\}$ over the choice of random point from the grid. In particular, over the Boolean cube ($S = \{0, 1\} \subseteq \mathbb{F}$), the lemma asserts non-zero polynomials are non-zero with probability at least 2^{-d}. In this work we extend the ODLSZ lemma optimally (up to lower-order terms) to "Boolean slices" i.e., points of Hamming weight exactly k. We show that non-zero polynomials on the slice are non-zero with probability $(t/n)^d(1 - o_n(1))$ where $t = \min\{k, n - k\}$ for every $d \leq k \leq (n - d)$. As with the ODLSZ lemma, our results extend to polynomials over Abelian groups. This bound is tight upto the error term as evidenced by multilinear monomials of degree d, and it is also the case that some corrective term is necessary. A particularly interesting case is the "balanced slice" ($k = n/2$) where our lemma asserts that non-zero polynomials are non-zero with roughly the same probability on the slice as on the whole cube.

The behaviour of low-degree polynomials over Boolean slices has received much attention in recent years. However, the problem of proving a tight version of the ODLSZ lemma does not seem to have been considered before, except for a recent work of Amireddy, Behera, Paraashar, Srinivasan and Sudan (SODA 2025), who established a sub-optimal bound of approximately $((k/n) \cdot (1 - (k/n)))^d$ using a proof similar to that of the standard ODLSZ lemma.

While the statement of our result mimics that of the ODLSZ lemma, our proof is significantly more intricate and involves spectral reasoning which is employed to show that a natural way of embedding a copy of the Boolean cube inside a balanced Boolean slice is a good sampler.

2012 ACM Subject Classification Theory of computation → Error-correcting codes

Keywords and phrases Low-degree polynomials, Boolean slices, Schwartz-Zippel Lemma

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.11

Category Track A: Algorithms, Complexity and Games

Related Version *Classification (Full Version)*: `Link to the full version`

Funding *Prashanth Amireddy*: Supported in part by Madhu Sudan's Simons Investigator Award and NSF Award CCF 2152413 and Salil Vadhan's Simons Investigator Award.
Amik Raj Behera: Supported by Srikanth Srinivasan's start-up grant from the University of Copenhagen.
Srikanth Srinivasan: Supported by European Research Council (ERC) under grant agreement no. 101125652 (ALBA).
Madhu Sudan: Supported in part by a Simons Investigator Award, NSF Award CCF 2152413 and AFOSR award FA9550-25-1-0112.

1 Introduction

The Ore-DeMillo-Lipton-Schwartz-Zippel (ODLSZ) [28, 12, 37, 33] lemma captures the basic algebraic fact that a low-degree polynomial does not have many roots on a "nice set" of points. The standard nice set for this lemma is a grid S^n (where S is a finite subset of a field) and a version of this lemma states that no non-zero degree-d polynomial can vanish on more that $d|S|^{n-1}$ points. This is easily seen to be tight: Take, for example, a univariate polynomial that has d roots in S.

There also exist useful variants of this lemma for the case where $|S| < d$. The example above shows that in general a degree-d polynomial can vanish over all of S^n and so some further condition is necessary. The most obvious condition is to simply force the polynomial to be non-zero on the grid S^n. In the setting of the Boolean cube, i.e. $S = \{0, 1\}$, which is the setting we study, this is equivalent to considering non-zero *multilinear* polynomials of degree d. In this setting (a variant of) the ODLSZ lemma states that a non-zero multilinear polynomial of degree d is non-zero on at least 2^{n-d} points of $\{0, 1\}^n$. Again, this is tight: Take, e.g., a multilinear monomial of degree d.

Though both these forms of the ODLSZ lemma are simple statements with easy inductive proofs, they have many different applications in the design of randomized algorithms [30], probabilistically checkable proofs [6, 5], pseudorandom constructions [15, 19], Boolean function analysis [26], data communication [2], small-depth circuit lower bounds [29, 23] and extremal combinatorics [31].

In this paper, we extend the ODLSZ lemma to a different nice set namely the *Boolean slice*, which is an important subset of the Boolean cube $\{0, 1\}^n$. For a parameter k, we use $\{0, 1\}^n_k$ to denote the kth Boolean slice, i.e., the set of points in the cube of Hamming weight exactly k. The behavior of low-degree polynomials on Boolean slices has received quite a bit of attention recently with motivations from learning theory [27], Boolean function analysis [36, 16, 18, 17], property testing [11, 24], circuit lower bounds [23], and local decoding algorithms [4]. However, as far as we know, the natural question of finding a tight version of the ODLSZ lemma over Boolean slices has not been considered before. This is the question we address in this paper.

More precisely, we consider the following question:

Given a polynomial P of degree at most d that does not vanish on $\{0, 1\}^n_k$, how many zeroes can have P have in this set?

This question makes sense when $d \leq t := \min\{k, n - k\}$, since any function on $\{0, 1\}^n_k$ can be expressed as a polynomial of degree t.

We give a near-optimal answer to this question for low-degree polynomials. More precisely, our main theorem is stated below. It holds for polynomials over any field and even in the case where the coefficients come from an *Abelian group*[1] (as is also true of the standard version of ODLSZ lemma over the Boolean cube).

▶ **Theorem 1** (Main Theorem). *There exists an absolute constant $\varepsilon > 0$ so that the following holds. Fix an arbitrary Abelian group G and a degree parameter $d \in \mathbb{N}$. For all natural numbers n and k such that $d \leq k \leq n - d$, the following holds whenever $d \leq t^\varepsilon$ where*

[1] A multilinear polynomial over an Abelian group G is of the form $\sum_{S \subseteq [n]} a_S \prod_{i \in S} x_i$ where $a_S \in G$ for each S. Polynomials over such domains appear naturally in applications to circuit complexity [7] and additive combinatorics [34].

$t = \min\{k, n-k\}$*:*
For any degree-d polynomial $P : \{0,1\}_k^n \to G$ *that does not vanish on* $\{0,1\}_k^n$*, we have*

$$\Pr_{\mathbf{x}\sim\{0,1\}_k^n}[P(\mathbf{x}) \neq 0] \;\geq\; \left(\frac{t}{n}\right)^d \cdot \left(1 - \frac{1}{t^\varepsilon}\right).$$

At a high level, the uniform distribution on $\{0,1\}_k^n$ is similar to the (k/n)-biased distribution, i.e., the distribution where each coordinate is independently 1 with probability k/n. With slight modifications to the proof of the ODLSZ lemma, one can show (see, for example, [14, Claim 6.8]) that the probability of sampling a nonzero point from (k/n)-biased distribution is $(t/n)^d$, where $t = \min\{k, n-k\}$. The bound given by Theorem 1 is equal to this bound up to small error terms.

Please refer to the full version of this paper for all the proofs.

Tightness

The bound given is easily seen to be nearly tight using essentially the same example as in the case of the Boolean cube. For $k \leq n/2$, the monomial $x_1 \cdots x_d$ is non-zero with probability approximately $(k/n)^d$, and there is a similar example for $k > n/2$. Moreover, it is also possible to see that for certain k, an error term is required. For example, assume that G is the finite field $\mathbb{F}_2$, $k = n/2$ and $d = 1$. Then the linear polynomial $x_1 + x_2 + 1$ is non-zero with probability $1/2 - \Theta(1/n)$, implying that the monomial does not yield exactly the optimal bound. In the case that the degree $d = 1$, we can improve the error parameter and show a bound of $t/n - 1/n$.

Proof Techniques

The standard proofs of the ODLSZ lemma follow a simple inductive strategy, using the obvious univariate case for both the base case and each inductive step of the argument. The recent work of Amireddy, Behera, Paraashar, Srinivasan and Sudan [4] used a similar idea to show the following sub-optimal bound. Unfortunately, it is not clear how to make the inductive strategy work for the slice to get a tight answer.

▶ **Lemma 2** (Suboptimal distance lemma for slices)**.** *[4, Lemma 5.1.6]. For every Abelian group* G *and non-negative integers* d, k, n *with* $n \geq 1$ *and* $d \leq k \leq n - d$ *the following holds: For every degree-d polynomial* $P : \{0,1\}_k^n \to G$ *that does not vanish on* $\{0,1\}_k^n$*, we have*

$$\Pr_{\mathbf{x}\sim\{0,1\}_k^n}[P(\mathbf{x}) \neq 0] \geq \binom{n-2d}{k-d} \Big/ \binom{n}{k}.$$

In particular, for $k = n/2$ *(for an even* n*), the above probability is at least* 4^{-d}*.*

A computation shows that for small d, the above implies that the fraction of points in $\{0,1\}_k^n$ where P does not vanish is at least $((k/n) \cdot (1 - (k/n)))^d$ (up to small error terms). When $k = n/2$, for example, this bound is 4^{-d}, which is quadratically worse than Theorem 1.

To get the tight bound, we use a very different approach. We start with the above suboptimal bound, but combine it with spectral techniques, which we elaborate on next. Note that if the slice $k = n/2$, i.e. the balanced slice, then we get a bound of nearly $1/2^d$, which is essentially the same as the ODLSZ lemma over the Boolean cube $\{0,1\}^n$ (Theorem 4).

To prove this, the high-level idea is to consider the process of choosing a random subcube in the balanced Boolean slice $\{0,1\}^n_{n/2}$ as follows: pair the n coordinates into $n/2$ pairs uniformly at random, and in each such pair $\{x_i, x_j\}$, identify x_i with the Boolean negation of x_j, i.e. $1 - x_j$. This gives us a random embedding of an $n/2$-dimensional cube in the slice $\{0,1\}^n_{n/2}$ and the polynomial P restricts to a degree-d polynomial Q on this subcube. If we could guarantee that Q was always non-zero, then the standard ODLSZ lemma on the cube would give us the desired statement. Unfortunately, there are subcubes on which P could be identically zero. The main technical lemma is to show that Q is non-zero with high probability: intuitively, this is because the random process above is a good *sampler* of the balanced slice, i.e. the points in the randomly chosen subcube behave essentially like independent samples of the balanced slice.

Formally, the technical lemma is a statement about the approximate pairwise independence of two random points of the chosen subcube. We show (see Lemma 9) that the probability that two random points of this subcube lie in a set of density ρ is roughly ρ^2. This is done by analyzing a natural weighted graph Γ on the balanced slice defined by the above sampling process. We show this via two arguments, depending on the regime of the degree parameter d.

For $d \leq C \log n$ for a constant $C > 0$, the main technical lemma follows from the use of the Expander mixing lemma [1], which implies such a statement using bounds on the second-largest eigenvalue of the graph. To analyze the second-largest eigenvalue of Γ, we show that it can be embedded (as an induced subgraph) in a Cayley graph defined on the subgroup of $\mathbb{F}_2^n$ defined by points of even Hamming weight. The latter is easier to analyze using Boolean Fourier analysis, and an application of the eigenvalue interlacing theorem allows us to bound the eigenvalues of Γ. See Section 3.1 for more details. This easier case of the lemma is already interesting: for instance, it yields a different (arguably easier) proof of a junta theorem on the Boolean slice [17], analogous to a well-known theorem of Nisan and Szegedy [26].

For $d = n^\gamma$ for a small constant $\gamma > 0$, we need to strengthen the guarantee of the sampler. To do so, we use the fact that the adjacency matrix for Γ can be spectrally upper-bounded by another matrix[2] that satisfies a *Hypercontractive inequality*. Intuitively, this is stronger than an eigenvalue bound, as the latter measures only the worst-case expansion of the underlying graph, while the former gives us stronger bounds on the expansion of smaller sets. Using this inequality alongside the Expander mixing lemma yields the desired pairwise independence. See the full version for more details.

For imbalanced slices, i.e., $k \neq n/2$, we reduce to the balanced case via a random restriction idea. The main conceptual idea is to obtain a basis for the space of polynomial functions on a slice. We note, essentially using an argument of Wilson [35], that for many distinct slices, the space of homogeneous multilinear monomials of degree d forms a basis for the space of polynomials of degree d on the slice. Unlike other known bases for this space [16], this idea also works over fields of positive characteristic and even over cyclic groups of prime power order. For such "good" slices $k \leq n/2$, we reduce to a $2k$-dimensional cube via a random restriction, which can easily be seen to leave the polynomial non-zero with probability $(2k/n)^d$. Invoking the balanced case now concludes the lemma for the good slices.

Finally, to extend the main theorem to all slices, we note that for any slice k, there is a good slice not too "far away" (in the range $[k - \mathcal{O}(d), k]$). By setting a few variables at random to 1, we are able to reduce to a good slice.

Please refer to the full version for complete details and all the proofs.

[2] The technically accurate descriptor for this matrix is the 'Noise operator in the Bose-Mesner algebra of the Johnson scheme.' See Section 3 for details.

Related Work

As mentioned above, the study of low-degree polynomials over Boolean slices has received much attention in recent years. Closely related to this work is the work of Filmus [16] that constructs a basis for the space of real-valued degree-d polynomial functions over general Boolean slices. A recent result of Kalai, Lifshitz, Minzer and Ziegler [24] constructs a *dense model* for the balanced slice $\{0,1\}^n_{n/2}$ under the Gowers norm U_d; in particular, this implies that there is a subset S of $\{0,1\}^n$ of constant density such that any polynomial of degree-d has the same density over S as it does over the balanced slice. In principle, both these works should be useful in order to prove a version of the ODLSZ lemma over Boolean slices. However, we note that each of these results is applicable over different domains ($\mathbb{R}$ or $\mathbb{F}_2$) while we prove a unified statement that holds over any Abelian group (and in particular over all fields).

1.1 Applications of Optimal Distance Lemma

To give some idea of the applicability of the ODLSZ lemma over the slice, we prove some variants of well-known theorems in combinatorics and Boolean function analysis.

Hyperplane covering

Given a subset S of the cube $\{0,1\}^n$, we define the *exact cover number of* S, denoted $\mathrm{ec}_n(S)$ to be the minimum number of hyperplanes (over some field $\mathbb{F}$) such that their union intersects $\{0,1\}^n$ exactly in the set S. A classical result of Alon and Füredi shows that for S being the cube with a single point removed, $\mathrm{ec}_n(S) = n$. This combinatorial result, which easily follows from with ODLSZ lemma over the cube, has seen many subsequent generalizations (e.g. [10, 32, 8]).

Using just the sub-optimal version of the ODLSZ lemma (Lemma 2), we immediately get an optimal version of the hyperplane covering over a Boolean slice $\{0,1\}^n_k$ with a missing point, instead of the whole Boolean cube $\{0,1\}^n$. More precisely, for $S \subseteq \{0,1\}^n_k$, let $\mathrm{ec}_{n,k}(S)$ be the minimum number of hyperplanes (over some fixed field $\mathbb{F}$) such that their union intersects $\{0,1\}^n_k$ exactly in the set S. Following the idea of [3], we have the following.

▶ **Theorem 3.** *Let n, k be natural numbers with $k \in [n]$. Fix an arbitrary point $\mathbf{a} \in \{0,1\}^n_k$. Then $\mathrm{ec}_{n,k}(\{0,1\}^n_k \setminus \{\mathbf{a}\}) = \min\{k, n-k\}$.*

Proof of Theorem 3. Without loss of generality, we assume that $k \le n/2$ and $\mathbf{a} = 1^k 0^{n-k}$. Let S denote $\{0,1\}^n_k \setminus \{\mathbf{a}\}$.

It is easy to see that $\mathrm{ec}_{n,k}(S) \le k$. The hyperplanes $H_i = \{x_i = 0\}$ for $i \in [k]$ cover exactly the points in S.

For the lower bound, assume for the sake of contradiction that there exists $m < k$ hyperplanes $H_i = \{\ell_i(\mathbf{x}) = 0\}$ (here $\ell_i(\mathbf{x})$ denotes a degree-1 polynomial and $i \in [m]$) covering exactly the points in S. Then the polynomial $P(\mathbf{x}) := \prod_{i=1}^m \ell_i(\mathbf{x})$ is non-zero at exactly one point of $\{0,1\}^n_k$.

However, by Lemma 2, P must be non-zero at at least $\binom{n-2m}{k-m} > 1$ points (since $m < k \le n/2$) of $\{0,1\}^n_k$. Hence we arrive at a contradiction. ◀

A junta theorem for the slice

Nisan and Szegedy [26] showed that any Boolean function on $\{0,1\}^n$ that has degree d over $\mathbb{R}$ depends on $\mathcal{O}(d2^d)$ variables, i.e. it is a $\mathcal{O}(d2^d)$-junta. Chiarelli, Hatami, and Saks [9] improved the bound to $\mathcal{O}(2^d)$. Filmus and Ihringer [17] extended this result to slices and

showed that for a suitable range of k, any degree-d (over $\mathbb{R}$) Boolean function on the slice k is a restriction of a degree-d function on $\{0,1\}^n$. Along with the result of [9], this implies that such a function is an $\mathcal{O}(2^d)$-junta. While the results of [26, 9] are fairly elementary, the theorem of [17] is more involved, relying on the Log-Sobolev inequality and Hypercontractivity for the Boolean slice [25, 13].

Using Theorem 1, we show that we can avoid the use of advanced analytic techniques[3] in the proof of [17], and give a direct proof (following the proof of [26]) of the fact that any degree-d Boolean function on the balanced slice $\{0,1\}^n_{n/2}$ depends on $\mathcal{O}(d2^d)$ variables (see Lemma 19). Plugging this into the proof of [17], we can again recover the optimal bound of $\mathcal{O}(2^d)$. More details can be found in Section 4.

2 Preliminaries

Notations

Let $(G,+)$ denote an Abelian group G with addition as the binary operation. For any $g \in G$, let $-g$ denote the inverse of $g \in G$. For any $g \in G$ and integer $a \geq 0$, $a \cdot g$ (or simply ag) is the shorthand notation of $g + \ldots + g$ (taken a times), and $-ag$ denotes $a \cdot (-g)$.

For any $\mathbf{x} \in \{0,1\}^n$, $|\mathbf{x}|$ denotes the Hamming weight of $\mathbf{x}$. For any $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, let $\Delta(\mathbf{x},\mathbf{y})$ denote the Hamming distance between $\mathbf{x}$ and $\mathbf{y}$, i.e. $\Delta(\mathbf{x},\mathbf{y}) = |\{i \in [n] \mid x_i \neq y_i\}|$. For natural numbers n and $k \leq n$, let $\{0,1\}^n_k$ denote the subset of strings in $\{0,1\}^n$ of Hamming weight exactly k.

We denote the set of functions $f : \{0,1\}^n \to G$ that can be expressed as a multilinear polynomial of degree d, with the coefficients being in G by $\mathcal{P}_d(n,G)$. We also consider functions $f : \{0,1\}^n_k \to G$. We denote the set of functions on $\{0,1\}^n_k$ that can be expressed as a multilinear polynomial of degree d with the coefficients in G by $\mathcal{P}_d(n,k,G)$. We will simply write $\mathcal{P}_d(n,k)$ when G is clear from the context.

For any natural numbers n and $k \leq n$, U_n denotes the uniform distribution on $\{0,1\}^n$ and $U_{n,k}$ denotes the uniform distribution on $\{0,1\}^n_k$. For a growing parameter n, $o_n(1)$ denotes a function that goes to 0 as n grows large.

Basic Tools

We start with the standard ODLSZ lemma over the Boolean cube.

▶ **Theorem 4** (ODLSZ lemma over $\{0,1\}^n$)**.** *Let G be any Abelian group and let $P \in \mathcal{P}_d(n,G)$ be any non-zero polynomial. Then*

$$\Pr_{\mathbf{x} \sim U_n}[P(\mathbf{x}) \neq 0] \geq \frac{1}{2^d}.$$

We will need the following standard facts about expanders and Cayley graphs. We refer the reader to the survey [21] for more details.

[3] We have two proofs of our main theorem. In the general case where d can be as large as $n^{\Omega(1)}$, our proof also relies on hypercontractivity. However, in the case that $d \leq C \log n$, which is also the main case of interest for junta theorems, our proof needs only basic Fourier analysis over the Boolean cube and the eigenvalue interlacing theorem.

▶ **Definition 5** (Weighted Cayley Graph). *Let $(G,+)$ be a finite Abelian group and $w : G \to \mathbb{R}^{\geq 0}$ be a* weight *function (we refer to the elements of non-zero weight as* generators*). We say that a weighted graph $\Gamma = \Gamma(G,w)$ defined as follows is a* weighted Cayley graph *over G.*

- *The vertices of Γ are the elements of G.*
- *For every $g, g' \in G$, we add an edge $(g, g+g')$ with weight $w(g')$ to Γ.*

The following lemma gives us a way of computing the eigenvalues of the adjacency matrix of weighted Cayley graphs over Abelian groups.

▶ **Lemma 6** (Eigenvalues of Cayley graphs, see e.g. [21]). *Let $\Gamma = \Gamma(G,w)$ be a weighted Cayley graph over a finite Abelian group G, where $w : G \to \mathbb{R}^{\geq 0}$ is the corresponding weight function. Let $\chi : G \to \mathbb{C}^{\times}$ be an arbitrary group homomorphism (which we will refer to as a* character*). Then, χ is an eigenvector of the adjacency matrix of Γ with eigenvalue equal to $\sum_{g \in G} w(g)\chi(G)$.*

Following is a consequence of the *expander mixing lemma*.

▶ **Lemma 7** (Expander mixing lemma, see e.g. [21] Lemma 2.5). *For a symmetric random walk matrix W over vertices V and every subset $S \subseteq V$, it holds that*

$$\Pr_{u \sim V, v \sim N(u)}[u \in S \text{ and } v \in S] \leq \left(\frac{|S|}{|V|}\right)^2 + \mu(W) \cdot \frac{|S|}{|V|},$$

where $N(u)$ denotes the distribution over V corresponding to taking a step from u according to W (i.e., the u-th row of W).

3 Distance Lemma for the Balanced Slice

In this section, we state the main technical lemma of our proof for Theorem 1. It is a statement on the "expansion" property of a graph $\{0,1\}^n_{n/2}$, where the edge weights are given by a random process. We start by describing a random process that maps a string in $\{0,1\}^n_{n/2}$ to a string in $\{0,1\}^n_{n/2}$. In this section, we will always assume that n is an even number.

▶ **Definition 8** (The map Γ). *Let $\mathbf{a} \in \{0,1\}^{n/2}$ and $\mathbf{u} \in \{0,1\}^n_{n/2}$. Let $\mathbf{u}^{-1}\{0\}$ denote the set of coordinates where $\mathbf{u}$ is 0, i.e. $\mathbf{u}^{-1}\{0\} = \{i \in [n] \mid u_i = 0\}$. Similarly we have $\mathbf{u}^{-1}\{1\}$. let $\mathbf{u}^{-1}\{1\} = \{i_1, \ldots, i_{n/2}\}$.*
For any perfect matching $\mathcal{M}$ between $\mathbf{u}^{-1}\{0\}$ and $\mathbf{u}^{-1}\{1\}$ ($\mathcal{M}$ is a bijection between these two sets), the function $\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a}))$ is a balanced string $\mathbf{v} \in \{0,1\}^n_{n/2}$ defined as follows: For every $k \in [n/2]$, $v_{i_k} = u_{i_k} \oplus a_k$ and $v_{\mathcal{M}(i_k)} = u_{\mathcal{M}(i_k)} \oplus a_k$.

In simple words, for every matching between the 0-coordinates and 1-coordinates of $\mathbf{u}$ and a string $\mathbf{a} \in \{0,1\}^{n/2}$, we get a new balanced string $\mathbf{v}$ by flipping the endpoints of a subset of matching edges. Here the subset of matching edges whose endpoints are flipped is given by the string $\mathbf{a}$. Following is an example for $n = 8$.

▶ **Example.** *Let $\mathbf{u} = 10101010$. Here $\mathbf{u}^{-1}\{0\} = \{2,4,6,8\}$ and $\mathbf{u}^{-1}\{1\} = \{1,3,5,7\}$. Let $\mathcal{M} = ((2,3),(6,1),(4,5),(8,7))$ and $\mathbf{a} = 0110$. Then $\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a})) = \mathbf{v} = 00110110$ (endpoints of the 2^{nd} matching edge $(6,1)$ and the 3^{rd} matching edge $(4,5)$ are flipped).*

Next, we define a weighted graph on all the balanced strings with weights representing the probability of going from one balanced string to another for a random matching $\mathcal{M}$ and a random string $\mathbf{a}$ (using the map Γ).

Let $n' = |\binom{n}{n/2}|$ denote the cardinality of the set of balanced strings $\{0,1\}^n_{n/2}$. Let G denote a weighted complete graph on n' vertices, where the vertices denote strings in $\{0,1\}^n_{n/2}$. For any two distinct balanced strings $\mathbf{u}, \mathbf{v} \in \{0,1\}^n_{n/2}$, the weight of the edge $(\mathbf{u}, \mathbf{v})$, denoted by $w(\mathbf{u}, \mathbf{v})$ is:

$$w(\mathbf{u}, \mathbf{v}) \;:=\; \Pr_{\mathcal{M}, \mathbf{a}}[\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a})) = \mathbf{v}],$$

where the probability is over the choice of a random perfect labeled matching $\mathcal{M}$ between $\mathbf{u}^{-1}\{0\}$ and $\mathbf{u}^{-1}\{1\}$, and a uniformly random string $\mathbf{a} \in \{0,1\}^{n/2}$. For every balanced string $\mathbf{u} \in \{0,1\}^n_{n/2}$, we will denote by $W(\mathbf{u})$ the distribution on $\{0,1\}^n_{n/2}$ where the probability of sampling $\mathbf{v}$ is equal to $w(\mathbf{u}, \mathbf{v})$. Let $W \in \mathbb{R}^{n' \times n'}$ denote the weighted adjacency matrix of G, i.e.

$$W[\mathbf{u}, \mathbf{v}] \;=\; w(\mathbf{u}, \mathbf{v}), \qquad \text{for all } \mathbf{u}, \mathbf{v} \in \{0,1\}^n_{n/2}$$

We are now ready to state the main technical lemma of our proof. It roughly says that if we sample a random vertex (which is a random balanced string) and its neighbour in the above-mentioned graph, then the two balanced strings behave "almost like pairwise-independent" points. In other words, the above-mentioned graph is a good sampler for the balanced slice $\{0,1\}^n_{n/2}$.

▶ **Lemma 9** (Main Lemma). *There exists a constant $\varepsilon > 0$ for which the following holds. Let G be the graph as mentioned above and let $S \subseteq \{0,1\}^n_{n/2}$ be an arbitrary subset of vertices with $|S| \geq 4^{-d} \cdot \binom{n}{n/2}$. Let ρ denote the density of the set S. Then,*

$$\Pr_{\substack{\mathbf{x} \sim U_{n,n/2} \\ \mathbf{y} \sim N(\mathbf{x})}}[\mathbf{x} \in S \;\; \textit{and} \;\; \mathbf{y} \in S] \;\leq\; \rho^2 \cdot \left(1 + \frac{1}{n^\varepsilon}\right).$$

We will give two proofs for Lemma 9, for two regimes of the degree d:

1. For degree $d \leq C \log n$ for some absolute constant $C > 0$, we give a simple argument using the spectral expansion properties of Cayley graphs and the expander mixing lemma. We prove this in Section 3.1.
2. For degree $d \leq n^\gamma$ for some absolute constant $\gamma > 0$, we rely on the spectrum of Johnson association schemes and use hypercontractivity for slice functions. See the full version.

We will also need a lower bound on the probability in Lemma 9. This will hold for all degree d. Combining the upper and lower bounds (i.e., Lemma 9 and Lemma 10 gives the final bound: see Theorem 16.)

▶ **Lemma 10** (The lower bound). *Let G be the graph as mentioned above and fix a degree parameter $d \in \mathbb{N}$. Let $P(\mathbf{x}) : \{0,1\}^n_{n/2} \to \mathbb{R}$ be a non-zero polynomial on the balanced slice $\{0,1\}^n_{n/2}$ with $\deg(P) \leq d$. If $S \subseteq \{0,1\}^n_{n/2}$ denote the set of non-zeroes of $P(\mathbf{x})$, then,*

$$\Pr_{\substack{\mathbf{x} \sim U_{n,n/2} \\ \mathbf{y} \sim W(\mathbf{x})}}[\mathbf{x} \in S \;\; \textit{and} \;\; \mathbf{y} \in S] \;\geq\; \frac{|S|}{\binom{n}{n/2}} \cdot \frac{1}{2^d}.$$

Proof of Lemma 10. Note that it is sufficient to show that

$$\Pr_{\mathbf{y} \sim W(\mathbf{x})}[\mathbf{y} \in S \,|\, \mathbf{x} \in S] \;\geq\; \frac{1}{2^d}, \qquad \text{for all } \mathbf{x} \in S.$$

Fix an arbitrary point $\mathbf{u} \in S$ and fix an arbitrary matching $\mathcal{M}$ between $\mathbf{u}^{-1}\{0\}$ and $\mathbf{u}^{-1}\{1\}$. We will show that for $1/2^d$-fraction of $\mathbf{a} \in \{0,1\}^{n/2}$, the string $\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a})) \in S$.

Define the polynomial $Q(z_1, \ldots, z_{n/2}) := P(\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{z})))$. Note that $\deg(Q) \leq \deg(P) \leq d$ and $Q(\mathbf{0}) = P(\mathbf{u}) \neq 0$. Now using the standard ODLSZ lemma (Theorem 4) on $Q(\mathbf{z})$, we get,

$$\Pr_{\mathbf{z} \sim \{0,1\}^{n/2}}[Q(\mathbf{z}) \neq 0] \geq \frac{1}{2^d} \quad \Rightarrow \quad \Pr_{\mathbf{a} \sim \{0,1\}^{n/2}}[\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a})) \in S] \geq \frac{1}{2^d}$$

Since the above lower bound holds for every matching $\mathcal{M}$ between $\mathbf{u}^{-1}\{1\}$ and $\mathbf{u}^{-1}\{0\}$, we have,

$$\Pr_{\mathcal{M}, \mathbf{a}}[\Gamma(\mathbf{u}, (\mathcal{M}, \mathbf{a})) \in S] \geq \frac{1}{2^d}$$

Since the above lower bound holds for arbitrary choice of $\mathbf{u} \in S$, this completes the proof of Lemma 10. ◀

Next, we observe that the random process mentioned above is "S_n-invariant"[4], i.e. the probabilities do not change even if we simultaneously permute the coordinates of $\mathbf{u}$ and $\mathbf{v}$ (using the same permutation for both of them).

▶ **Observation 11.** *For any $\mathbf{u}, \mathbf{v} \in \{0,1\}^n_{n/2}$, the weight $w(\mathbf{u}, \mathbf{v})$ depends only on[5] $\Delta(\mathbf{u}, \mathbf{v})$. We have,*

$$w(\mathbf{u}, \mathbf{v}) = \frac{\Delta!(n/2 - \Delta)!}{2^{n/2} \cdot (n/2)!} = \frac{1}{2^{n/2} \cdot \binom{n/2}{\Delta}}, \qquad \text{where } 2\Delta = \Delta(\mathbf{u}, \mathbf{v}) \in [0, n]$$

To see the above probability, observe that the $\frac{1}{2^{n/2}}$ factor corresponds to sampling the right $\mathbf{a}$ and the $\frac{\Delta!(n/2-\Delta)!}{(n/2)!}$ factor corresponds to picking a matching $\mathcal{M}$ that results in the output $\mathbf{v}$).

Note that the above observation in particular implies that the weighted adjacency matrix W is a real symmetric matrix, and thus has real eigenvalues. Both of our proofs for Lemma 9 will be based on upper bounding the eigenvalues of W.

3.1 Simple Proof Using Cayley Graphs

In this section, we prove a version of Lemma 9 using a simple and (mostly) self-contained argument. This version holds for degrees $d \leq C \log n$ for some absolute constant C.

Let $1 = \mu_1 \geq \mu_2 \geq \cdots \geq \mu_{n'}$ be the eigenvalues of W and let $\mu(W)$ denote the second largest eigenvalue in absolute value, i.e. $\mu(W) := \max(|\mu_2|, |\mu_{n'}|)$. A small value of $\mu(W)$ suggests that the random walk represented by W is "expanding" (see Lemma 7). The main lemma of this subsection is the following, which shows that $\mu(W)$ is small, i.e. W is a good expander. In the rest of the subsection, let $m = n/2$.

▶ **Lemma 12** (W is a good expander). *Let W denote the $n' \times n'$ matrix as described before. Then,*

$$\mu(W) \leq \mathcal{O}\left(\frac{\log n}{\sqrt{n}}\right).$$

[4] S_n is the group of permutations on n elements.
[5] Recall that $\Delta(\cdot, \cdot)$ represents the Hamming distance.

We now prove Lemma 12, i.e., we show that W is a good expander. The idea of the proof is to show that one can turn W into a (weighted) Cayley graph by adding additional edges and vertices and deduce that the original graph is an expander by using the expansion of the Cayley graph. In particular, we will show that W is an induced subgraph of a Cayley graph and use the interlacing of eigenvalues to prove the expansion.

Proof of Lemma 12. For the proof, we will assume that m is even; the odd case is handled similarly. Let $\{0,1\}^n_{odd}$ and $\{0,1\}^n_{even}$ denote the sets of points in $\{0,1\}^n$ that are of odd Hamming weight and even Hamming weight respectively. We will now define the weighted Cayley graph.

Let $V' = \{0,1\}^n_{even}$ (note that $\{0,1\}^n_{n/2} \subseteq V'$ as n is assumed to be even). Note that V' is an Abelian group with addition defined by performing coordinate sums modulo 2 (in particular, we may identify {0,1} with $\mathbb{F}_2$). We shall define a weighted Cayley graph W' over vertices V' by specifying its generators (and their weights) as follows. The set of generators is $\mathcal{S} = \{0,1\}^n_{even}$ and a generator $\mathbf{x} \in \mathcal{S}$ has weight

$$w(\mathbf{x}) \;=\; \frac{1}{2^m \cdot \binom{m}{\Delta}}, \quad \text{where } |\mathbf{x}| = 2\Delta \quad \text{and} \quad 0 \le \Delta \le m$$

With the above definition for V' and W', we note that the induced subgraph of W' when restricted to the balanced slice $\{0,1\}^n_{n/2} \subseteq V'$, is identical to W. Hence, by applying the eigenvalue interlacing theorem, we have the following.

▷ Claim 13 (Eigenvalue interlacing, see e.g. [22]). Let $\mu'_1 \ge \mu'_2 \ge \cdots \ge \mu'_{|V'|}$ be the eigenvalues of W'. Then $\mu_2 \le \mu'_2$ and $\mu'_{|V'|} \le \mu_{n'}$. Hence, $\mu(W) \le \max(|\mu'_2|, |\mu'_{|V'|}|)$.

The above claim allows us to bound $\mu(W)$ by bounding the absolute values of the eigenvalues of W' (except the largest). To do this, we will first fix an eigenbasis for W'. The characteristic vectors of the first $(n-1)$ variables forms such an eigenbasis (because if $\mathbf{x} \in V'$, then x_n can be expressed as a $\mathbb{F}_2$-linear combination of $x_1, \ldots, x_{n-1}$). That is, for $A \subseteq [n-1]$, the characteristic vector $\chi_A \in \mathbb{R}^{2^{n-1}}$ is defined as $\chi_A(\mathbf{x}) := (-1)^{\sum_{i \in A} x_i}$ for $\mathbf{x} \in V'$. The corresponding eigenvalue of χ_A is denoted by μ'_A (with a slight abuse of notation), and by Lemma 6, is equal to

$$\mu'_A = \sum_{\mathbf{y} \in \mathcal{S}} w(\mathbf{y}) \chi_A(\mathbf{y}).$$

It will be convenient to normalize the weights of the generators in $\mathcal{S}$ to make it a probability distribution. More formally, let $\mathcal{D}$ be the probability distribution over $\mathcal{S}$ where the probability of sampling a point $\mathbf{x} \in \mathcal{S}$ is equal to $w(\mathbf{x}) / \sum_{\mathbf{y} \in \mathcal{S}} w(\mathbf{y})$. Thus, $\mu'_\emptyset = \sum_{\mathbf{y} \in \mathcal{S}} w(\mathbf{y})$ and $\mu'_A / \mu'_\emptyset = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\chi_A(\mathbf{x})]$.

We will now show that $0 < \mu'_\emptyset \le \mathcal{O}(\sqrt{m})$ and $|\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\chi_A(\mathbf{x})]| = |\mu'_A / \mu'_\emptyset| \le O\big(\frac{\log m}{m}\big)$ for all *non-empty* $A \subseteq [2m-1]$. This would in turn give that $\mu(W) \le \max_{A \neq \emptyset}(|\mu'_A|) \le \mathcal{O}\left(\frac{\sqrt{m} \log m}{m}\right) = \mathcal{O}(\log m / \sqrt{m})$, finishing the proof of Lemma 12.

The proof of Claim 14 follows from a simple counting argument, and we defer it to the end.

▷ Claim 14. $0 < \mu'_\emptyset \le \mathcal{O}(\sqrt{m})$.

It remains to show that $|\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\chi_A(\mathbf{x})]| \le \mathcal{O}\big(\frac{\log m}{m}\big)$ for all *non-empty* $A \subseteq [2m-1]$. Note that the distribution $\mathcal{D}$ has some symmetry in the sense that all the points of a given Hamming weight have the same probability. Furthermore, two given points, one of weight

2Δ and the other of weight $(2m-2\Delta)$ also have the same probability mass (for arbitrary $0 \le \Delta \le m$). This leads to $\mathbb{E}_{\mathbf{x}\sim\mathcal{D}}[\chi_A(\mathbf{x})]$ being equal to 0 if $|A| = 1$ or $2m-1$, and hence it suffices to focus on the regime $2 \le |A| \le 2m-2$.

We show the following concentration inequality for the distribution $\mathcal{D}$. The proof of Claim 15 can we defer it to the end.

▷ Claim 15. $\Pr_{\mathbf{x}\sim\mathcal{D}}[||\mathbf{x}| - m| > \sqrt{50m\log m}] \le \mathcal{O}(1/m^2)$.

Assuming Claim 15, it suffices to show that $\mathbb{E}_{\mathbf{x}\sim\mathcal{D}}[\chi_A(\mathbf{x}) \mid |\mathbf{x}| \in m \pm \sqrt{50m\log m}] \le \mathcal{O}\big(\frac{\log m}{m}\big)$ to conclude that $\mathbb{E}_{\mathbf{x}\sim\mathcal{D}}[\chi_A(\mathbf{x})] \le \mathcal{O}\big(\frac{\log m}{m}\big)$. We will show that this holds even conditioned on $|\mathbf{x}| = 2\Delta$ for every $\Delta \in (m \pm \sqrt{50m\log m})/2$. However, recall that $\mathcal{D}$ is uniform when restricted to $\{0,1\}^{2m}_{2\Delta}$. Therefore, we can equivalently upper bound the quantity $|\mathbb{E}_{\mathbf{x}\in\{0,1\}^{2m}_{2\Delta}}[\chi_A(\mathbf{x})]|$ to conclude the proof. Now, we note that $\mathbb{E}_{\mathbf{x}\in\{0,1\}^{2m}_{2\Delta}}[\chi_A(\mathbf{x})] = \mathbb{E}_{B\sim\binom{[2m]}{|A|}}[\chi_B(\mathbf{c})]$, where $\mathbf{c}$ is an arbitrary point in $\{0,1\}^{2m}_{2\Delta}$ (we will fix it to be $0^{2m-2\Delta}1^{2\Delta}$). Hence, it suffices to show that

$$\left|\mathbb{E}_{B\sim\binom{[2m]}{k}}[\chi_B(\mathbf{c})]\right| \le \mathcal{O}\left(\frac{\log m}{m}\right), \tag{1}$$

for every $2 \le k \le (2m-2)$ (since we assumed that $2 \le |A| \le 2m-2$). We may further assume that $k \le m$ without loss of generality, as $\chi_B(\mathbf{c}) = \chi_{\overline{B}}(\mathbf{c})$.

To help with the analysis, we will choose $B \sim \binom{[2m]}{k}$ by first choosing a subset C of $[2m]$ of size $(k-2)$ (which is non-negative) uniformly at random and then choosing two elements $b_1 \ne b_2$ from $\overline{C} = [2m] \setminus C$ uniformly at random.

For a subset $C \subseteq [2m]$, we will use the notation $\mathrm{wt}(C)$ to denote the number of 1's in $\mathbf{c}$ when restricted to the coordinates indexed by C. Let $p := \mathrm{wt}([2m])/(2m)$ denote the fractional Hamming weight of $\mathbf{c}$. We say that a subset $C \subseteq [2m]$ is *good*, if $\big||\overline{C}| - 2\mathrm{wt}(\overline{C})\big| \le \sqrt{2000\log m}$. We claim that $\Pr_{C\sim\binom{[2m]}{k-2}}[C \text{ is not good}] \le \mathcal{O}(1/m^2)$. This essentially follows from standard tail bounds for the hypergeometric distribution but needs some care to handle small k. We divide this analysis into two cases.

- **Case 1:** $k \le \sqrt{50m\log m}$**.** We note that $|\overline{C}| = 2m - k + 2 \in 2m \pm \sqrt{50m\log m}$, and similarly $\mathrm{wt}(\overline{C}) \in \mathrm{wt}([2m]) \pm \sqrt{50m\log m} \subseteq m \pm 2\sqrt{50m\log m}$. Using these bounds, it follows that $\big||\overline{C}| - 2\mathrm{wt}(\overline{C})\big| \le \sqrt{2000\log m}$ for sufficiently large m (for every choice of $C \in \binom{[2m]}{k-2}$). Hence all choices of C are good in this case.
- **Case 2:** $k > \sqrt{50m\log m}$**.** We note that $\mathrm{wt}(C)$ is distributed according to a hypergeometric distribution – it corresponds to the number of successes in $k-2$ draws with replacement from a population of $2m$ total states and $\mathrm{wt}([2m]) = 2d$ success states. Using a standard tail bound [20], we obtain that $\Pr_{C\sim\binom{[2m]}{k-2}}\big[\big|\frac{\mathrm{wt}(C)}{k-2} - p\big| > \sqrt{\frac{4\log k}{k}}\big] = \mathcal{O}(1/k^4) = \mathcal{O}(1/m^2)$. Using $\big|p - \frac{1}{2}\big| \le \sqrt{\frac{50\log m}{m}}$, we thus get that $||C| - 2\mathrm{wt}(C)| \le 4\sqrt{50m\log m}$ with probability at least $1 - \mathcal{O}(1/m^2)$. Because $2m = |C| + \overline{C}$ and $2d = \mathrm{wt}(C) + \mathrm{wt}(\overline{C})$, with probability at least $1 - \mathcal{O}(1/m^2)$, we have $||\overline{C}| - 2\mathrm{wt}(\overline{C})| \le 6\sqrt{50m\log m}$, i.e., C is good.

We now show that conditioned on C being good, the expectation of $\chi_B(\mathbf{c})$ is upper bounded by $\mathcal{O}\big(\frac{\log m}{m}\big)$ in absolute value. This would then prove (1). For ease of notation, let $n_0 = |\overline{C}| = 2m - k + 2 \ge m$ and $d_0 = \mathrm{wt}(\overline{C}) \in \frac{n_0}{2} \pm \Theta(\sqrt{n_0\log n_0})$ (since C is good). We now note that $\chi_B(\mathbf{c}) = \chi_C(\mathbf{c}) \cdot \chi_{\{b_1,b_2\}}(\mathbf{c})$, so it suffices to bound $|\mathbb{E}_{b_1,b_2}[(-1)^{c_{b_1}+c_{b_2}}]|$. The idea now is that c_{b_1} and c_{b_2} almost behave like two independent draws, so the expectation is

roughly the square of $|\mathbb{E}_b[(-1)^{c_b}]|$ for a uniformly random coordinate $b \in \overline{C}$, which is equal to $\left|\frac{n_0-2d_0}{n_0}\right| \le O(\sqrt{\frac{\log n_0}{n_0}}) \le O(\sqrt{\frac{\log m}{m}})$ as $n_0 \ge m$. More precisely, we have the following:

$$\begin{aligned}
|\mathbb{E}_{b_1,b_2}[(-1)^{c_{b_1}+c_{b_2}}]| &= \frac{|\binom{d_0}{2} + \binom{n_0-d_0}{2} - d_0(n_0-d_0)|}{\binom{n_0}{2}} \\
&= \frac{|(n_0-2d_0)^2 - n_0|}{n_0(n_0-1)} \\
&\le \mathcal{O}\left(\frac{\log n_0}{n_0}\right) \le \mathcal{O}\left(\frac{\log n}{n}\right). \qquad \text{(as } d_0 \in n_0/2 \pm \Theta(\sqrt{n_0 \log n_0})\text{)}
\end{aligned}$$

Now we prove Claim 14 and Claim 15.

Proof of Claim 14. By the definition of the weight function w of the generators, we have

$$\begin{aligned}
\mu'_\emptyset = \sum_{\mathbf{y}\in\{0,1\}^{2m}_e} w(\mathbf{y}) = \sum_{d=0}^{m} \binom{2m}{2d} \cdot \frac{1}{2^m \cdot \binom{m}{d}} = \sum_{d=0}^{m} \frac{\binom{2m}{2d}}{\binom{m}{d}^2} \cdot \frac{\binom{m}{d}}{2^m} &\le \max_{d\in[0..m]} \left(\frac{\binom{2m}{2d}}{\binom{m}{d}^2}\right) \cdot \sum_{d=0}^{m} \frac{\binom{m}{d}}{2^m} \\
&= \max_{d\in[0..m]} \left(\frac{\binom{2m}{2d}}{\binom{m}{d}^2}\right).
\end{aligned}$$

Now, for every $d \in [0..m]$,

$$\begin{aligned}
\frac{\binom{2m}{2d}}{\binom{m}{d}^2} = \frac{(2m)!d!^2(m-d)!^2}{(2d)!(2m-2d)!m!^2} &= \frac{\binom{2m}{m}}{\binom{2d}{d}\binom{2m-2d}{m-d}} \\
= O\left(\frac{2^{2m}}{\sqrt{m}} \cdot \frac{\sqrt{d}}{2^{2d}} \cdot \frac{\sqrt{m-d}}{2^{2m-2d}}\right) &= O\left(\sqrt{\frac{d(m-d)}{m}}\right) \le O(\sqrt{m}),
\end{aligned}$$

where the last inequality uses the AM-GM inequality. Therefore, $0 < \mu'_\emptyset \le O(\sqrt{m})$. ◁

Finally, we prove Claim 15.

Proof of Claim 15. Letting $B := \{d \in [0..m] \mid |2d-m| > \sqrt{50m\log m}\}$, we can explicitly express the probability as

$$\Pr_{\mathbf{x}\sim\mathcal{D}}[||\mathbf{x}|-m| > \sqrt{50m\log m}] = \sum_{d\in B} \binom{2m}{2d} \cdot \frac{1}{2^m\binom{m}{d}} = \sum_{d\in B} \frac{\binom{2m}{2d}}{\binom{m}{d}^2} \cdot \frac{\binom{m}{d}}{2^m} \le O(\sqrt{m}) \cdot \sum_{d\in B} \frac{\binom{m}{d}}{2^m},$$

by using the bound $\frac{\binom{2m}{2d}}{\binom{m}{d}^2} \le O(\sqrt{m})$ from the proof of Claim 14. To bound the second factor $\sum_{d\in B} \frac{\binom{m}{d}}{2^m}$, we use a Chernoff bound for the sum of m i.i.d. copies of a uniformly random Boolean variable. In particular, we get $\sum_{d\in B} \frac{\binom{m}{d}}{2^m} \le O(1/m^3)$. Hence $\Pr_{\mathbf{x}\sim\mathcal{D}}[||x|-m| > \sqrt{50m\log m}] \le O(1/m^2)$. ◁

This finishes the proof of Lemma 12. ◀

Proof of Lemma 9 for $d \le C\log n$. We use Lemma 12 to prove an upper bound for Lemma 9 that holds for all $d \le C\log n$ for an absolute constant $C > 0$. Using the expander mixing lemma (Lemma 7), we obtain

$$\Pr_{\mathbf{x}\sim U_{n,n/2},\mathbf{y}\sim W(\mathbf{x})}[\mathbf{x} \in S \text{ and } \mathbf{y} \in S] \le \rho^2 + \rho \cdot \mathcal{O}\left(\frac{\log n}{\sqrt{n}}\right),$$

where S denotes the non-zeroes of P in $\{0,1\}^n_{n/2}$ and $\rho = |S|/\binom{n}{n/2}$. Thus, assuming $d \leq C \log n$ for small enough constant C and using $\rho \geq 4^{-d}$ (Lemma 2), we get that the above probability is at most $\rho^2(1+1/n^\varepsilon)$ for sufficiently small constant ε. Hence, this finishes the proof of Lemma 9 in the regime $d \leq C \log n$. ◀

A distance lemma for the balanced slice

Using Lemma 9, we now prove a distance lemma for non-zero degree-d polynomials over the balanced slice $\{0,1\}^n_{n/2}$. Following we give a proof in the setting when $d = \mathcal{O}(\log n)$. For the setting when $d = n^\varepsilon$, please refer to the full version.

▶ **Theorem 16** (Distance lemma over the balanced slice). *There exists an absolute constant $C > 0$ so that the following holds. Fix an arbitrary Abelian group G and fix a degree parameter $d \in \mathbb{N}$ where $d \leq C \log n$. For every even natural number n, and for every non-zero degree-d polynomial $P(\mathbf{x}) \in \mathcal{P}_d(n, n/2, G)$,*

$$\Pr_{\mathbf{x} \sim U_{n,n/2}}[P(\mathbf{x}) \neq 0] \;\geq\; \frac{1}{2^d} \cdot \left(1 - \frac{1}{n^{\Omega(1)}}\right)$$

Proof of Theorem 16. Letting $S \subseteq \{0,1\}^n_{n/2}$ denote the set of points on the balanced slice on which P evaluates to a non-zero value. From Lemma 2, we know that $|S|/n' \geq 4^{-d}$. By combining Lemma 9 and Lemma 10, we obtain

$$\frac{|S|}{n'} \cdot \frac{1}{2^d} \leq \Pr_{\substack{\mathbf{x} \sim U_{n,n/2} \\ \mathbf{y} \sim W(\mathbf{x})}}[\mathbf{x} \in S \;\text{ and }\; \mathbf{y} \in S] \leq \left(\frac{|S|}{n'}\right)^2 \cdot \left(1 + \frac{1}{n^\varepsilon}\right),$$

where ε is a sufficiently small constant. Hence, $|S|/n' \geq \frac{1}{2^d} \cdot \left(1 - \frac{1}{n^{\Omega(1)}}\right)$. ◀

4 Low-degree Functions Over Slices

In this section we will give a simple proof of a lemma of Filmus and Ihringer [17] following the proof idea of Nisan and Szegedy [26]. We will first give a couple of definitions and set the notations for this section.

For a function $f(x_1, \ldots, x_n)$ on the slice $\{0,1\}^n_k$ with coefficients in $\mathbb{R}$, and for any two coordinates $i, j \in [n] \times [n]$, define $f^{(ij)}$ to be the function where we swap the i^{th} variable with the j^{th} variable in the function f.

▶ **Definition 17** (Influence). *Let $f : \{0,1\}^n \to \mathbb{R}$ be a function on the slice $\{0,1\}^n_k$. For $(i,j) \in [n] \times [n]$, the $(i,j)^{th}$-influence of f, denoted by $\mathrm{Inf}_{ij}(f)$, is defined as,*

$$\mathrm{Inf}_{ij}(f) \;:=\; \frac{1}{4} \Pr_{\mathbf{x} \sim U_{n,k}}[f(\mathbf{x}) \neq f^{(ij)}(\mathbf{x})]$$

The total influence of f, denoted by $\mathrm{Inf}(f)$, is defined as,

$$\mathrm{Inf}(f) \;:=\; \frac{1}{n} \sum_{1 \leq i < j \leq n} \mathrm{Inf}_{ij}(f)$$

Note that if $i = j$ in the above definition, then $\mathrm{Inf}_{ii}(f) = 0$ and it does not contribute anything towards the total influence.

A key lemma in the proof of [17] is a lower bound on every non-zero influence (see [17, Lemma 3.1]). They showed that there exists a constant α such that every non-zero influence of a degree-d polynomial on the balanced slice is at least α^d. The proof of this lemma in [17] uses analytic techniques such as the Log-Sobolev inequality on the Boolean slice [25] and the Hypercontractive inequality [13]. Using our distance lemma for the balanced slices (Theorem 16), we can improve the lower bound to almost $1/2^d$ (which is easily seen to be tight up to constant factors). Note that the main result of this section only holds for degree $d \leq C \log n$ for some absolute constant $C > 0$, it suffices to use the simpler proof of Theorem 16. We state the lemma below.

▶ **Lemma 18** (Lower bound on influences). *Let $f(x_1, \ldots, x_n)$ be a non-zero degree-d function on the balanced slice $\{0,1\}^n_{n/2}$. Then for every $(i,j) \in [n] \times [n]$ for which $\mathrm{Inf}_{ij}(f) > 0$, the following holds:*

$$\mathrm{Inf}_{ij}(f) \;\geq\; \frac{1}{4} \cdot \frac{1}{2^d} \cdot \left(1 - \frac{1}{n^\varepsilon}\right)$$

for some absolute constant $\varepsilon > 0$.

Proof. Fix some pair $(i,j) \in [n] \times [n]$ with $\mathrm{Inf}^{(ij)}(f) > 0$ and consider the polynomial g_{ij} on the balanced slice, defined as follows: $g_{ij}(\mathbf{x}) := f(\mathbf{x}) - f^{(ij)}(\mathbf{x})$.
Observe that since f is a degree-d polynomial, $f^{(ij)}$ is also a degree-d polynomial, which means g_{ij} is also a degree-d polynomial on the balanced slice. Since the influence $\mathrm{Inf}_{ij}(f) > 0$, this means that g_{ij} is non-zero on the slice $\{0,1\}^n_{n/2}$. Now using our distance lemma on the balanced slice Theorem 16,

$$\Pr_{\mathbf{x} \sim U_{n,n/2}}[f(\mathbf{x}) \neq f^{(ij)}(\mathbf{x})] \;=\; \Pr_{\mathbf{x} \sim U_{n,n/2}}[g_{ij}(\mathbf{x}) \neq 0] \;\geq\; \frac{1}{2^d} \cdot \left(1 - \frac{1}{n^\varepsilon}\right)$$
$$\Rightarrow \mathrm{Inf}_{ij}(f) \;\geq\; \frac{1}{4} \cdot \frac{1}{2^d} \cdot \left(1 - \frac{1}{n^\varepsilon}\right)$$

for some absolute constant $\varepsilon > 0$. ◀

Filmus and Ihringer [17, Lemma 3.3] use this lower bound on non-zero influences to get a bound on junta of degree-d polynomials on the balanced slice. Using the above-mentioned improved lower bound on non-zero influence, we can also improve the bounds in [17, Lemma 3.3].

▶ **Lemma 19.** *There exists an absolute constant $C > 0$ such that for all degree parameters $d \in \mathbb{N}$ such that $d \leq C \log n$, the following holds. Every degree-d polynomial on the slice $\{0,1\}^n_{n/2}$ is a $\eta(d)$-junta, where*

$$\eta(d) = \mathcal{O}(d \cdot 2^d).$$

Proof. The proof is essentially the same proof as in [17], except for one inequality which can be improved using Lemma 18. Let $f(\mathbf{x}) \in \mathcal{P}_d(n, n/2, \mathbb{R})$ and let G be a graph on the vertex set $[n]$ where (i,j) is an edge if $\mathrm{Inf}_{ij}(f) \geq 1/2^d \cdot (1 - 1/n^\varepsilon)$, where ε is the absolute constant from Lemma 18. Let M be a maximal matching in G. We now proceed similar to the proof in [17, Lemma 3.3] and we request the reader to refer [17] as we just highlight the changes in the proof here.

Using Lemma 18, we get the following two inequalities upper and lower bounding the influence:

$$\left(\frac{1}{4}\cdot\frac{1}{2^d}\cdot\left(1-\frac{1}{n^\varepsilon}\right)\right)\cdot\left(1-\frac{1}{n}\right)\cdot M \;\leq\; \mathrm{Inf}(f) \;\leq\; d$$
$$\Rightarrow M \leq \mathcal{O}(d\cdot 2^d),$$

where we used the assumption $d \leq C\log n$ in upper bounding $1/n^\varepsilon$ by $\frac{1}{10}\cdot\frac{1}{2^d}$. Following the argument of [17], this gives us that f is a $2M$-junta, i.e., a $\mathcal{O}(d\cdot 2^d)$-junta. ◀

As already noted in the introduction, a stronger upper bound of $\eta(d) = \mathcal{O}(2^d)$ follows from the work of [17, 9] (and can also be obtained by plugging Lemma 18 in place of [17, Lemma 3.3] in the proof of [17]). The advantage here is the relatively simple proof following exactly the template of [26].

References

1. N. Alon and F.R.K. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1):15–19, 1988. doi:10.1016/0012-365X(88)90189-6.
2. Noga Alon, Ernest E. Bergmann, Don Coppersmith, and Andrew M. Odlyzko. Balancing sets of vectors. *IEEE Trans. Inf. Theory*, 34(1):128–130, 1988. doi:10.1109/18.2610.
3. Noga Alon and Zoltán Füredi. Covering the cube by affine hyperplanes. *Eur. J. Comb.*, 14:79–83, 1993. doi:10.1006/EUJC.1993.1011.
4. Prashanth Amireddy, Amik Raj Behera, Manaswi Paraashar, Srikanth Srinivasan, and Madhu Sudan. Low Degree Local Correction Over the Boolean Cube . *Electron. Colloquium Comput. Complex.*, TR24-164, 2024. URL: https://eccc.weizmann.ac.il/report/2024/164.
5. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. doi:10.1145/278298.278306.
6. László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In Cris Koutsougeras and Jeffrey Scott Vitter, editors, *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 21–31. ACM, 1991. doi:10.1145/103418.103428.
7. Abhishek Bhrushundi, Kaave Hosseini, Shachar Lovett, and Sankeerth Rao. Torus polynomials: An algebraic approach to ACC lower bounds. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 13:1–13:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.ITCS.2019.13.
8. Anurag Bishnoi, Simona Boyadzhiyska, Shagnik Das, and Tamás Mészáros. Subspace coverings with multiplicities. *Combinatorics, Probability and Computing*, 32(5):782–795, 2023. doi:10.1017/S0963548323000123.
9. John Chiarelli, Pooya Hatami, and Michael Saks. An asymptotically tight bound on the number of relevant variables in a bounded degree boolean function. *Combinatorica*, 40(2):237–244, April 2020. doi:10.1007/s00493-019-4136-7.
10. Alexander Clifton and Hao Huang. On almost k-covers of hypercubes. *Combinatorica*, 40(4):511–526, 2020. doi:10.1007/S00493-019-4221-Y.
11. Roee David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. Direct sum testing. *SIAM Journal on Computing*, 46(4):1336–1369, 2017. doi:10.1137/16M1061655.
12. Richard A. DeMillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978. doi:10.1016/0020-0190(78)90067-4.
13. Persi Diaconis and Laurent Saloff-Coste. Logarithmic sobolev inequalities for finite markov chains. *The Annals of Applied Probability*, 6(3):695–750, 1996.

14 Irit Dinur, Yuval Filmus, and Prahladh Harsha. Agreement tests on graphs and hypergraphs. *Electron. Colloquium Comput. Complex.*, TR17, 2017. URL: https://api.semanticscholar.org/CorpusID:452567.

15 Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to kakeya sets and mergers. *SIAM J. Comput.*, 42(6):2305–2328, 2013. doi:10.1137/100783704.

16 Yuval Filmus. An orthogonal basis for functions over a slice of the boolean hypercube. *The Electronic Journal of Combinatorics*, 23:1, 2016. doi:10.37236/4567.

17 Yuval Filmus and Ferdinand Ihringer. Boolean constant degree functions on the slice are juntas. *Discrete Mathematics*, 342(12):111614, 2019. doi:10.1016/j.disc.2019.111614.

18 Yuval Filmus, Guy Kindler, Elchanan Mossel, and Karl Wimmer. Invariance principle on the slice. *ACM Trans. Comput. Theory*, 10(3):11:1–11:37, 2018. doi:10.1145/3186590.

19 Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. Essential coding theory (book draft), 2023. URL: http://www.cse.buffalo.edu/atri/courses/coding-theory/book.

20 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, pages 409–426, 1994.

21 Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43:439–561, 2006.

22 Roger A Horn and Charles R Johnson. Topics in matrix analysis, 1991. *Cambridge University Presss, Cambridge*, 37:39, 1991.

23 Pavel Hrubes, Sivaramakrishnan Natarajan Ramamoorthy, Anup Rao, and Amir Yehudayoff. Lower bounds on balancing sets and depth-2 threshold circuits. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 72:1–72:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.ICALP.2019.72.

24 Gil Kalai, Noam Lifshitz, Dor Minzer, and Tamar Ziegler. A dense model theorem for the boolean slice, 2024. To appear in the Proceedings of the 65th IEEE Symposium of Foundations of Computer Science (FOCS) 2024, Chicago, USA. arXiv:2402.05217.

25 Tzong-Yow Lee and Horng-Tzer Yau. Logarithmic sobolev inequality for some models of random walks. *The Annals of Probability*, 26(4):1855–1873, 1998.

26 Noam Nisan and Mario Szegedy. On the degree of boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1992. URL: https://api.semanticscholar.org/CorpusID:6919144.

27 Ryan O'Donnell and Karl Wimmer. Kkl, kruskal-katona, and monotone nets. *SIAM J. Comput.*, 42(6):2375–2399, 2013. doi:10.1137/100787325.

28 Øystein Ore. Über höhere kongruenzen. *Norsk Mat. Forenings Skrifter*, 1(7):15, 1922.

29 Ramamohan Paturi and Michael E. Saks. On threshold circuits for parity. In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I*, pages 397–404. IEEE Computer Society, 1990. doi:10.1109/FSCS.1990.89559.

30 Michael O Rabin and Vijay V Vazirani. Maximum matchings in general graphs through randomization. *Journal of Algorithms*, 10(4):557–567, 1989. doi:10.1016/0196-6774(89)90005-9.

31 Shubhangi Saraf and Madhu Sudan. An improved lower bound on the size of Kakeya sets over finite fields. *Analysis and PDE*, 1(3):375–379, 2008. doi:10.2140/apde.2008.1.375.

32 Lisa Sauermann and Yuval Wigderson. Polynomials that vanish to high order on most of the hypercube. *Journal of the London Mathematical Society*, 106(3):2379–2402, 2022. doi:10.1112/jlms.12637.

33 Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980. doi:10.1145/322217.322225.

34 Terence Tao and Tamar Ziegler. The inverse conjecture for the gowers norm over finite fields in low characteristic. *Annals of Combinatorics*, 16(1):121–188, 2012.

35 Richard M. Wilson. A diagonal form for the incidence matrices of t-subsets vs.k-subsets. *European Journal of Combinatorics*, 11(6):609–615, 1990. doi:10.1016/S0195-6698(13)80046-7.

36 Karl Wimmer. Low influence functions over slices of the boolean hypercube depend on few coordinates. In *IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014*, pages 120–131. IEEE Computer Society, 2014. doi:10.1109/CCC.2014.20.

37 Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation*, pages 216–226. Springer Berlin Heidelberg, 1979. doi:10.1007/3-540-09519-5_73.