





Eigenvalue Bounds for Symmetric Markov Chains on Multislices with Applications

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

We consider random walks on “balanced multislices” of any “grid” that respects the “symmetries” of the grid, and show that a broad class of such walks are good spectral expanders. (A grid is a set of points of the form \mathcal{S}^n for finite \mathcal{S} , and a balanced multi-slice is the subset that contains an equal number of coordinates taking every value in \mathcal{S} . A walk respects symmetries if the probability of going from $u = (u_1, \dots, u_n)$ to $v = (v_1, \dots, v_n)$ is invariant under simultaneous permutations of the coordinates of u and v .) Our main theorem shows that, under some technical conditions, every such walk where a single step leads to an almost $\mathcal{O}(1)$ -wise independent distribution on the next state, conditioned on the previous state, satisfies a non-trivially small singular value bound.

We give two applications of our theorem to error-correcting codes: (1) We give an analog of the Ore-DeMillo-Lipton-Schwartz-Zippel lemma for polynomials, and junta-sums, over balanced multislices. (2) We also give a local list-correction algorithm for d -junta-sums mapping an arbitrary grid \mathcal{S}^n to an Abelian group, correcting from a near-optimal $(\frac{1}{|\mathcal{S}|^d} - \varepsilon)$ fraction of errors for every $\varepsilon > 0$, where a d -junta-sum is a sum of (arbitrarily many) d -juntas (and a d -junta is a function that depends on only d of the n variables).

Our proofs are obtained by exploring the representation theory of the symmetric group and merging it with some careful spectral analysis.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains

Keywords and phrases Markov Chains, Random Walk, Multislices, Representation Theory of Symmetric Group, Local Correction, Low-degree Polynomials, Polynomial Distance Lemma

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.34

Category RANDOM

Related Version *Full Version:* <https://eccc.weizmann.ac.il/report/2025/093/>

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.

Acknowledgements We would like to thank the anonymous reviewers of RANDOM 2025 for many valuable comments, including pointers to crucial papers in the literature (specifically [7]), that significantly improved some of our proofs.



© Prashanth Amireddy, Amik Raj Behera, Srikanth Srinivasan, and Madhu Sudan;
licensed under Creative Commons License CC-BY 4.0
Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques
(APPROX/RANDOM 2025).

Editors: Alina Ene and Eshan Chattopadhyay; Article No. 34; pp. 34:1–34:12



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Consider the following natural random walk whose states are the balanced vectors of $\{0, 1\}^n$, i.e., the balanced Boolean slice with an equal number of 0s and 1s, where a single step of the random walk takes a state u to a state v at Hamming distance exactly $n/2$ from it. One would expect this random walk to mix extremely rapidly, and indeed this is known. The underlying graph here is a special case of a Johnson graph whose entire eigenspectrum is well known [8] and, in particular, implies that the second eigenvalue of this graph is $o_n(1)$.

Now consider the following variant of the above random walk: The states now are elements of the “balanced multislice” in $\{-1, 0, 1\}^n$, i.e. vectors with exactly $1/3$ rd fraction of the letters -1 , 0 and 1 , and in a single step from a balanced vector u to a random balanced vector v obtained by flipping exactly $1/3$ fraction of each of the letters of u to -1 , 0 and 1 . (So for a single coordinate i , v_i is uniform in $\{-1, 0, 1\}$ conditioned on u_i .) It is intuitive to believe that such a random walk should also converge to the uniform distribution over all balanced vectors extremely fast, but, as far as we know, it was not even known that the second-eigenvalue of this random walk (or its transition probability matrix) has value $o_n(1)$.

The gap in the understanding between the Boolean and non-Boolean cases in such problems can be significant for fundamental reasons. For example, for the alternate version of the random walk where the transition is defined by a uniformly random transposition of coordinates, it took a decade after optimal bounds on the mixing time were proved in the Boolean case [10] to prove similar results in the non-Boolean setting [19]. We refer the reader to the work of Filmus, O’Donnell, and Wu [14] for a nice overview of the challenges posed by the non-Boolean setting in such problems. Some of these obstructions have to do with associated representations of the symmetric group that play a role in the corresponding proofs; these representations are simpler (‘multiplicity-free’) in the Boolean setting than in the non-Boolean setting. This also creates difficulties in resolving the questions we consider.

The main contribution of this work is to address some of the challenges alluded to above. In particular, we show that the variant random walk described in the second paragraph also has fast mixing, specifically by giving a $o_n(1)$ bound on its second eigenvalue. Indeed, we study this question in more generality for balanced multislices, with “nearly balanced moves”. We believe the questions carry intrinsic interest and should find broad applications in the field. We justify this belief partially by describing two applications in coding theory:

- The first gives a near-tight distance bound on codes obtained by evaluations of polynomials on balanced multislices.
- The second gives a local list-correction algorithm for subclasses of polynomials evaluated on *grids*. (Note that the second application does not refer to balanced multislices in the problem definition – the multislices arise naturally in the design and analysis of the local correction algorithm!)

We elaborate on our setting and results, the applications, and the technical challenges below. In this extended abstract, we only include an overview of the proof techniques. For full proofs, see the full version of this paper [3].

1.1 Multislices and Random Walks

By a *grid*, we refer to sets of the form \mathcal{S}^n for some finite set \mathcal{S} and positive integer n . (Usually we think of $s := |\mathcal{S}|$ as a constant and study the growth of relevant parameters as a function of n). The *balanced multislice* of a grid \mathcal{S}^n is the set

$$\mathcal{S}_\mu^n := \left\{ \mathbf{a} \in \mathcal{S}^n \mid \forall \sigma \in \mathcal{S}, |\{i \in [n] \mid a_i = \sigma\}| = \frac{n}{s} \right\}.$$

(Note that a multislice is non-empty if and only if s divides n . We will drop the term “balanced” in the future and simply refer to multislices to keep the term short.)

The random walks we consider have the multislice of some grid as their state space. Recall that such a random walk can be described by a $\mathcal{S}_\mu^n \times \mathcal{S}_\mu^n$ matrix W with $W(\mathbf{a}, \mathbf{b})$ denoting the probability of transition from state \mathbf{a} to \mathbf{b} . We consider walks where every step of the walk makes a “nearly balanced move”. To elaborate, let us define the *generalized Hamming distance*¹ $\Delta(\mathbf{a}, \mathbf{b})$, for vectors $\mathbf{a}, \mathbf{b} \in \mathcal{S}^n$, to be the $\mathcal{S} \times \mathcal{S}$ matrix given by $\Delta_{\sigma, \tau}(\mathbf{a}, \mathbf{b}) = |\{i \in [n] \mid a_i = \sigma, b_i = \tau\}|$. We say that a generalized Hamming distance parameter $\Delta \in \mathbb{Z}^{\mathcal{S} \times \mathcal{S}}$ determines a random walk matrix, denoted W_Δ , if for each vertex \mathbf{a} , the random step corresponding to W_Δ is obtained by picking, uniformly at random, a vertex \mathbf{b} on the multislice such that $\Delta(\mathbf{a}, \mathbf{b}) = \Delta$.

For constant $C < \infty$, we say that a generalized Hamming distance parameter $\Delta \in \mathbb{Z}^{\mathcal{S} \times \mathcal{S}}$ is C -balanced if each entry of Δ is $\frac{m}{s} \pm (C\sqrt{m \log m})$ where $m = n/s$. In other words, all the entries of Δ are equal up to a difference of at most $2C\sqrt{m \log m}$. Informally, when considering $n \rightarrow \infty$ we refer to Δ , as also a random walk matrix W_Δ determined by Δ , as “nearly balanced” if Δ is C -balanced for some constant C . Here, we note that W_Δ is a well-defined random walk matrix over the multislice only if Δ/m is a doubly-stochastic matrix (i.e., every row and every column of Δ sums to m).

Note that the mixing time of a random walk matrix W is closely tied to the second largest singular value, which we denote by $\sigma_2(W)$. (In particular, the singular values satisfy $1 = \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N \geq 0$ and we let $\sigma_2(W) = \sigma_2$. If the walk is symmetric, then this captures the second eigenvalue. Specifically, if the eigenvalues are $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N \geq -1$ where $N = |\mathcal{S}_\mu^n|$, then $\sigma_2(W) = \max\{|\lambda_2|, |\lambda_N|\}$.) Our main goal is to bound the value of $\sigma_2(W)$ by some function $o_n(1)$ that tends to 0 with growing n for a broad class of random walk matrices W over the multislice \mathcal{S}_μ^n . In general, it is desirable to have such singular value bounds, and such random walk matrices are said to have good “spectral expansion” or “fast mixing”.

The following theorem gives such a fast mixing result on the balanced multislice for all nearly balanced walks that “respect symmetries”. More formally, for a permutation $\pi \in \text{Sym}_n$ and $\mathbf{a} \in \mathcal{S}^n$, let $\pi(\mathbf{a})$ denote the action of π on \mathcal{S}_μ^n , i.e., $\pi(\mathbf{a}) := (a_{\pi^{-1}(1)}, \dots, a_{\pi^{-1}(n)})$. For a stochastic matrix $M \in \mathbb{R}^{\mathcal{S}_\mu^n \times \mathcal{S}_\mu^n}$, we say M *respects symmetries* if for all permutations $\pi \in \text{Sym}_n$ and for all $\mathbf{a}, \mathbf{b} \in \mathcal{S}_\mu^n$ we have $M(\pi(\mathbf{a}), \pi(\mathbf{b})) = M(\mathbf{a}, \mathbf{b})$. Our main theorem shows that walks that respect symmetries and are nearly balanced have fast mixing.

► **Theorem 1** (Singular value bound for nearly balanced walks). *For every $s \geq 2$ and $C < \infty$, there exists $\tau > 0$ such that for every finite set \mathcal{S} of size s and sufficiently large $n \in \mathbb{N}$, the following holds:*

If W is a stochastic matrix over the multislice \mathcal{S}_μ^n that respects symmetries, and satisfies the condition that

$$W(\mathbf{a}, \mathbf{b}) > 0 \quad \Rightarrow \quad \Delta(\mathbf{a}, \mathbf{b}) \text{ is } C\text{-balanced} \quad \forall \mathbf{a}, \mathbf{b} \in \mathcal{S}_\mu^n,$$

then $\sigma_2(W) \leq 1/n^\tau$.

The above result implies that the random walk on the balanced multislice mentioned earlier in this section (which corresponds to $s = 3$ and C -balanced generalized distance parameter with $C = 0$) has its second largest eigenvalue polynomially bounded. In fact, Theorem 1 is more general and covers multislices over any grid \mathcal{S} of constant size (i.e., for every $|\mathcal{S}| = \mathcal{O}(1)$). Additionally, it is robust to perturbations of transition probabilities as long as the transition probabilities are nearly balanced.

¹ Sometimes, this is also called as a “meet table” in algebraic combinatorics.

Indeed Theorem 1 follows from our main technical theorem, stated as Theorem 3 below, which abstracts the main properties that suffice to prove the bound on the second largest singular value. Specifically Theorem 3 shows that, in addition to the symmetries respected by the matrix W , the important features that suffice to prove fast mixing are:

1. The next state of the random walk is “almost $\mathcal{O}(1)$ -wise independent” conditioned on the current state
2. The Frobenius norm of W is polynomially bounded in n .

We elaborate on these conditions below before stating our main technical result Theorem 3. For a distribution D supported on \mathcal{S}^n and set $T \subseteq [n]$, we let D_T denote the marginal distribution supported on \mathcal{S}^T induced by projecting a random variable $x \sim D$ to its coordinates in T . Recall that a distribution D is k -wise independent if for every set $T \subseteq [n]$ with $|T| \leq k$ we have D_T is the uniform distribution on \mathcal{S}^T . Recall further that D is ε -almost k -wise independent if for every set $T \subseteq [n]$ with $|T| \leq k$ we have D_T is ε -close in total variation distance to the uniform distribution on \mathcal{S}^T .

In the following definition we view the rows of a stochastic matrix $M \in \mathbb{R}^{\mathcal{S}_\mu^n \times \mathcal{S}_\mu^n}$, denoted $M(\mathbf{a}) := (M(\mathbf{a}, \mathbf{b}) | \mathbf{b} \in \mathcal{S}_\mu^n)$ for $\mathbf{a} \in \mathcal{S}_\mu^n$, as distributions supported on \mathcal{S}^n (which have zero support outside \mathcal{S}_μ^n).

► **Definition 2** (ε -almost k -wise independent matrix). *For parameter $k \in \mathbb{N}$ and $\varepsilon > 0$ we say that a stochastic matrix $M \in \mathbb{R}^{\mathcal{S}_\mu^n \times \mathcal{S}_\mu^n}$ is ε -almost k -wise independent if for every row $\mathbf{a} \in \mathcal{S}_\mu^n$, the distribution $M(\mathbf{a})$ is ε -almost k -wise independent.*

Finally we recall that for a matrix $M \in \mathbb{R}^{N \times N}$, its Frobenius norm, denoted $\|M\|_F$, is the quantity $\sqrt{\sum_{(i,j) \in N \times N} M(i,j)^2}$.

We now state the main theorem of our work.

► **Theorem 3** (Singular Value Bound for Markov Chains on Balanced Multislice). *For every $\kappa > 0$, and $s \in \mathbb{N}$ with $\kappa \geq s$ there exists $c_1, c_2, c_3 < \infty$ such that for every $\varepsilon > 0$ and every sufficiently large $n \in \mathbb{N}$ that is divisible by s , the following holds: Suppose \mathcal{S} is a set of size s , and $M \in \mathbb{R}^{\mathcal{S}_\mu^n \times \mathcal{S}_\mu^n}$ is a stochastic matrix that satisfies the following three conditions:*

1. *The matrix M respects symmetries.*
2. *$\|M\|_F \leq c_1 \cdot n^\kappa$.*
3. *The matrix M is ε -almost k -wise independent for $k = 10s\kappa$.*

Then we have $\sigma_2(M) \leq \max\{c_2/n, c_3 \cdot \varepsilon\}$.

If the Markov chain is symmetric, then the singular values correspond to the eigenvalues, and hence Theorem 3 yields eigenvalue bounds for symmetric Markov chains satisfying the properties mentioned above. Theorem 3 is proved in [3, Section 3]. The proof involves many standard and some new elements of representation theory for the symmetric group. We elaborate on this in Section 1.2.2. We also note that Theorem 1 immediately follows from Theorem 3, modulo some calculations that verify that Condition (2) above applies to C -balanced matrices. For more details, see [3, Section 3.5].

To illustrate the applicability of Theorem 1, we give two examples, both related to coding theoretic aspects of polynomials and other polynomial-like functions that we refer to as *junta-sums*. These results extend corresponding works in the Boolean setting [1, 2], obtaining natural generalizations to non-Boolean settings.

Distance of polynomials and junta-sums on multislices

A function $f : \mathcal{S}^n \rightarrow G$ is called a d -junta if it depends on only d of the n variables, i.e., there exists a set $I \subseteq [n]$, $|I| \leq d$ and a function $g : \mathcal{S}^I \rightarrow G$ such that for all $\mathbf{a} \in \mathcal{S}^n$, $f(\mathbf{a}) = g(\mathbf{a}|_I)$ where $\mathbf{a}|_I$ is the projection of \mathbf{a} to the coordinates in I . Here we could allow G to be any set, though in this work G will denote an Abelian group. We say f is a *degree d junta-sum* (or simply a *d -junta-sum*) if there exists d -junta's $f_1, \dots, f_k : \mathcal{S}^n \rightarrow G$ such that $f = f_1 + \dots + f_k$.

When $G = \mathbb{F}$ is a field and $\mathcal{S} \subseteq \mathbb{F}$, then degree d junta-sums are closely related to the notion of degree d polynomials. In particular, every degree- d polynomial is also a degree- d junta-sum, and degree- d junta-sums are polynomials of degree at most $(s-1)d$ where $s = |\mathcal{S}|$. Junta-sums come up naturally when studying questions related to testing direct sums and low-degree polynomials [11, 6, 4].

A well-studied question about degree- d polynomials is: How often can a non-zero polynomial be zero on a grid? The well-known and oft-discovered Ore-DeMillo-Lipton-Schwartz-Zippel lemma [18, 9, 23, 20] (henceforth ODLSZ lemma) asserts that a non-zero degree- d polynomial over a field \mathbb{F} is non-zero with probability at least $s^{-d/(s-1)}$ over the uniform distribution over \mathcal{S}^n . When $\mathcal{S} = \mathbb{F}$, the precise bound is $\delta(q, d) = (1 - \beta/q)q^{-\alpha}$, where α and β are the quotient and remainder respectively when d is divided by $q-1$. The former bound immediately implies that a degree- d junta-sum is non-zero with probability at least s^{-d} over \mathcal{S}^n (and the claim even extends to arbitrary \mathcal{S} and Abelian groups G).

A natural related question then becomes – *how do these bounds change when considering natural subsets T that are not grids (or more generally product sets)?* Recent work has begun to address such questions [2, 17]. In this work, we consider the case of the balanced multislice i.e., $T = \mathcal{S}_\mu^n$. Despite the simple nature of these questions, the answer does not seem to have been pinned down before, with the exception of the Boolean case that was resolved recently [2]. We are able to show a clean connection between \mathcal{S}^n and \mathcal{S}_μ^n that allows us to show that these probabilities (in the worst case) differ by at most $\lambda_2(W)$ for some nearly balanced walk over the multislice. This allows us to prove the following theorem, which generalizes the work of [2] beyond the Boolean case.

► **Theorem 4** (Polynomial distance over multislice). *For every finite field $\mathbb{F} = \mathbb{F}_q$, if a degree d polynomial $P(\mathbf{x})$ is such that $P(\mathbf{a}) \neq 0$ for some $\mathbf{a} \in \mathbb{F}_\mu^n$ on the balanced multislice, then*

$$\Pr_{\mathbf{b} \sim \mathbb{F}_\mu^n} [P(\mathbf{b}) \neq 0] \geq \delta(q, d) - \frac{1}{n^{\Omega_q(1)}},$$

where $\delta(q, d) = (1 - \beta/q)q^{-\alpha}$, where α and β are the quotient and remainder respectively when d is divided by $q-1$.

Please refer to [3, Section 4] for the proof.

Note that $\delta(q, d)$ is exactly the distance of the space of degree- d polynomials on the field \mathbb{F}_q and hence the above theorem says that the distance of the space of polynomials on the balanced multislice is nearly exactly what it is in the grid \mathbb{F}_q^n .² An analogous statement can also be made for junta-sums, getting a bound that almost matches the bound over grids, i.e., $1/s^d$ (see [3, Theorem 4.2]).

² An important subtlety here is that there are polynomials that are non-zero in the grid \mathbb{F}_q^n but are zero at all points on the multislice. That is the reason this theorem is only stated for polynomials that are non-zero as functions on the multislice. This is analogous to similar restrictions we place on polynomials in the setting of grids (e.g., in the setting of the Boolean cube, we only consider non-zero *multilinear* polynomials).

Following the proof idea of [2], both cases are handled by a similar proof technique that first proves a quantitatively weak bound on the probability that f is non-zero³(see [3, Corollary 5.11]), and then randomly identifies a small grid inside the multislice. On each such grid, we can apply the ODLSZ lemma as a black-box to assert that if f is non-zero within the randomly identified small grid, then it is non-zero with the “correct” probability (either $\delta(q, d)$ or s^{-d}). It suffices, therefore, to prove that f is non-zero on most of the grids, which is where the main technical theorem regarding the expansion of the walk on the balanced multislice comes into play. We use our eigenvalue bounds along with the quantitatively weak bound already obtained to establish that all but an $n^{-\Omega(1)}$ -fraction of the grids satisfy this property. See Section 1.2.3 for the proof overview and [3, Section 4] for a formal proof.

Local Correction of Junta-Sums

Our next application considers the local correction problem for junta-sums over *grids*. Here a (possibly randomized) corrector is given oracle access to a function f that is known to be δ -close (in normalized Hamming distance) to some degree- d junta-sum P , and also given a point $\mathbf{a} \in \mathcal{S}^n$ and needs to output $P(\mathbf{a})$ (with high probability) while making few oracle queries to f .

In the list-correction setting, the amount of error δ may be too high for P to be defined uniquely by f and δ , but it may be known a priori that the list size is bounded. In the local list-correction problem, the goal for the corrector is to make a few queries to f to produce several “oracle” algorithms, such that for every degree d junta-sum P that is δ -close to f , there is an algorithm with oracle access to f that computes P . We refer the reader to [3, Section 2] for more formal definitions.

Local correction algorithms for low-degree polynomials have played a central role in complexity theory, for example [15, 22]. While most of the early works like [16, 5] considered the setting where $\mathcal{S} = G = \mathbb{F}$, some recent works have considered the setting of $\mathcal{S} = \{0, 1\}$ and general abelian G such as [1] (Note that when $|\mathcal{S}| = 2$, then degree- d polynomials are the same as degree- d junta-sums.)

For general \mathcal{S} and Abelian group G , even the list-decoding radius was not completely understood till this work. We prove that for $\delta = |\mathcal{S}|^{-d} - \epsilon$ there are most $\mathcal{O}_\epsilon(1)$ degree d junta-sums P that are δ -close to any given function f . (This bound is tight in that for $\delta = |\mathcal{S}|^{-d}$ the number of junta-sums grows with n .) This motivates the corresponding local list-correction problem, which we solve tightly in this work. We state an informal version below and point to [3, Theorem 5.1] for the more precise version.

► **Theorem 5** (Local list-correction of junta-sums (Informal)). *For every set \mathcal{S} , every Abelian group G , every integer d and $\epsilon > 0$, there exists an $L = L(\epsilon, d, \mathcal{S})$ such that the following holds.*

There exists an algorithm \mathcal{A} that on oracle access to a function $f : \mathcal{S}^n \rightarrow G$, outputs L oracle algorithms ψ_1, \dots, ψ_L such that for every degree d junta-sum $P : \mathcal{S}^n \rightarrow G$ that is $(1/s^d - \epsilon)$ -close to f , there exists $i \in [L]$ such that $\psi_i^f(\cdot)$ computes P (with high probability for every input).

The query complexity of \mathcal{A} and ψ_1, \dots, ψ_L is $\text{poly}(\log n)$.

³ We prove these bounds by an adaptation of the standard inductive strategy used to prove the standard ODLSZ lemma. Unfortunately, we are unable to use this strategy to prove a tight bound.

This theorem is formalized as [3, Theorem 5.1] with more explicit bounds on the error probability and query complexity, and is proved in [3, Section 5].

This theorem generalizes a theorem of Amireddy, Behera, Paraashar, Srinivasan, and Sudan ([1, Theorem 1.3.4]) who solved the corresponding problem over the Boolean cube $\{0, 1\}^n$. (Note that in the Boolean setting, junta-degree is the same as algebraic degree, and their result is thus expressed in terms of the latter phrase.) Our extension follows the same sequence of steps as employed in [1]. Their work ultimately ends up using the expansion properties of Boolean multislices, which, as we've noted earlier, is well-understood. Extending their work to general grids requires a number of changes that we elaborate on in Section 1.2.4, with the most significant change being the use of Theorem 1 instead of the expansion results on the Boolean slice.

1.2 Techniques and Proof Overview

In this section, we first review known methods for bounding the singular values of walks that respect symmetries and explain where there is a gap in knowledge. We then show how we overcome these challenges by overviewing the proof of our main theorem Theorem 3 in Section 1.2.2. Next, we give an overview of the proof of the ODLSZ theorem for multislices, Theorem 4, in Section 1.2.3. Finally, we discuss the proof of the local correction theorem for grids, Theorem 5 in Section 1.2.4.

1.2.1 Prior Approaches and Obstructions

We describe some prior cases where random walk matrices respecting symmetries (i.e., the first condition of Theorem 3) have been studied and explain the special properties in play there.

Boolean Hypercube and Cayley graphs

A broad class of examples bounding eigenvalues of highly symmetric graphs are the bounds on the eigenvalues of Cayley graphs over abelian groups - this captures random walks on the Boolean hypercube and many more general settings. Here it is well known that the random walk matrix is *diagonalizable*⁴ and the eigenvectors of the random walk matrix depend only on the group (and not the set of generators). This makes it possible to determine the entire eigenspectrum for many basic groups using Fourier analysis. We note that Cayley graphs over some non-abelian groups have been studied, but general results are mostly lacking. In these cases, the random walk matrix is typically not diagonalizable, but can be made block diagonal, using the representation theory of the underlying group. This is a complex tool, and many basic questions are unanswered as we elaborate below.

Boolean slices

One well-studied setting that happens to be the special case of $s = 2$ of our problem is the setting of Boolean slices. Here $\mathcal{S} = \{0, 1\}$ and \mathcal{S}_μ^n is the balanced Boolean slice (all points in $\{0, 1\}^n$ of Hamming weight exactly $n/2$). This setting has particular relevance to the analysis of Boolean functions and combinatorics; see, e.g. [8, 12, 13]. The random walk

⁴ A matrix $M \in \mathbb{R}^{N \times N}$ is said to be diagonalizable if there is a unitary matrix $U \in \mathbb{R}^{N \times N}$ such that UMU^T is a diagonal matrix.

matrices in this setting lie in the *Johnson scheme*, which is an algebra of symmetric matrices that commute with one another. This implies that all such matrices can be diagonalized *simultaneously*, i.e., there exists one unitary matrix U such that for every random walk matrix M on the Boolean slice that respects symmetries, we have that UMU^T is diagonal. This implies that all such matrices M have the same eigenvectors. The works [12, 21] gave explicit descriptions of the common eigenspaces. This can be quite useful when analyzing the spectrum of such matrices and in a recent example [2] used this description to show that a particular random walk matrix on the balanced Boolean slice is a good spectral expander (see [2, Lemma 3.2]).

1.2.2 Spectral Expansion of Multislice Walks

Turning to our setting, our matrix M is not *diagonalizable* and so the techniques from the analysis of Cayley graphs on abelian groups as well as the random walk on the Boolean slice, do not work in this setting. We have to resort to the use of representation theory, but here as we alluded to earlier, the knowledge is not complete. In what follows, we explain what representation theory implies for our setting and how we build on it.

Summary of known facts from representation theory

The fact that our matrix M respects symmetries allows us to invoke results from the representation theory of the symmetric group Sym_n . We cover these results in detail in [3, Proposition 3.1] and [3, Theorem 3.3] (Parts (1) and (2)). Essentially, we can use representation theory to show that our matrix M can be block-diagonalized with relatively few blocks.


Specifically⁵ there is an orthonormal matrix $U = U_{\mathcal{S},n}$ such that for every M respecting symmetries the matrix UMU^T is block diagonal with blocks M_0, M_1, \dots, M_t where t and the “shape” of the blocks is known from standard representation theory. In particular, $M_0 = [1]$ is just a 1×1 matrix that contributes the top singular value (which is 1), and M_1, \dots, M_t determine $\sigma_2(M) < 1$.

Now let us understand the structure of M_i ’s in more detail. Each block M_i is a Kronecker-/tensor product of a “small” matrix $A_i \in \mathbb{R}^{m(i) \times m(i)}$ with a somewhat larger identity matrix i.e.,

$$M_i = A_i \otimes \text{Id}_{k(i)}, \quad \text{where } \text{Id}_k \text{ is the } k \times k \text{ identity matrix.}$$

See Figure 1 for an informal pictorial description. Both the quantities $m(i)$ and $k(i)$ are known from the representation theory of Sym_n (and in particular only depend on \mathcal{S} and n and are independent of the particular matrix M). However, the small matrices A_i ’s do depend on M , and more importantly, the matrix U is not too well-understood. (In particular, we need to understand the effect of U on A_i and this is not clear.) In particular, if we were to arrange i such that $m(i)$ ’s are non-decreasing, then $k(i)$ ’s are also non-decreasing. Intuition from Fourier analysis in the abelian world would suggest that λ_2 comes from $M_1 = A_1 \otimes \text{Id}_{k(1)}$ but, as far as we are aware, even this is *not known*.

⁵ This conclusion only requires that M respects symmetries (see Theorem 3).



■ **Figure 1** An informal visualization for the block diagonalization of our matrix M .

Our analysis

Given that the A_i 's are not determined by only \mathcal{S} and n , and U is not explicitly understood, we need to find some crude ways to bound the singular values of M . We give such an analysis in [3, Section 3] and summarize the essence here. We start with the following informal observation.

► **Observation 6.** *If for some $i \in [t]$, the quantity $k(i)$ is a polynomial larger than the Frobenius norm of M , then every singular value corresponding to the block M_i must necessarily be small.*

The observation follows from the fact that the Frobenius norm is lower bounded by the sum of the squares of the singular values of M_i , each of which repeats at least $k(i)$ times, so each singular value must be small.

In the context of our goal, the observation above immediately allows us to eliminate all the large blocks in the block diagonalization of M and turns the focus to the small blocks – where $k(i)$ is bounded by some polynomial in n . Here, standard facts of representations of Sym_n imply that the corresponding matrix A_i is of constant size (dependent on \mathcal{S} and the exponent of n in $k(i)$ but not on n). However, this does not immediately translate to bounds on the singular values, since these depend on the actual matrix A_i and its entries. Ideally, we would have liked to get our hands on the singular vectors of M corresponding to singular vectors of the A_i 's (after a change of basis according to U), but such vectors were not known (and we do not get such vectors either). Fortunately, a collection of vectors that span the space corresponding to the singular vectors of A_i 's is known. In particular, we use a description given by Dafni, Filmus, Lifshitz, Lindzey, and Vinyals [7] – see [3, Definition 3.8]. Our main contribution adds two observations about this collection of vectors, called “special vectors” below.

Our main contribution here does get something almost as good, for our purposes (i.e., to show that each M_i has top singular value $o_n(1)$):

We observe that the special vectors given in [3, Definition 3.8] are “weakly orthogonal” in the sense that they have $\Omega(1)$ volume (in the space they span). We further observe that these functions are junta-like and so shrink significantly when acted on by ε -almost k -wise independent matrices for sufficiently large constant k . (See [3, Theorem 3.3], Parts 3(c) and 3(d)).

More specifically, the work of [7] yields $\dim(A_i) = m(i)$ many vectors (see [3, Section 3.3]) of M that are supported on coordinates of one of the blocks of M_i after transforming the basis according to U . We show that while these vectors do not form an orthogonal basis, they are sufficiently divergent to ensure their determinantal volume is large (see [3, Lemma 3.12]). Thus, bounding the length of the vectors obtained by applying the linear map A_i by $o_n(1)$ suffices to bound the spectral norm of A_i and hence also M_i . Details of this part may be found in [3, Section 3.3]. We give some more insight in the next paragraph.

To show our singular value bounds, we use the fact that the vectors in the basis correspond to $\mathcal{O}(1)$ -junta's. Specifically, note that a vector that M acts on can be viewed as a function f from \mathcal{S}_μ^n to \mathbb{R} , which can in turn be viewed as a partial function from \mathcal{S}^n to \mathbb{R} . We show that this function depends on only $\mathcal{O}(1)$ -coordinates of the input vector. (See [3, Section 3.3].) This property now combines nicely with our third condition in Theorem 3 which asserts that after multiplication by M any vector f looks essentially random when projected on $\mathcal{O}(1)$ -coordinates and so has little correlation left with f – this immediately translates into an upper bound on the singular value of M corresponding to coordinates in M_i , and yields a proof of Theorem 3.

Why does our Theorem 1 hold only for nearly-balanced walks?

The reason is related to Observation 6. We note that the Frobenius norm condition is not completely natural, and indeed the natural matrices in our applications do not satisfy this condition (and we have to find workarounds). The Frobenius norm restriction is satisfied by nearly-balanced walks as considered in Theorem 1, and indeed it is one of the reasons why Theorem 1 is restricted to such nearly balanced walks.

1.2.3 Distance Lemma over Balanced Multislice

In this subsection, we discuss the proof overview for Theorem 4. The strategy is a generalization of the proof for [2, Lemma 3.2]. The idea is to find a random copy of $\mathcal{S}^{n/s}$ inside the balanced multislice \mathcal{S}_μ^n such that it is a good *sampler* for \mathcal{S}_μ^n , i.e. if we choose points from this subgrid $\mathcal{S}^{n/s}$ at random, then the corresponding points in \mathcal{S}_μ^n behave like random samples. As we explain now, this guarantee essentially allows us to move from balanced multislice to subgrid, where we have a complete understanding of distance. For every non-zero d -junta-sum $P : \mathcal{S}_\mu^n \rightarrow G$, we choose a random copy of $\mathcal{S}^{n/s}$ inside \mathcal{S}_μ^n and restrict P to this copy. With the sampling guarantee, we can argue that the restricted d -junta-sum is also non-zero on the subgrid $\mathcal{S}^{n/s}$, and we get the claimed bound by applying [3, Claim 2.6] on this restricted polynomial. Next, we explain the process of finding a random copy of the n/s -dimensional subgrid inside the balanced multislice.

The key step in our proof is to show that we can find a random copy of $\mathcal{S}^{n/s}$ inside \mathcal{S}_μ^n , which is a sampler for \mathcal{S}_μ^n . We do it by randomly grouping the coordinates x_1, \dots, x_n into n/s buckets of size s and in each bucket, we randomly assign distinct values to the s coordinates. We prove that for two random points in $\mathcal{S}^{n/s}$, their corresponding points in the balanced multislice \mathcal{S}_μ^n are almost pairwise independent. We show this via the second moment method and the expander mixing lemma. We use our main theorem Theorem 3 to show that the random walk on \mathcal{S}_μ^n arising from the above random process has good spectral expansion, making the expander mixing lemma applicable in this context.

1.2.4 Local List Correction for Junta-Sums

Our local list corrector (see [3, Theorem 5.1]) is a generalization of [1, Theorem 1.3.4] to d -junta-sums and arbitrary grids \mathcal{S}^n (instead of degree- d *polynomials* and *Boolean* cube). We do not dwell on the algorithm here, but only highlight and discuss the key technical difference in our work and the previous work of [1]. We request the reader to please refer to [3, Section 5] for more details on the algorithm.

An important step of our local list corrector involves a random restriction of \mathcal{S}^{sk} to a subgrid \mathcal{S}^k , as follows: We randomly group the sk coordinates into k groups of size s , and identify all the s coordinates in a group together by a single new coordinate. To show that our local list corrector has small error probability, we need the following guarantee from the above random restriction: If a d -junta-sum $P \in \mathcal{J}_d(\mathcal{S}^{sk})$ is non-zero on the balanced multislice, then with high probability, it continues to be a non-zero junta-sum on \mathcal{S}^k after the random restriction. For this, we show that the above-mentioned random process can be interpreted as finding a random copy of the balanced multislice in \mathcal{S}^k inside the balanced multislice in \mathcal{S}^{sk} . Similar to the distance lemma for multislices ([3, Theorem 4.2]), we show that we get a good sampler using Theorem 3.

We now briefly touch upon some of the additional challenges in going from the Boolean case of [1] to junta-sums over grids \mathcal{S}^n , in the context of local list-correction. For the local correction algorithm in the unique decoding regime, the main idea is to reduce the problem to the Boolean case but over a biased distribution instead of the uniform one; the proof then proceeds by a mostly straightforward generalization of the local corrector from [1] for the uniform distribution. The overall template for proving the combinatorial bound is also similar to that over the Boolean cube, except now we will need more general anti-concentration lemmas and distance lemmas for junta-sums. As already described in the above paragraph, going from the combinatorial bound for list-decodability to the local list-corrector is the main technical challenge we overcome in this work by making use of the fact that a certain random embedding of the multislice of \mathcal{S}^k inside the multislice of \mathcal{S}^n is a good sampler.

References

- 1 Prashanth Amireddy, Amik Raj Behera, Manaswi Paraashar, Srikanth Srinivasan, and Madhu Sudan. *Low Degree Local Correction Over the Boolean Cube*, pages 5504–5511. Society for Industrial and Applied Mathematics, 2025. doi:10.1137/1.9781611978322.187.
- 2 Prashanth Amireddy, Amik Raj Behera, Srikanth Srinivasan, and Madhu Sudan. A Near-Optimal Polynomial Distance Lemma over Boolean Slices. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, volume 334 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:17, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2025.11.
- 3 Prashanth Amireddy, Amik Raj Behera, Srikanth Srinivasan, and Madhu Sudan. Eigenvalue bounds for symmetric markov chains on multislices with applications. *Electron. Colloquium Comput. Complex.*, TR25-093, 2025. URL: <https://eccc.weizmann.ac.il/report/2025/093>.
- 4 Prashanth Amireddy, Srikanth Srinivasan, and Madhu Sudan. Low-degree testing over grids. In *Approximation, Randomization, and Combinatorial Optimization (RANDOM)*, volume 275, pages 41:1–41:22, 2023. doi:10.4230/LIPIcs.APPROX/RANDOM.2023.41.
- 5 Abhishek Bhowmick and Shachar Lovett. The list decoding radius for Reed-Muller codes over small fields. *IEEE Trans. Inf. Theory*, 64(6):4382–4391, 2018. doi:10.1109/TIT.2018.2822686.

- 6 Andrej Bogdanov and Gautam Prakriya. Direct sum and partitionability testing over general groups. In *International Colloquium on Automata, Languages, and Programming, (ICALP)*, volume 198, pages 33:1–33:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.33.
- 7 Neta Dafni, Yuval Filmus, Noam Lifshitz, Nathan Lindzey, and Marc Vinyals. Complexity Measures on the Symmetric Group and Beyond. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 87:1–87:5, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2021.87.
- 8 Ph. Delsarte. Hahn polynomials, discrete harmonics, and t-designs. *SIAM Journal on Applied Mathematics*, 34(1):157–166, 1978. URL: <http://www.jstor.org/stable/2100864>.
- 9 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.
- 10 Persi Diaconis and Mehrdad Shahshahani. Time to reach stationarity in the bernoulli-laplace diffusion model. *SIAM Journal on Mathematical Analysis*, 18(1):208–218, 1987. doi:10.1137/0518016.
- 11 Irit Dinur and Konstantin Golubev. Direct sum testing: The general case. In *Approximation, Randomization, and Combinatorial Optimization (RANDOM)*, volume 145, pages 40:1–40:11. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.40.
- 12 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.
- 13 Yuval Filmus. Junta threshold for low degree boolean functions on the slice. *The Electronic Journal of Combinatorics*, 30, 2023. doi:10.37236/11115.
- 14 Yuval Filmus, Ryan O'Donnell, and Xinyu Wu. Log-Sobolev inequality for the multislice, with applications. *Electron. J. Probab.*, 27:Paper No. 33, 30, 2022. doi:10.1214/22-ejp749.
- 15 Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *ACM Symposium on Theory of Computing (STOC)*, pages 25–32, 1989. doi:10.1145/73007.73010.
- 16 Parikshit Gopalan, Adam R. Klivans, and David Zuckerman. List-decoding Reed-Muller codes over small fields. In *ACM Symposium on Theory of Computing (STOC)*, pages 265–274, 2008. doi:10.1145/1374376.1374417.
- 17 Swastik Kopparty, Mrinal Kumar, and Harry Sha. High rate multivariate polynomial evaluation codes. *CoRR*, abs/2410.13470, 2024. doi:10.48550/arXiv.2410.13470.
- 18 Øystein Ore. Über höhere kongruenzen. *Norsk Mat. Forenings Skrifter*, 1(7):15, 1922.
- 19 Fabio Scarabotti. Time to reach stationarity in the Bernoulli-Laplace diffusion model with many urns. *Adv. in Appl. Math.*, 18(3):351–371, 1997. doi:10.1006/aama.1996.0514.
- 20 Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980. doi:10.1145/322217.322225.
- 21 Murali K. Srinivasan. Symmetric chains, gelfand-tsetlin chains, and the terwilliger algebra of the binary hamming scheme. *Journal of Algebraic Combinatorics*, 34, 2011. doi:10.1007/s10801-010-0272-2.
- 22 Madhu Sudan, Luca Trevisan, and Salil P. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001. doi:10.1006/JCSS.2000.1730.
- 23 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.