

High Dimensional Expanders: Eigenstripping, Pseudorandomness, and Unique Games

Mitali Bafna* Max Hopkins† Tali Kaufman‡ Shachar Lovett§

November 3, 2021

Abstract

Higher order random walks (HD-walks) on high dimensional expanders (HDX) have seen an incredible amount of study and application since their introduction by Kaufman and Mass (ITCS 2016), yet their broader combinatorial and spectral properties remain poorly understood. We develop a combinatorial characterization of the spectral structure of HD-walks on two-sided local-spectral expanders (Dinur and Kaufman FOCS 2017), which offer a broad generalization of the well-studied Johnson and Grassmann graphs. Our characterization, which shows that the spectra of HD-walks lie tightly concentrated in a few combinatorially structured strips, leads to novel structural theorems such as a tight ℓ_2 -characterization of edge-expansion, as well as to a new understanding of local-to-global graph algorithms on HDX.

Towards the latter, we introduce a novel spectral complexity measure called *Stripped Threshold Rank*, and show how it can replace the (much larger) threshold rank as a parameter controlling the performance of algorithms on structured objects. Combined with a sum-of-squares proof for the former ℓ_2 -characterization, we give a concrete application of this framework to algorithms for unique games on HD-walks, where in many cases we improve the state of the art (Barak, Raghavendra, and Steurer FOCS 2011, and Arora, Barak, and Steurer JACM 2015) from nearly-exponential to polynomial time (e.g. for sparsifications of Johnson graphs or of slices of the q -ary hypercube). Our characterization of expansion also holds an interesting connection to hardness of approximation, where an ℓ_∞ -variant for the Grassmann graphs was recently used to resolve the 2-2 Games Conjecture (Khot, Minzer, and Safra FOCS 2018). We give a reduction from a related ℓ_∞ -variant to our ℓ_2 -characterization, but it loses factors in the regime of interest for hardness where the gap between ℓ_2 and ℓ_∞ structure is large. Nevertheless, our results open the door for further work on the use of HDX in hardness of approximation and their general relation to unique games.

*Department of Computer Science, Harvard University, MA 02138. Email: mitalibafna@g.harvard.edu. Supported in part by the Simons Investigator Award to Madhu Sudan.

†Department of Computer Science and Engineering, UCSD, CA 92092. Email: nmhopkin@eng.ucsd.edu. Supported by NSF Award DGE-1650112.

‡Department of Computer Science, Bar-Ilan University. Email: kaufmant@mit.edu. Supported by ERC and BSF.

§Department of Computer Science and Engineering, UCSD, CA 92092. Email: slovett@cs.ucsd.edu. Supported by NSF Award CCF-1953928.

Contents

1 Introduction	3
1.1 Contributions	3
2 Our Results	5
2.1 Eigenstripping and ST-Rank	6
2.2 Characterizing (non)-Expansion in HD-Walks	7
2.3 Playing Unique Games	9
2.4 Connections to Hardness and the Grassmann Graphs	10
3 Proof Overview	11
3.1 Eigenstripping and the HD-Level-Set Decomposition	11
3.2 Characterizing Expansion	12
3.3 Playing Unique Games	14
4 Related Work	15
5 Preliminaries and Notation	16
5.1 Local-Spectral Expanders and Higher Order Random Walks	16
5.1.1 Two-Sided Local-Spectral Expanders	16
5.1.2 Higher Order Random Walks	17
5.1.3 Expansion of HD-Walks	18
6 Approximate Eigendecompositions and Eigenstripping	18
7 The Spectra of HD-walks	22
8 Pseudorandomness and the HD-Level-Set Decomposition	26
9 Expansion of HD-walks	29
10 Playing Unique Games on HD-Walks	34
10.1 Background for Unique Games and SoS	35
10.2 The Algorithm	36
10.3 Analysis of Algorithm 10.2	38
10.4 Proof of Proposition 10.1	40
11 Acknowledgements	44
A Proof of Lemma 7.1	48
B Orthogonality and the HD-Level-Set Decomposition	50
C Unique Games	51
C.1 Random-walks and Weighted graphs	51
C.2 The Shift-Partition Potential	52
C.3 Sum of Squares and the HD-Level-Set Decomposition	53
C.4 Remaining Proofs from Section 10.3	55
C.5 Remaining Proofs from Section 10.4	55

1 Introduction

Since their introduction by Kaufman and Mass [KM16a] in 2016, higher order random walks (HD-walks) on high dimensional expanders (HDX) have seen an explosion of research and application throughout theoretical computer science, perhaps most famously in approximate sampling [ALOV19, AL20, ALO20, CLV20, CLV21, CGSV21, FGYZ21, JPV21, Liu21, BCC⁺21], but also in CSP-approximation [AJT19], error correction [DHK⁺19, DDHRZ20, JQST20, JST21], and agreement testing [DK17, DD19, KM20]. These breakthroughs, while evidence of the importance of HD-walks, tend to have a fairly narrow focus in their analysis of the walks themselves—most rely only on proving (often one-sided) bounds on spectral expansion. On the other hand, there are many reasons to study combinatorial and spectral structure beyond the second eigenvalue. Such results are often useful, for instance, when designing graph algorithms (e.g. techniques relying on threshold rank [BRS11, GS11, OT13]), and have even shown up in recent breakthroughs in hardness of approximation [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMMS18, KMS18], where such an analysis of the Johnson and Grassmann graphs (themselves HD-walks) proved crucial for the resolution of the 2-2 Games Conjecture. Despite so many interesting connections and the recent flurry of work on HDX, these structures remain relatively unexplored and poorly understood.

Building on work of [KO20, DDFH18], we make progress on this problem, building a combinatorial characterization of the spectral structure of HD-walks on *two-sided local-spectral expanders*, a variant of HDX introduced by Dinur and Kaufman [DK17] to generalize the Johnson and Grassmann schemes¹ (and their corresponding agreement tests). We show how our characterization leads to a new understanding of local-to-global algorithms on HDX, including the introduction of a novel spectral parameter generalizing threshold rank, and give a concrete application to efficient algorithms for Unique Games on HD-walks. While the structural theorems we develop lie in a different regime than the one needed to recover hardness results like the proof of 2-2 Games, we open the door for future work connecting HDX and hardness of approximation.

1.1 Contributions We start with an informal overview of our contributions, broken down into three main sections. We give a more thorough exposition of these results in Section 2 along with some relevant background, and overview their proofs in Section 3. The remainder of the paper is devoted to background, proof details, and further discussion.

Eigenstripping and ST-Rank: We develop the interplay of spectral and combinatorial structure on HD-walks, and introduce a spectral parameter called *Stripped Threshold Rank* (ST-Rank) that generalizes threshold rank and controls the performance of local algorithms on such graphs. In more detail, we prove that the spectrum of any HD-walk can be decomposed into small, disjoint intervals we call “eigenstrips,” where the corresponding eigenvectors in each strip share the same combinatorial structure.

THEOREM 1.1. (EIGENSTRIPPING (INFORMAL)) *Let M be an HD-walk on k -sets of an γ -HDX. Then the spectrum of M lies in $k+1$ eigenstrips of width $O_{k,M}(\gamma^{1/2})$,² each corresponding combinatorially to a level in the underlying HDX.³*

The ST-Rank of an HD-walk (Definition 2.1) then measures the number of spectral strips with eigenvalues above some smallness threshold. This generalizes the concept of standard threshold rank, which measures the *total* number of such eigenvalues, to coarser eigendecompositions. Threshold rank itself is well known to control the performance of many graph algorithms [ABS15, BRS11, GS11, OT13, HST20]. At their core, these techniques often boil down to various methods of enumerating over eigenvectors with large eigenvalues. On structured objects like HD-walks, we argue one should instead enumerate over *strips* of eigenvectors with matching combinatorial structure, and therefore that *ST-rank* is the relevant parameter. Since many important objects share similar combinatorial and spectral characteristics to HD-walks (e.g. noisy hypercube, q -ary hypercube...), we expect ST-Rank to have far-reaching applications beyond those studied in this work.

¹The Johnson Scheme consists of matrices indexed by k -sets of $[n]$ which depend only on intersection size. HD-walks on the complete complex are exactly the non-negative elements of the Johnson Scheme. The Grassmann scheme is an analogous object over vector spaces.

²This was recently improved to $O_k(\gamma)$ by Zhang [Zha20].

³There are $k+1$ rather than k strips since 0 is a dimension of the complex (corresponding to the empty set).

Edge-expansion in HD-walks: The combinatorial and spectral machinery we develop allows us to characterize a fundamental graph property on HD-walks: edge-expansion. Given a graph $G_M = (V, E)$ (corresponding to a random walk M), the edge expansion of a subset $S \subseteq V$ measures the expected probability of leaving S after a single application of M . It is hard to overstate the importance of edge-expansion throughout theoretical computer science. While the most widely used characterizations of expansion are for families where all sets expand (expander graphs) or all *small* sets expand (small set expanders), more involved characterizations for objects such as the Johnson and Grassmann graphs have seen recent use for both hardness of approximation [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMS18] and algorithms [BBK⁺21]. We prove that HD-walks exhibit similar expansion characteristics to the Johnson graphs in two key aspects: the expansion of *locally-structured* sets called links, and the *structure* of non-expanding sets.

Applying our machinery to the former reveals a close connection between local expansion and ST-rank:

THEOREM 1.2. (LOCAL EXPANSION VS. EIGENSTRIPPING (INFORMAL)) *The (non)-expansion of any i -link⁴ is almost exactly the eigenvalue corresponding to the i th level's eigenstrip.*

This connection is particularly useful in the construction and analysis of spectral ‘local-to-global’ graph algorithms, where it helps tie performance guarantees to ST-rank rather than threshold rank. Surprisingly, to our knowledge, Theorem 1.2 is novel even for the Johnson graphs (though in this case it follows easily from known results). In fact, this special case alone already gives a better understanding of recent local-to-global algorithm for unique games [BBK⁺21].

A corollary of this result, on the other hand, extends a very well known fact on the Johnson graphs: *locally structured sets expand poorly*. This raises a natural question: are *all* (small) non-expanding sets explained by local structure? Before answering, we address a subtle point: what exactly does it mean to be “explained” by local structure? There are a number of reasonable ways to formalize this notion, each with its own use. Following [BBK⁺21], we resolve an ℓ_2 -variant of this conjecture stating:

THEOREM 1.3. (CHARACTERIZING EXPANSION IN HD-WALKS (INFORMAL ℓ_2 -VARIANT)) *Any non-expanding set in an HD-walk has **high variance** across low-level links.*

Similar results used in hardness of approximation [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMMS18, KMS18], on the other hand, rely on an ℓ_∞ -variant that replaces variance with *maximum*. While our bound is exactly tight in the former regime, it (necessarily) loses a factor in the latter in cases where there is a significant gap between ℓ_2 and ℓ_∞ structure. Proving a tight bound directly on the ℓ_∞ -variant for HD-walks remains an important open question given their close connections to the structures used in hardness of approximation.

Playing Unique Games: Unique Games are a well-studied class of 2-CSP that underlie a central open question in computational complexity and algorithms called the *Unique Games Conjecture* (UGC). The UGC stipulates that distinguishing between almost satisfiable (value $\geq 1 - \varepsilon$) and highly unsatisfiable (value $\leq \varepsilon$) instances of Unique Games is NP-hard. It implies optimal hardness of approximation for a host of optimization problems such as Max-Cut [KKMO07], Vertex Cover [KR08], and in fact all CSPs [Rag08], but is still not known to be true or false. In a recent breakthrough, Khot, Minzer, and Safra [KMS18] (following a long line of work [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMMS18]) used an ℓ_∞ characterization of expansion on the Grassmann graphs to prove a weaker variant known as the 2-2 Games Conjecture: it is NP-hard to distinguish $(\frac{1}{2} - \varepsilon)$ -satisfiable instances from ε -satisfiable instances of unique games. On the other hand, in contrast to hard CSPs such as 3SAT, a long line of research establishes sub-exponential time approximation algorithms for unique games [ABS15] and furthermore shows that approximating unique games is easy on various restricted classes of graphs such as expanders [MM10], perturbed random graphs [KMM11], certifiable small-set expanders and Johnson graphs [BBK⁺21].

While our ℓ_2 -characterization of expansion may not be as useful for hardness of approximation as the ℓ_∞ -variant, it gets its chance to shine in the latter context: *algorithms*. For most classes of constraint graphs, the best known algorithms for unique games depend on the threshold rank [ABS15, BRS11, GS11] of the graph. Using a sum-of-squares variant of our ℓ_2 -characterization and the connection between local expansion and stripped eigenvalues, we build a local-to-global algorithm for unique games on HD-walks (based off the recent paradigm of [BBK⁺21]) that depends instead on ST-rank.

⁴Links are grouped by level of the complex. The set of links at each level give increasingly finer decompositions of the HDX into local parts. See Section 2 for details.

THEOREM 1.4. (PLAYING UNIQUE GAMES ON HD-WALKS (INFORMAL)) *For any $\varepsilon \in (0, .01)$, there exists an algorithm \mathcal{A} with the following guarantee. If \mathcal{I} is a $(1 - \varepsilon)$ -satisfiable instance of unique games whose constraint graph is an HD-walk M^5 on k -sets of an HDX, then:*

1. \mathcal{A} outputs a $\text{poly}(k^{-r}, \varepsilon)$ -satisfying solution.
2. \mathcal{A} runs in time at most $N^{\tilde{O}((\frac{k}{r})\varepsilon^{-1})}$

where $r = R_{1-O(\varepsilon)}(M)$ is the ST-rank of M above threshold $1 - O(\varepsilon)$, and there are N total k -sets.

When $k = O(1)$, the ST-rank above any threshold τ is substantially smaller than the threshold rank above τ for (non-expanding) HD-walks, and Theorem 1.4 therefore obtains a strict improvement over previous results based on threshold rank for this family. In many cases, such as sparsifications of (constant-level) Johnson graphs or related objects like slices of the noisy q -ary hypercube, our algorithm reduces the best-known runtime from nearly-exponential to polynomial. Besides its independent algorithmic interest, the result carries a number of connections to other open problems both in and outside of the study of unique games. The algorithm sheds some (limited) light,⁶ for instance, on requisite structure for candidate hardness reductions that use direct product testing [KM16b] or potential attempts to use agreement tests based on local-spectral expanders [DK17]. It gives some hope for better unique games algorithms on graphs like the hypercube [AKKT15] as well, which can be viewed as an HD-walk on level $k = O(\log N)$ on a weaker (one-sided) notion of high dimensional expansion than we study. Outside of unique games the result has some further connections to error correcting codes, where approximation algorithms for general CSPs on HDX underlie recent breakthroughs in efficient decoding for locally testable codes [JQST20, JST21]. In fact, algorithms specifically for unique games over expanders have already seen similar use for efficient decoding of direct-product codes based on HDX structure [DHK⁺21].

2 Our Results

We now move to a more in-depth exposition of our main results. First, however, we cover some background regarding local-spectral expanders and higher order random walks. For a full treatment of these and related objects, see Section 5. Local-spectral expansion is a robust notion of local connectivity on **pure simplicial complexes** introduced by Dinur and Kaufman [DK17]. A d -dimensional pure simplicial complex is the downward closure of a d -uniform hypergraph, that is a collection:

$$X = X(0) \cup \dots \cup X(d)$$

where $X(d) \subseteq \binom{[n]}{d}$ is a d -uniform hypergraph, and $X(i)$ consists of all $\tau \in \binom{[n]}{i}$ such that $\tau \subset T$ for some $T \in X(d)$ (and $X(0) = \{\emptyset\}$). We note that this notation is off-by-one from much of the HDX literature, where $X(i)$ is instead given by sets of size $i + 1$, a notation rooted in the topological view of simplicial complexes. The former definition is more common in combinatorial works (e.g. the sampling literature [ALO20]), and more natural for our purposes.

Simplicial complexes come equipped with natural local structure called **links**. For every i and $\tau \in X(i)$, the link of τ is the restriction of the complex to faces containing τ , that is:

$$X_\tau = \{\sigma : \sigma \cap \tau = \emptyset, \sigma \cup \tau \in X\}.$$

We call the link of an i -face an i -link, and when clear from context, will also use X_τ to denote the set of faces at a given level (e.g. in $X(d)$) containing τ . Following Dinur and Kaufman [DK17], we say a complex is a **two-sided γ -local-spectral expander** if the graph underlying every link is a γ -spectral expander (that is all non-trivial eigenvalues are smaller than γ in absolute value). Finally, it is worth noting that we actually study the more general set of **weighted pure simplicial complexes** (see Section 5), but for simplicity restrict to the unweighted case for the moment as it requires essentially no modification.

Higher order random walks [KM16a] are analogues of the walk associated with standard graphs (given by the normalized adjacency matrix) that moves from vertex to vertex via an edge. In a higher order random walk,

⁵Note that self-adjoint random walks can be equivalently viewed as undirected weighted graphs, and therefore also as underlying constraint graphs for unique games.

⁶The result certainly does not rule out reductions using such structure, but states that one must be careful in doing so that the resulting constraint graphs fall outside the parameter regime for our algorithm.

one applies a similar process at any level of the complex—moving for instance between two edges via a triangle, or two triangles via a pyramid. We call walks between k -faces of a complex **k -dimensional HD-walks** (see Definition 5.7 for formal definition), and study a broad set of walks that capture important structures such as sparsifications of the Johnson and Grassmann schemes.⁷ For simplicity, throughout the introduction we will often focus on two natural classes of HD-walks which see the most use in the literature: the **canonical walks** N_k^i which walk between k -faces via a neighboring $(k+i)$ -face, and the **partial-swap walks** S_k^i which do the same but only move between k -faces of fixed intersection size $k-i$. While this latter class may seem less natural at first, notice that on the complete complex they are exactly the well-studied Johnson graph $J(n, k, k-i)$ (the graph on $\binom{[n]}{k}$ where $(v, w) \in E$ iff $|v \cap w| = k-i$).

2.1 Eigenstripping and ST-Rank With background out of the way, we start our results in earnest with a more in-depth discussion of eigenstripping and ST-rank: the spectral and combinatorial structure of HD-walks. We prove that the spectra of k -dimensional HD-walks lie in $k+1$ tightly concentrated “eigenstrips”, where the i th strip corresponds combinatorially to functions lifted from the i th level of the complex by averaging.

THEOREM 2.1. (SPECTRUM OF HD-WALKS (INFORMAL THEOREM 6.1 + COROLLARY 7.1 + PROPOSITION 7.3))
Let M be an HD-walk on the k th level of a two-sided γ -local-spectral expander. Then the spectra of M is highly concentrated in $k+1$ strips:

$$\text{Spec}(M) \in \{1\} \cup \bigcup_{j=1}^k [\lambda_i(M) - e, \lambda_i(M) + e]$$

where the $\lambda_i(M)$ are decreasing constants depending only on M , and the error term satisfies $e \leq O_{k,M}(\sqrt{\gamma})$. Moreover, the span of eigenvectors with eigenvalues in the strip $\lambda_i(M) \pm e$ are (approximately) functions in $X(i)$ lifted to $X(k)$ by averaging.

Recently, Zhang [Zha20] provided a quantitative improvement of Theorem 2.1 such that $e \leq O_k(\gamma)$. Theorem 2.1 can be seen as a marriage (and generalization) of previous results studying N_k^1 of Kaufman and Oppenheim [KO20], and Dikstein, Dinur, Filmus, and Harsha (DDFH) [DDFH18]. The former prove that N_k^1 exhibits spectral eigenstripping, while the latter introduce the corresponding combinatorial structure but lack the machinery to tie it directly to these strips. We fill in the gap by proving a general linear algebraic theorem of independent interest.

THEOREM 2.2. (APPROXIMATE EIGENDECOMPOSITIONS IMPLY EIGENSTRIPPING (INFORMAL THEOREM 6.1))
Let M be a self-adjoint operator over an inner product space V , and $V = V^1 \oplus \dots \oplus V^k$ a decomposition satisfying $\forall 1 \leq i \leq k, f_i \in V^i$:

$$\|Mf_i - \lambda_i f_i\| \leq c_i \|f_i\|$$

for some family of constants $(\{\lambda_i\}_{i=1}^k, \{c_i\}_{i=1}^k)$. Then as long as the c_i are sufficiently small, the spectra of M is concentrated around each λ_i :

$$\text{Spec}(M) \subseteq \bigcup_{i=1}^k [\lambda_i - e, \lambda_i + e] = I_{\lambda_i},$$

where $e = O_{k,\lambda}(\sqrt{\max_i \{c_i\}})$. This was recently improved to $e \leq O_k(\max_i \{c_i\})$ by Zhang [Zha20].

Theorem 2.1 then follows from proving that DDFH’s decomposition is an approximate eigendecomposition for all HD-walks (DDFH only show this holds for N_k^1).

In the analysis of graphs, it is often useful to bound the number of eigenvalues above some smallness threshold. This parameter, called the **threshold rank**, often plays a key role, for instance, in graph algorithms, including many methods for generic 2CSP approximation [ABS15, BRS11, GS11, OT13, HST20]. Painted in broad strokes, these algorithms usually boil down to some method of eigenvalue enumeration, whether performed directly or implicitly in the proof of key structural lemmas. Given the spectral structure of HD-walks (or generally any operator with a natural approximate eigendecomposition), we argue that it is generally more natural to enumerate over *strips* with large eigenvalues instead. This motivates a natural spectral complexity measure we call **stripped threshold rank**.

⁷We note that the Grassmann scheme comes from applying HD-walks to the Grassmann poset, which is not a simplicial complex.

DEFINITION 2.1. (STRIPPED THRESHOLD RANK) *Let M be a linear operator over a vector space V with decomposition $V = \bigoplus_i V^i$ denoted \mathcal{D} , where each V^i is the span of some set of eigenvectors. Given $\delta \in \mathbb{R}$, the stripped threshold rank (ST-Rank) with respect to δ and \mathcal{D} is:*

$$R_\delta(M, \mathcal{D}) = |\{V^i \in \mathcal{D} : \exists f \in V^i, Mf = \lambda f, \lambda > \delta\}|.$$

In this work, \mathcal{D} will always correspond to the eigenstrips given by Theorem 2.1 so we drop it from the notation.

The ST-Rank of any (non-expanding) k -dimensional HD-walk is always substantially smaller than its standard threshold rank (which is at least $\text{poly}(n)$, and often as large as $n^{\Omega(k)}$). As a general algorithmic paradigm, if one can leverage the combinatorial structure of eigenstrips to replace eigenvalue enumeration, it is possible to achieve substantially better performance on structured graphs. In Section 2.3, we discuss a concrete instantiation of this framework to unique games, where we gain substantial improvements over state of the art algorithms [BRS11, ABS15] over such constraint graphs.

Before moving on to such results, however, it is natural to ask whether we can give a finer-grained characterization of ST-rank for HD-walks. In Section 7 we show how to explicitly compute the approximate eigenvalue $\lambda_i(M)$ corresponding to each strip based upon the structure of M . While the full result requires further background, its specification to canonical and partial-swap walks has a nice combinatorial interpretation. Recall that the canonical walk N_k^i walks between k -faces through a shared $(k+i)$ -face. We say it has **depth** $i/(k+i)$ since it traverses i of the $k+i$ relevant levels of the complex. Similarly, we say the partial-swap walk S_k^i has depth i/k , since it swaps i out of k elements (up to factors in γ , this can also be viewed as walking between k faces via a shared $k-i$ face, where one traverses i out of the k relevant levels). We prove that the stripped eigenvalues corresponding to the canonical and partial-swap walks decay exponentially fast with a base rate dependent on depth.

THEOREM 2.3. (ST-RANK OF HD-WALKS (COROLLARY 7.2 + COROLLARY 7.3)) *Let M be a canonical or partial-swap walk of depth $0 \leq \beta \leq 1$ on a sufficiently strong two-sided local-spectral expander. Then the eigenvalues corresponding to the eigenstrips of M decay exponentially fast:*

$$\lambda_i(M) \leq e^{-\beta i}.$$

The ST-Rank $R_\delta(M)$ is then at most:

$$R_\delta(M) \leq \frac{\ln(\frac{1}{\delta})}{\beta}$$

In other words, walks that reach deep into the complex (e.g. $N_k^{k/2}$) have constant ST-Rank, whereas shallow walks like N_k^1 have ST-rank $\Omega(k)$.

2.2 Characterizing (non)-Expansion in HD-Walks We now move to our main structural application, characterizing edge expansion in HD-walks. Edge expansion is a fundamental combinatorial property of graphs with applications across many areas of theoretical computer science, including (as we will soon discuss) both hardness and algorithms for unique games. For a subset of vertices $S \subseteq V$, edge expansion measures the (normalized) fraction of edges which leave S .

DEFINITION 2.2. (EDGE EXPANSION) *Given a graph $G(V, E)$, the edge expansion of a subset $S \subset V$ is:*

$$\phi(G, S) = \frac{E(S, V \setminus S)}{E(S, V)}$$

where $E(S, T)$ counts the number of edges crossing from S to T (double-counting edges in the intersection). When convenient, we denote non-expansion, $1 - \phi(G, S)$, as $\bar{\phi}(G, S)$, and drop G from the notation when clear from context.

It is often useful to characterize exactly which sets in a graph expand. The two most widely used characterizations are when all⁸ sets expand (expanders), or when all *small* sets expand (small-set expanders). However, many

⁸By this we really mean all sets of size at most $|V|/2$.

important structures fall outside such a simple characterization. The Johnson and Grassmann graphs, for instance, are well known to have small non-expanding sets. Characterizing the structure of non-expansion in these graphs was crucial not only for the resolution of the 2-2 Games Conjecture [KMS18], but also for recent algorithms for unique games [BBK⁺21].

Using the spectral machinery developed in the previous section, we give a tight characterization of expansion in HD-walks. As discussed, we focus mainly on two key aspects: the expansion of *links*, and the structure of non-expanding sets. We start with the former, where links have long been the prototypical example of small, non-expanding sets for the Johnson graphs. In fact, we prove this stems from a much stronger connection between local expansion and spectral structure in HD-walks.

THEOREM 2.4. (LOCAL (NON)-EXPANSION = GLOBAL SPECTRUM (INFORMAL THEOREM 9.1)) *Let M be a k -dimensional HD-walk on a sufficiently strong d -dimensional two-sided γ -local-spectral expander with $d > k$,⁹ and let $\lambda_i(M)$ be approximate eigenvalues corresponding to M 's $k+1$ eigenstrips. For all $0 \leq i \leq k$ and $\tau \in X(i)$, let X_τ denote the set of k -faces that contain τ . Then X_τ is expanding if and only if $\lambda_i(M)$ is small:*

$$\bar{\phi}(X_\tau) \in \lambda_i(M) \pm O_k(\gamma).$$

Since HD-walks (like the Johnson graphs) are generally poor spectral expanders, Theorem 2.4 implies the existence of small, non-expanding sets (namely 1-links¹⁰). Our stronger characterization of local expansion is of independent interest beyond this simple corollary, however, due to its important connections to local-to-global algorithms and ST-rank. In particular, essentially all work on high-dimensional expanders relies on a strategy known as the *local-to-global paradigm*, where a global property (e.g. mixing, agreement testing, etc.) is reduced to examining a local version on links. These arguments often reduce to showing that, in some relevant sense, interaction *between* links is minimal. Theorem 2.4 offers exactly such a statement for levels of the complex which correspond to eigenspaces with large eigenvalues—since links at these levels are non-expanding, they don't have much interaction. This gives a holistic local-to-global approach for spectral graph algorithms on HDX, since the eigenspaces corresponding to *small* eigenvalues are traditionally easy to handle through other means. Later we will see an algorithmic application of this approach to unique games, where the non-expansion of links allows us to patch together local solutions, and the connection to global spectra ties performance of the algorithm to ST-rank.

Returning to the topic of expansion, Theorem 2.4 shows that as long as M is not a spectral expander, there always exist small, non-expanding sets in the form of links. This raises a natural question: are *all* non-expanding sets explained by links? Before answering, let's spend a little time formalizing this. Given a subset $S \subset X(k)$, let $L_{S,i}$ be a function on i -faces measuring the deviation of S from its average across links at level i :

$$\forall \tau \in X(i) : L_{S,i}(\tau) = \mathbb{E}_{X_\tau}[\mathbb{1}_S] - \mathbb{E}[\mathbb{1}_S].$$

The statement “non-expansion is explained by links” then really amounts to saying that we expect $L_{S,i}$ to be far from 0 in some sense if S is non-expanding. There are a number of reasonable ways to formalize this notion, each with its own use. Following [BBK⁺21], we mainly focus on an ℓ_2 -variant that is particularly useful for designing local-to-global algorithms: if S is non-expanding, then $\|L_{S,i}\|_2^2$ (or equivalently the *variance* of S across i -links) must be large. It is perhaps easier to think about this result in terms of its contrapositive. Call a set **ℓ_2 -pseudorandom** if its variance across i -links is *small*.

DEFINITION 2.3. (ℓ_2 -PSEUDORANDOM SETS) *Let X be a pure simplicial complex. We call a subset $S \subset X(k)$ (ε, ℓ) - ℓ_2 -pseudorandom if*

$$\forall i \leq \ell : \|L_{S,i}\|_2^2 \leq \varepsilon \mathbb{E}[\mathbb{1}_S]$$

We remark briefly on the choice of normalization by $\mathbb{E}[\mathbb{1}_S]$ on the right. While mostly a formality, we will see this is in fact a natural choice both in comparison to the ℓ_∞ -variant (which differs in some natural sense by at least a factor of $\mathbb{E}[\mathbb{1}_S]$), and when considering expansion which is itself normalized by $\mathbb{E}[\mathbb{1}_S]$. Indeed with this in mind, we prove that (ε, ℓ) - ℓ_2 -pseudorandom sets expand near-perfectly.

⁹The lower bound on non-expansion still holds for $d = k$, which is sufficient for most applications of this result. It's also worth noting that the condition $d > k$ itself is generally trivially satisfied, since all known two-sided local-spectral expanders are in fact cutoffs of larger complexes.

¹⁰Since we generally consider the regime where $|X(1)| \gg k$, a basic averaging argument gives that there must exist a $o(1)$ -size 1-link. One may also note that in a γ -local-spectral expander, no 1-link can have more than $O_k(\gamma)$ mass, so all links are small if $\gamma \leq o(1)$.

THEOREM 2.5. (ℓ_2 -PSEUDORANDOM SETS EXPAND (INFORMAL THEOREM 9.2)) *Given a k -dimensional HD-walk M on a sufficiently strong two-sided γ -local-spectral expander and small constants $\alpha, \delta, \varepsilon > 0$, we have that any $(\varepsilon, R_\delta(M))$ -pseudorandom set S of density α expands near-perfectly:*

$$\phi(S) \geq 1 - \alpha - \delta - O_k(\varepsilon) - O_k(\gamma)$$

We note that (the formal version of) this result is exactly tight (in the limit of $\gamma \rightarrow 0$). Passing back to the original (now contrapositive) form of the statement, Theorem 2.5 immediately implies that any non-expanding set must have large variance across links at a level determined by the ST-rank of M . On the algorithmic side, this informally translates to the statement that most interesting structure lies on low-level links, which is crucial to any local-to-global algorithm.

Before formalizing this algorithmic intuition in the case of unique games, it's worth taking a moment to consider the implication of Theorem 2.5 to a different formalization of this problem with recent applications to hardness of approximation: the ℓ_∞ -variant. In this characterization, the (squared) ℓ_2 -norm of $L_{S,i}$ is replaced with its maximum. In other words, the ℓ_∞ -variant characterization posits that every non-expanding set must be highly concentrated in some *individual* link. We prove that Theorem 2.5 actually holds in this regime as well by a simple reduction. If we analogously define ℓ_∞ -pseudorandom sets (see Definition 8.2), it is not hard to show that any ℓ_∞ -pseudorandom set must also be ℓ_2 -pseudorandom. This results in a tight version of Theorem 2.5 for the ℓ_∞ -regime, but only when $\max(L_{S,i})$ is close to $\frac{1}{\mathbb{E}[\mathbb{1}_S]} \|L_{S,i}\|^2$. Practically, the interesting regime in hardness of approximation is when these two quantities are far apart. In this case our ℓ_∞ to ℓ_2 reduction necessarily loses important factors, so we cannot recover any results from the hardness of approximation literature.

Nevertheless, it is worth stating that as an immediate corollary of our ℓ_2 -based analysis we get an ℓ_∞ -characterization of expansion in HD-walks: any non-expanding set must be non-trivially concentrated in a link.

COROLLARY 2.1. (NON-EXPANDING SETS CORRELATE WITH LINKS (INFORMAL COROLLARY 9.1)) *Let M be a k -dimensional HD-walk on a sufficiently strong two-sided γ -local-spectral expander. Then if $S \subset X(k)$ is a set of density α and expansion at most:*

$$\phi(S) < 1 - \alpha - \delta - O_k(\gamma)$$

for some $\delta > 0$, S must be non-trivially correlated with some i -link for $1 \leq i \leq r = R_{\delta/2}(M)$:

$$\exists 1 \leq i \leq r, \tau \in X(i) : \mathbb{E}_{X_\tau}[\mathbb{1}_S] \geq \alpha + \frac{c_{\delta,r}}{\binom{k}{i}}$$

where $c_{\delta,r}$ depends only on δ and the ST-rank r .

Proving a k -independent ℓ_∞ -characterization for HD-walks, either directly or via a stronger k -dependent reduction to the ℓ_2 -regime, remains an interesting open problem. Such a characterization is known on the Johnson graphs (and a number of related objects [KLLM19, FKLM20]), and may shed light on similar structure in the Grassmann graphs used to prove the 2-2 Games Conjecture [KMS18]. We discuss some additional subtleties in this direction in Section 2.4.

2.3 Playing Unique Games One motivation for studying spectral structure and non-expansion in HD-walks stems from a simple class of 2-CSPs known as unique games, a central object of study in hardness-of-approximation since Khot's introduction of the Unique Games Conjecture (UGC) [Kho02] nearly 20 years ago. We study affine unique games which are known to be as hard as unique games [KKMO07]:

DEFINITION 2.4. (AFFINE UNIQUE GAMES) *An instance $I = (G, \mathcal{S})$ of affine unique games over alphabet $\Sigma = \{0, \dots, m-1\}$ is a weighted undirected graph $G(V, E)$, and set of affine shifts $\mathcal{S} = \{s_{uv} \in \Sigma\}_{(u,v) \in E}$. The value of I , $\text{val}(I)$, is the maximum fraction of satisfied constraints over all possible assignments Σ^V :*

$$\max_{X \in \Sigma^V} \Pr_{(u,v) \sim E} [X_u - X_v = s_{uv} \pmod{|\Sigma|}],$$

where edges are drawn corresponding to their weight. For an individual assignment X , we refer to this expectation as $\text{val}_I(X)$. Any weighted undirected graph G is uniquely associated with a self-adjoint random-walk matrix M where $M(u, v) = \Pr_E[(u, v)] / \sum_w \Pr_E[(u, w)]$ (see Appendix C.1). We will usually view the constraint graph in this manner instead, referring to unique games instances over random-walks as $I = (M, \Pi)$.

Informally, the UGC states that for sufficiently small constants ε, δ , there exists an alphabet size such that distinguishing between instances of unique games with value $1 - \varepsilon$ and δ is NP-hard. A positive resolution to the UGC would resolve the hardness-of-approximation of many important combinatorial optimization problems, including CSPs [Rag08], vertex-cover [KR08], and a host of others such as [Kho02, KKMO07, GMR08, CJHJ11, KTW14, KV15].

Building on the recent framework of Bafna, Barak, Kothari, Schramm, and Steurer [BBK⁺21], we show how our structural theorems combine with the notion of ST-rank to give efficient algorithms for unique games over HD-walks.

THEOREM 2.6. (PLAYING UNIQUE GAMES ON HD-WALKS (INFORMAL THEOREM 10.1)) *For any $\varepsilon \in (0, .01)$, there exists an algorithm \mathcal{A} with the following guarantee. If $\mathcal{I} = (M, \mathcal{S})$ is an instance of affine unique games with value at least $1 - \varepsilon$ over M , a complete k -dimensional HD-walk on a d -dimensional two-sided γ -local-spectral expander with $\gamma \ll o_k(1)$ and $d > k$ ¹¹ then \mathcal{A} outputs a $\text{poly}(\tau)$ -satisfying assignment in time $|X(k)|^{\text{poly}(1/\tau)}$, where $\tau = \left(\frac{\varepsilon}{(r(\varepsilon))^k}\right)$ and $r(\varepsilon) = R_{1-O(\varepsilon)}(M)$ is the ST-rank of M .*

It is worth giving a quantitative comparison of this result to the best known algorithms based on threshold rank [ABS15, BRS11, GS11]. These algorithms run in time roughly exponential in the $(1 - O(\varepsilon))$ -threshold rank of G to get similar soundness guarantees when $k = O(1)$ ¹². Due to the poor threshold rank of HD-walks, this generally amounts to nearly exponential time $(2^{|X(k)|^{\text{poly}(\varepsilon)}})$, whereas our algorithm runs in time roughly $|X(k)|^{2^{O(k)}}$. In the regime of constant k , Theorem 2.6 therefore gives a polynomial time algorithm for such instances achieving an exponential improvement over threshold-rank based algorithms. We describe a few examples of such families below.

Since Theorem 2.6 can be a bit hard to interpret without additional knowledge of the high dimensional expansion literature, we end our discussion of algorithms for unique games with a few concrete examples. Perhaps the most basic example generalizing the Johnson graphs that fits into our framework are slices of the q -ary noisy hypercube. For constant-level slices, we improve over BRS from nearly exponential to polynomial. Further, the result is robust in the sense that it continues to hold even when the underlying complex is perturbed. This results in algorithms, for instance, for dense random sparsifications of these slices, or slices taken from a negatively correlated distribution over \mathbb{Z}_q^n (or more generally slices of any spectrally-independent spin-system [ATO20]). Another interesting class of graphs we see significant speed-ups on are algebraic sparsifications of the Johnson graphs stemming from bounded-degree constructions of two-sided local-spectral expanders such as (cutoffs of) Lubotsky, Samuels, and Vishne's [LSV05] Ramanujan complexes, or Kaufman and Oppenheim's [KO18] coset complex expanders. The resulting speedups on this class of graphs is somewhat more surprising than the above, since they exhibit substantially different structure in other aspects (they are, for instance, of bounded degree unlike the Johnson graphs).

2.4 Connections to Hardness and the Grassmann Graphs We finish the discussion of our results by taking a deeper look at connections with recent progress on the UGC, and argue that our framework opens an avenue for further progress. The resolution of the 2-2 Games Conjecture [KMS18] hinged on a characterization of non-expanding sets on the Grassmann graph not dissimilar to what we have shown for two-sided local-spectral expanders. While we have focused above on HDX which are simplicial complexes, our work extends to a broader set of objects introduced by DDFH [DDFH18] called expanding posets which includes class of objects includes expanding subsets of the Grassmann poset we call q -posets (we refer the reader to [DDFH18] for definitions and further discussion).

Theorem 2.5 and Corollary 2.1 extend naturally to HD-walks on q -posets. We state the latter result here since it follows without too much difficulty from the same arguments as in this paper, but the full details (and further generalizations to expanding posets) will appear in a companion paper.

¹¹Formally, we note γ also has dependence on M (see Theorem 10.1), though this can be removed in special cases like the Johnson scheme. We also note again that $d > k$ is generally a trivial condition for strong two-sided local-spectral expanders.

¹²This corresponds to the setting in which the dimension of the HD-walk is fixed, and the number of vertices in the complex grows.

COROLLARY 2.2. (NON-EXPANSION IN q -EPOSET) *Let (X, Π) be a two-sided γ - q -eposet with γ sufficiently small, M a k -dimensional HD-walk on (X, Π) . Then if $S \subset X(k)$ is a set of density α and expansion:*

$$\phi(S) < 1 - \alpha - O_{q,k}(\gamma) - \delta$$

for some $\delta > 0$ and $r = R_{\delta/2}(M)$, S must be non-trivially correlated with some i -link for $1 \leq i \leq r$:

$$\exists 1 \leq i \leq r, \tau \in X(i) : \mathbb{E}_{X_\tau} [\mathbb{1}_S] \geq \alpha + \frac{c_{\delta,r}}{\binom{k}{i}_q}$$

where $c_{\delta,r}$ depends only on δ and the ST-rank r , and $\binom{k}{i}_q$ is the standard q -binomial coefficient:

$$\binom{k}{i}_q = \prod_{j=0}^{i-1} \frac{q^{k-j} - 1}{q^{i-j} - 1},$$

Since the Grassmann graphs are simply partial-swap walks on the Grassmann poset.¹³ Corollary 2.2 provides a direct connection to the proof of the 2-2 Games Conjecture [KMS18]. Unfortunately, due to the dependence on k , this result is too weak to recover the proof.

The dependence of both Corollary 2.1 and Corollary 2.2 on k is a subtle but important point, so we'll finish the section by discussing it in a bit more detail. The particular dependence we get in the ℓ_2 -regime, $\binom{k}{i}$ for the Johnson and $\binom{k}{i}_q$ for the Grassmann, is tight and can be understood by examining the variance of i -links. While i -links are the prototypical example of an $\Omega(1)$ -pseudorandom set in the ℓ_∞ regime, one can show that they are actually about $1/\binom{k}{i}$ -pseudorandom in the ℓ_2 -regime (or $1/\binom{k}{i}_q$ for the Grassmann). Since links are non-expanding, any bound in the ℓ_2 -regime must have matching dependence on k to make up for this fact. Indeed, one can use the same argument to show that a k -independent bound cannot exist in ℓ_2 -regime (even if we relax dependence on pseudorandomness, see Proposition 9.2). As a result, any ℓ_∞ to ℓ_2 reduction like ours (which has no dependence on k) will always result in a final bound depending on k .

Finally, it's worth noting that while the ℓ_∞ -regime escapes this particular issue since links are $\Omega(1)$ -pseudorandom (and indeed that for the Johnson graphs, a k -independent bound is known [KMMMS18]), there is an additional consideration for the Grassmann graphs: there exist small, non-expanding ℓ_∞ -pseudorandom sets [DKK⁺18a]. The proof of the 2-2 Games Conjecture therefore relies on a finer-grained definition of local structure than links (called zoom-in zoom-outs) [KMS18]. While we cannot hope to apply exactly the same techniques to analyze this variant, we view our method's generality and simplicity as evidence that a deeper understanding of higher order random walks may be key to further progress on the UGC.

3 Proof Overview

In this section we give a proof overview of our main results. Throughout the section we will assume the complex is endowed with a uniform distribution, and will (usually) ignore error terms in the spectral parameter γ . The full details for weighted complexes and a careful treatment of the error terms is given in the main body along with a number of further generalizations.

3.1 Eigenstripping and the HD-Level-Set Decomposition We start with a discussion of the techniques underlying Theorem 2.1 (Eigenstripping). At its core, this (and indeed all of) our results rely upon now-standard machinery for working on simplicial complexes called the averaging operators.

DEFINITION 3.1. (THE AVERAGING OPERATORS) *Let X be a d -dimensional pure simplicial complex. For any $0 \leq k \leq d$, denote the space of functions $f : X(k) \rightarrow \mathbb{R}$ by C_k . The “Up” operator lifts $f \in C_k$ to $U_k f \in C_{k+1}$:*

$$\forall y \in X(k+1) : U_k f(y) = \frac{1}{k+1} \sum_{x \in X(k) : x \subset y} f(x).$$

¹³Seeing that they are HD-walks is non-trivial, and follows from the q -analog of work in [AJT19].

The “Down” operator lowers functions $f \in C_{k+1}$ to $D_{k+1}f \in C_k$:

$$\forall x \in X(k) : D_{k+1}f(x) = \frac{1}{n-k} \sum_{y \in X(k+1) : y \supset x} f(y).$$

Since it is often useful to compose these operators, we will use the shorthand $U_i^k = U_{k-1} \dots U_i$ and $D_i^k = D_{i+1} \dots D_k$ to denote the composed averaging operators which raise and lower functions between C_i and C_k by averaging.

In fact, the averaging operators are crucial even to defining higher order random walks. We discuss this definition in greater detail in Section 5, and for now settle for noting that a basic class of HD-walks are given simply by composing U and D . For instance, the canonical walk N_k^i is the composition $D_k^{k+i} U_k^{k+i}$, and the partial swap walks S_k^i turn out to be an affine combinations of the N_k^i [AJT19].

The proof of Theorem 2.1 relies on a useful decomposition for functions on simplicial complexes based upon the averaging operators we call the *HD-level-set Decomposition* recently introduced by Dikstein, Dinur, Filmus, and Harsha [DDFH18]. The idea is to break functions into components coming from each level of the complex, lifted by averaging to the top level.

THEOREM 3.1. (HD-LEVEL-SET DECOMPOSITION, THEOREM 8.2 [DDFH18]) *Let X be a d -dimensional two-sided γ -local-spectral expander, $\gamma < \frac{1}{d}$, $0 \leq k \leq d$, and let:*

$$H^0 = C_0, H^i = \text{Ker}(D_i), V_k^i = U_i^k H^i.$$

Then:

$$C_k = V_k^0 \oplus \dots \oplus V_k^k.$$

In other words, every $f \in C_k$ has a unique decomposition $f = f_0 + \dots + f_k$ such that $f_i = U_i^k g_i$ for $g_i \in \text{Ker}(D_i)$.

The HD-Level-Set Decomposition is particularly useful not only for its rigid combinatorial structure, but also its spectral properties. Namely, one familiar with the Johnson Scheme might notice that when X is the complete complex, this decomposition exactly gives the eigenspaces of the Johnson Scheme. DDFH gave an approximate extension of this result to the basic “upper” walk $D_{k+1}U_k$, proving that the V_k^i are approximate eigenspaces.

The proof of Theorem 2.1 follows from combining an extension of this result to all HD-walks (see Corollary 7.1) with Theorem 2.2 which states that any approximate eigendecomposition strictly controls the spectrum of the underlying operator. However, since the proof of Theorem 2.2 is purely linear algebraic (and is essentially irrelevant for understanding our core results on expansion and unique games), we defer detailed discussion of the result to Section 6.

3.2 Characterizing Expansion We now take a look at how this combinatorial understanding of the spectra of higher order random walks allows us to characterize their edge-expansion. We start with a standard observation, the expansion of a set can be written as the inner product of its indicator function. In other words, if X is a simplicial complex and M is an HD-walk on $X(k)$, the expansion of any $S \subset X(k)$ with respect to M may be written as:

$$\phi(M, \mathbb{1}_S) = 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \langle \mathbb{1}_S, M \mathbb{1}_S \rangle.$$

The key is then to notice that by the bilinearity of inner products, we can expand the right-hand side in terms of the HD-Level-Set Decomposition. In particular, writing $\mathbb{1}_S = \sum_{i=0}^k \mathbb{1}_{S,i}$ for $\mathbb{1}_{S,i} \in V_k^i$, we have:

$$\phi(M, \mathbb{1}_S) = 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \langle \mathbb{1}_S, M \mathbb{1}_{S,i} \rangle.$$

Finally, since we know each V_k^i is approximately an eigenstrip concentrated around some λ_i , up to additive error in γ we can simplify this to:

$$(3.1) \quad \phi(M, \mathbb{1}_S) \approx 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \lambda_i \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle.$$

Thus we see that understanding the expansion of S comes down to the interplay between its projection onto each level of the complex and their corresponding approximate eigenvalues. Our characterization of expansion combines this observation with the combinatorial and spectral structure of the HD-Level-Set Decomposition. We devote the rest of the section to sketching the proofs of Theorem 2.4 (local expansion vs global spectra), Theorem 2.5 (pseudorandom sets expand), and our reduction from the ℓ_∞ -variant to ℓ_2 -variant.

Proof sketch of Theorem 2.4: Recall that Theorem 2.4 shows a tight inverse relation between the expansion of links and the spectra of M . Given $\tau \in X(j)$, we wish to examine the indicator function $\mathbb{1}_{X_\tau}$ of the link X_τ . First, notice that since $\mathbb{1}_{X_\tau} = \binom{k}{j} U_j^k \mathbb{1}_\tau$, it's easy to see that $\mathbb{1}_{X_\tau} \in V_k^0 \oplus \dots \oplus V_k^j$. We prove something stronger, that $\mathbb{1}_{X_\tau}$ in fact lies almost entirely in V_k^j . In particular, we show that for every $i \neq j$, $\mathbb{1}_{X_\tau}$ has almost no projection onto V_k^i :

$$\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau, i} \rangle \leq O_k(\gamma) \mathbb{E}[\mathbb{1}_{X_\tau}].$$

Since the HD-Level-Set Decomposition is approximately orthogonal [DDFH18] (more generally, this is true of any approximate eigendecomposition, see Lemma 6.1), this implies that the mass on level j is around $(1 \pm O_k(\gamma)) \mathbb{E}[\mathbb{1}_{X_\tau}]$, and plugging these observations into Equation (3.1) we get:

$$\begin{aligned} \phi(M, \mathbb{1}_S) &\approx 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \lambda_i \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \\ &\approx 1 - \frac{\lambda_i}{\mathbb{E}[\mathbb{1}_S]} \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \\ &\approx 1 - \lambda_i \end{aligned}$$

where we have ignored some $O_k(\gamma)$ error terms.

Proof sketch of Theorem 2.5: Proving that pseudorandom functions expand is a bit more involved, but still at its core revolves around the analysis of Equation (3.1). In particular, since the approximate eigenvalues of M monotonically decrease (i.e. $\forall i \geq j$, $\lambda_i(M) \leq \lambda_j(M)$) then for any $\delta > 0$ we can simplify Equation (3.1) to:

$$(3.2) \quad \phi(M, \mathbb{1}_S) \lesssim 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^{R_\delta(M)-1} \lambda_i \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle + \delta,$$

where we recall $R_\delta(M)$ is the ST-rank, denoting the number of eigenstrips of M with eigenvalues greater than δ . Using Theorem 2.1, it is possible to show (see Proposition 7.3) that essentially all HD-walks of interest satisfy this property. With this in hand, proving Theorem 2.5 boils down to characterizing the projection of a set S onto low levels of the complex by its behavior on links.

THEOREM 3.2. (PSEUDORANDOMNESS BOUNDS LOW-LEVEL WEIGHT (INFORMAL THEOREM 8.2)) *Let X be a sufficiently strong d -dimensional γ -local-spectral expander. For any $i \leq k \leq d$ let $S \subset X(k)$ be a (ε, i) - ℓ_2 -pseudorandom set. Then for all $j \leq i$:*

$$\langle \mathbb{1}_S, \mathbb{1}_{S,j} \rangle \lesssim \binom{k}{i} \varepsilon \mathbb{E}[\mathbb{1}_S]$$

Theorem 2.5 follows immediately from plugging this result into Equation (3.2). Perhaps surprisingly, Theorem 3.2 itself follows without too much difficulty from combining analysis of the averaging operators and our spectral analysis of the HD-Level-Set Decomposition (namely Theorem 2.1). In particular, given a set $S \subset X(k)$, the idea is to examine the lowered indicator function $D_j^k \mathbb{1}_S$, which exactly gives the expectation of S over j -links, that is for any $\tau \in X(j)$:

$$D_j^k \mathbb{1}_S(\tau) = \mathbb{E}_{X_\tau}[\mathbb{1}_S].$$

Proving Theorem 2.5 then corresponds to lower bounding $\text{Var}(D_j^k \mathbb{1}_S)$ by some function of the projection of S onto level j . In fact, it is possible to exactly understand this connection. The idea is to reduce the problem to a

spectral analysis of HD-walks using the adjointness of D_j^k and U_j^k (see e.g. [DDFH18]):

$$\begin{aligned}\text{Var}(D_j^k \mathbb{1}_S) &= \langle D_j^k \mathbb{1}_S, D_j^k \mathbb{1}_S \rangle - \mathbb{E}[D_j^k \mathbb{1}_S]^2 \\ &= \langle \mathbb{1}_S, U_j^k D_j^k \mathbb{1}_S \rangle - \mathbb{E}[\mathbb{1}_S]^2 \\ &= \sum_{\ell=0}^k \langle \mathbb{1}_S, U_j^k D_j^k \mathbb{1}_{S,\ell} \rangle - \mathbb{E}[\mathbb{1}_S]^2 \\ &= \sum_{\ell=1}^k \langle \mathbb{1}_S, U_j^k D_j^k \mathbb{1}_{S,\ell} \rangle.\end{aligned}$$

Since $U_j^k D_j^k$ is an HD-walk, Theorem 2.1 implies that $U_j^k D_j^k \mathbb{1}_{S,\ell} \approx \lambda_\ell \mathbb{1}_{S,\ell}$ for some approximate eigenvalue λ_ℓ . In the formal version of Theorem 2.1 we compute the approximate eigenvalues of such walks, and in this case a simple computation shows that $\lambda_\ell = \frac{\binom{j}{\ell}}{\binom{k}{\ell}}$. As a result, we get a tight connection between $\text{Var}(D_j^k \mathbb{1}_S)$ and the projection of S onto low levels of the complex:

$$\text{Var}(D_j^k \mathbb{1}_S) \approx \sum_{\ell=1}^k \frac{\binom{j}{\ell}}{\binom{k}{\ell}} \langle \mathbb{1}_S, \mathbb{1}_{S,\ell} \rangle$$

where we have ignored factors in γ . The proof then follows from noting that the right-hand side is approximately lower bounded by $\frac{1}{\binom{k}{j}} \langle \mathbb{1}_S, \mathbb{1}_{S,j} \rangle$ by near orthogonality of the HD-Level-Set decomposition [DDFH18].

ℓ_∞ to ℓ_2 reduction: We end with a simple observation that generalizes our results to the ℓ_∞ -variant characterization via reduction. The idea is that the 2-norm may be re-written as a (normalized) expectation over a modified distribution:

$$\frac{1}{\mathbb{E}[\mathbb{1}_S]} \text{Var}(D_j^k \mathbb{1}_S) = \mathbb{E}_{\Pi_S} [D_j^k \mathbb{1}_S] - \mathbb{E}[\mathbb{1}_S] \leq \max(D_j^k \mathbb{1}_S - \mathbb{E}[\mathbb{1}_S])$$

where Π_S is the distribution on $X(j)$ naturally induced by restricting Π to the support of S (see Section 8 for further details). It immediately follows that any (ε, ℓ) - ℓ_∞ -pseudorandom set is also (ε, ℓ) - ℓ_2 -pseudorandom.

3.3 Playing Unique Games We end the section with a discussion of Theorem 2.6, an application of our structural theorems to algorithms for unique games. We follow the algorithmic framework of [BBK⁺21] for playing unique games on Johnson graphs, a special case of HD-walks on the complete complex. We generalize their algorithm and its analysis to HD-walks over all (sufficiently strong) two-sided local spectral expanders by abstracting the algorithm into the broader local-to-global paradigm for HDX:

1. For some r , break the UG instance down into sub-instances over r -links of the complex.
2. Solve the local sub-instances and patch together the solutions into a good global solution.¹⁴

For their underlying algorithmic machinery, BBKSS rely heavily on a well-studied paradigm known as the Sum-of-Squares (SoS) semidefinite programming hierarchy. We won't go into too much detail about this paradigm here (see Section 10.1 for details), and for the moment it suffices to say that analysing SoS algorithms relies on porting the proofs of relevant inequalities (specific to the problem at hand) into the Sum-of-Squares proof system, a restricted proof system for proving polynomial inequalities. We leverage the BBKSS framework to solve unique games on HDX by observing that their algorithm's analysis can be generalized to rely on two core structural properties true of the graphs underlying HD-walks:

1. A low-degree Sum-of-Squares proof that non-expanding sets have high variance in size across links.
2. For every small enough ε , the existence of some $r = r(\varepsilon)$ such that:

¹⁴This is only informally speaking, as the actual algorithm is iterative and repeats these two steps many times to obtain a good solution for the whole graph

(a) The $(r + 1)$ -st largest (distinct) stripped-eigenvalue of the constraint graph is small:

$$\lambda_r \leq 1 - \Omega(\varepsilon)$$

(b) The expansion of any r -link is small as well:

$$\forall \tau \in X(r) : \Phi(X_\tau) \leq O(\varepsilon).$$

In essence, properties (1) and (2a) ensure that there are good local solutions at level r , and property (2b) ensures that these solutions can be patched together without too much loss. Given the above properties, the proof of Theorem 2.6 follows the BBKSS analysis framework, but is more technical as properties that hold for Johnson graphs generally only hold approximately for HD-walks.

The novelty in our analysis lies mostly in proving these two properties. Luckily, we have already done most of the work! Property (1) is simply an SoS variant of Theorem 2.5, which we prove by developing SoS versions of the now standard machinery for the HD-Level-Set Decomposition from [DDFH18]. The proof then follows essentially as discussed in Section 3.2. The second property is slightly more subtle. This parameter, found in [BBK⁺21] by direct computation on the Johnson graphs, determines both the soundness and runtime of their algorithm. In fact, our framework completely demystifies its existence: since λ_r and the expansion of r -links are inversely correlated by Theorem 2.4, r is exactly the ST-Rank of the underlying constraint graph! As a result, using these properties and generalizing the analysis of BBKSS results in Theorem 2.6, an algorithm for unique games on HD-walks with approximation and runtime guarantees dependent on ST-Rank.

Organization

In Section 4 we discuss related work on HD-walks, Unique Games, and CSP approximation. In Section 5 we give requisite background and notation. In Section 6 we prove that approximate eigendecompositions tightly control spectral structure. In Section 7 we show that the combinatorial decomposition of [DDFH18] is an approximate eigendecomposition for any HD-walk and compute its corresponding approximate spectra, thereby determining the spectral structure of HD-walks. In Section 8 we prove Theorem 3.2, showing how pseudorandomness controls projections onto any HD-walks eigendecomposition. In Section 9 we use this fact to give a tight characterization of edge expansion in HD-walks. Finally, in Section 10 we give a sum-of-squares variant of our results and show they imply efficient algorithms for unique games over HD-walks.

4 Related Work

Higher Order Random Walks: The spectral structure of higher order random walks has seen significant study in recent years, starting with the work of Kaufman and Oppenheim [KO20] who proved bounds on the spectra of N_k^1 on one-sided local-spectral expanders. Their result lead not only to the resolution of the Mihail-Vazirani conjecture [ALOV19], but to a number of further breakthroughs in sampling algorithms via a small but consequential improvement on their bound by Alev and Lau [AL20]. The spectral structure of N_k^1 on the stronger two-sided local-spectral expanders was further studied by DDFH [DDFH18] who introduced the HD-Level-Set Decomposition, and Kaufman and Oppenheim [KO20] who introduced a distinct approximate eigendecomposition with the benefit of orthogonality (though this came at the cost of additional combinatorial complexity). In recent work, Kaufman and Sharakanski [KS20] claim that these two decompositions are equivalent on sufficiently strong two-sided γ -local-spectral expanders, but their proof relies on [KO20, Theorem 5.10] which has a non-trivial error. Indeed, it is possible to construct arbitrarily strong two-sided local-spectral expanders for which the HD-Level-Set Decomposition is not orthogonal (see Appendix B), so their result cannot hold.¹⁵ Finally, Alev, Jeronimo, and Tulsiani [AJT19] showed that the HD-Level-Set Decomposition is an approximate eigendecomposition (in a weaker sense than we require) for general HD-walks, a result we strengthen in Section 7. For further information on these prior works and their applications, the interested reader should see [Ale20].

Unique Games: The study of unique games has played a central role in hardness-of-approximation since Khot's [Kho02] introduction of the Unique Games Conjecture. One line of work towards *refuting* the UGC focuses on building efficient algorithms for unique games for restricted classes of constraint graphs based off of

¹⁵It is worth noting that the main results of [KO20, KS20] are unaffected by this error, as an approximate version of [KO20, Theorem 5.10] remains true and is sufficient for their purposes.

spectral or spectrally-related properties; these include works employing spectral expansion [AKK⁺08, MM10], threshold rank [ABS15, BRS11, GS11, Kol11], hypercontractivity [BBH⁺12], and certified small-set-expansion or characterized non-expansion [BBK⁺21]. Our work continues to expand this direction with polynomial-time algorithms for (affine) unique games over HD-walks and the introduction of ST-Rank. On the other hand, recent work towards *proving* the UGC has focused on characterizing non-expanding sets in structures such as the Grassmann [KMS18, KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMM18, KMS18] and Shortcode [BKS18, KMS18] graphs. Our spectral framework based on HD-walks and the HD-Level-Set Decomposition provides a more general method to approach this direction than previous Fourier analytic machinery.

CSPs on HDXs: Finally, it is worth noting a related, recent vein of work connecting high dimensional expansion, Sum of Squares, and CSP-approximation. In particular, Alev, Jeronimo, and Tulsiani [AJT19] recently showed that for $k > 2$, certain natural k -CSP's on two-sided local-spectral expanders can be efficiently approximated by Sum of Squares. Conversely, Dinur, Filmus, Harsha, and Tulsiani [DFHT20] later used cosystolic expanders (a stronger variant) to build explicit instances of 3-XOR that are hard for SoS. Both these works focus on k -CSPs for $k \geq 3$, where the variables of the CSP are level 1 of the complex and the constraints are defined using the higher levels of the complex; therefore these works do not encompass unique games. At a high level, the techniques of [AJT19] are based on [BRS11], and they generalize the BRS algorithm for 2-CSPs to k -CSPs. While these works are not directly related to ours since their definition do not encompass unique games, we see a similar pattern where high dimensional expanding structure is useful both for hardness of and algorithms for CSP-approximation.

5 Preliminaries and Notation

5.1 Local-Spectral Expanders and Higher Order Random Walks We now overview the theory of local-spectral expanders and higher order random walks in more formality than our brief treatment in Section 2.

5.1.1 Two-Sided Local-Spectral Expanders Two-sided local-spectral expanders are a generalization of spectral expander graphs to weighted, uniform hypergraphs, which we will think of as simplicial complexes.

DEFINITION 5.1. (WEIGHTED, PURE SIMPLICIAL COMPLEX) *A d -dimensional, pure simplicial complex X on n vertices is a subset of $\binom{[n]}{d}$. We will think of X as the downward closure of these sets, and in particular define the level $X(i)$ as:*

$$X(i) = \left\{ s \in \binom{[n]}{i} \mid \exists t \in X, s \subseteq t \right\}.$$

We call the elements of $X(i)$ i -faces.¹⁶ A simplicial complex is weighted if its top level faces are endowed with a distribution Π . This induces a distribution over each $X(i)$ by downward closure:

$$(5.3) \quad \Pi_i(x) = \frac{1}{i+1} \sum_{y \in X(i+1): y \supseteq x} \Pi_{i+1}(y),$$

where $\Pi_d = \Pi$.

Two-sided local-spectral expanders are based upon a phenomenon called *local-to-global* structure, which looks to propagate information on local neighborhoods of a simplicial complex called *links* to the entire complex.

DEFINITION 5.2. (LINK) *Given a weighted, pure simplicial complex (X, Π) , the **link** of an i -face $s \in X(i)$ is the sub-complex containing s , i.e.*

$$X_s = \{t \setminus s \in X \mid t \supseteq s\}.$$

Π induces a distribution over X_s by normalizing over top-level faces which we denote by Π_s . When considering a function on the k -th level of a complex, we also use X_s to denote the k -faces which contain s as long as it is clear from context, and refer to X_s as an i -link if $s \in X(i)$.

Two-sided local-spectral expansion simply posits that the graph underlying every link¹⁷ must be a two-sided spectral expander.

¹⁶We differ here from much of the HDX literature where an i -face is often defined to have $i+1$ elements. Since our work is mostly combinatorial rather than topological or geometric, defining an i -face to have i elements ends up being the more natural choice.

¹⁷The underlying graph of a simplicial complex X is its 1-skeleton $(X(1), X(2))$.

DEFINITION 5.3. (LOCAL-SPECTRAL EXPANSION) A weighted, pure simplicial complex (X, Π) is a two-sided γ -local-spectral expander if for every $i \leq d - 2$ and every face $s \in X(i)$, the underlying graph of X_s is a two-sided γ -spectral expander.¹⁸

5.1.2 Higher Order Random Walks Weighted simplicial complexes admit a natural generalization of the standard vertex-edge-vertex walk on graphs known as *higher order random walks* (HD-walks). The basic idea is simple: starting at some k -set $S \subset X(k)$, pick at random a set $T \in X(k+1)$ such that $T \supset S$, and then return to $X(k)$ by selecting some $S' \subset T$. Let the space of functions $f : X(k) \rightarrow \mathbb{R}$ be denoted by C_k . Formally, higher order random walks are a composition of two averaging operators: the “Up” operator which lifts a function $f \in C_k$ to $U_k f \in C_{k+1}$:

$$\forall y \in X(k+1) : U_k f(y) = \frac{1}{k+1} \sum_{x \in X(k) : x \subset y} f(x),$$

and the “Down” operator which lowers a function $f \in C_{k+1}$ to $D_{k+1} f \in C_k$:

$$\forall x \in X(k) : D_{k+1} f(x) = \frac{1}{k+1} \sum_{y \in X(k+1) : y \supset x} \frac{\Pi_{k+1}(y)}{\Pi_k(x)} f(y).$$

These operators exist for each level of the complex, and composing them gives a basic set of higher order random walks we call *pure* (following [AJT19]).

DEFINITION 5.4. (k -DIMENSIONAL PURE WALK) Given a weighted, simplicial complex (X, Π) , a k -dimensional pure walk $Y : C_k \rightarrow C_k$ on (X, Π) (of height $h(Y)$) is a composition:

$$Y = Z_{2h(Y)} \circ \cdots \circ Z_1,$$

where each Z_i is a copy of D or U .

We call an affine combination¹⁹ of pure walks which start and end on $X(k)$ a k -dimensional HD-walk.

DEFINITION 5.5. (HD-WALK) Let (X, Π) be a pure, weighted simplicial complex. Let \mathcal{Y} be a family of pure walks $Y : C_k \rightarrow C_k$ on (X, Π) . We call an affine combination

$$M = \sum_{Y \in \mathcal{Y}} \alpha_Y Y$$

a k -dimensional HD-walk on (X, Π) as long as it is self-adjoint and remains a valid walk (i.e. has non-negative transition probabilities).

Previous work on HD-walks mainly focuses on two natural classes: canonical walks (introduced in [KM16a, DK17]), and partial-swap walks (introduced in [AJT19, DD19]).

DEFINITION 5.6. (CANONICAL WALK) Given a d -dimensional weighted, pure simplicial complex (X, Π) , and parameters $k + j \leq d$, the canonical walk N_k^j is:

$$N_k^j = D_k^{k+j} U_k^{k+j},$$

where $U_\ell^k = U_{k-1} \dots U_\ell$, and $D_\ell^k = D_{\ell+1} \dots D_k$.

In other words, the canonical walk N_k^j takes j steps up and down the complex via the averaging operators. Partial-swap walks are a similar process, but after ascending the complex, we restrict to returning to faces with a given intersection from the starting point.

¹⁸A weighted graph $G(V, E)$ with edge weights Π_E is a two-sided γ -spectral expander if the vertex-edge-vertex random walk with transition probabilities proportional to Π_E has second largest eigenvalue in absolute value at most γ .

¹⁹An affine combination is a linear combination whose coefficients sum to 1.

DEFINITION 5.7. (PARTIAL-SWAP WALK) *The partial-swap walk S_k^j is the restriction of N_k^j to faces with intersection $k - j$. In other words, if $|s \cap s'| \neq k - j$, $S_k^j(s, s') = 0$, and otherwise $S_k^j(s, s') = \alpha_s N_k^j(s, s')$, where*

$$\alpha_s = \left(\sum_{s':|s \cap s'|=k-j} N_k(s, s') \right)^{-1}$$

is the appropriate normalization factor.

It is not hard to see that partial-swap walk S_k^t on the complete complex $J(n, d)$ (all d -subsets of $[n]$ endowed with the uniform distribution) is exactly the Johnson graph $J(n, k, k - t)$. While it is not immediately obvious that the partial-swap walks are HD-walks, Alev, Jeronimo, and Tulsiani [AJT19] showed this is the case by expressing them as an alternating hypergeometric sum of canonical walks.

5.1.3 Expansion of HD-Walks In this work, we study the combinatorial edge expansion of HD-walks, a fundamental property of graphs with strong connections to many areas of theoretical computer science, including both hardness and algorithms for unique games. Given a weighted graph $G = ((V, E), (\Pi_V, \Pi_E))$ where Π_V is a distribution over vertices, and Π_E is a set of non-negative edge weights, the expansion of a subset $S \subset V$ is the average edge-weight leaving S .

DEFINITION 5.8. (WEIGHTED EDGE EXPANSION) *Given a weighted, directed graph $G = ((V, E), (\Pi_V, \Pi_E))$, the weighted edge expansion of a subset $S \subset V$ is:*

$$\phi(G, S) = \mathbb{E}_{v \sim \Pi_V|_S} [E(v, V \setminus S)],$$

where

$$E(v, V \setminus S) = \sum_{(v, y) \in E: y \in V \setminus S} \Pi_E((v, y))$$

is the total weight of edges between vertex v and the subset $V \setminus S$, and $\Pi_V|_S$ is the re-normalized restriction of Π_V to S . In the context of a k -dimensional HD-Walk M on a weighted simplicial complex (X, Π) , we will always have $V = X(k)$, $\Pi_V = \Pi_k$, and E, Π_E given by M . Thus when clear from context, we will simply write $\phi(S)$.

Edge expansion in a weighted graph is closely related to the spectral structure of its adjacency matrix. Given a set $S \subset V$ of density $\alpha = \mathbb{E}[\mathbb{1}_S]$, we may write

$$\phi(G, S) = 1 - \frac{1}{\alpha} \langle \mathbb{1}_S, A_G \mathbb{1}_S \rangle_{\Pi_V},$$

where A_G is the adjacency matrix with weights given by Π_E , and $\langle f, g \rangle_{\Pi_V}$ is the expectation of fg over Π_V . When considering such an inner product over a weighted simplicial complex (X, Π) , the associated distribution will always be Π_k , so we will drop it from the corresponding notation. Notice that the right-hand side of this equivalence may be further broken down via a spectral decomposition of $\mathbb{1}_S$ with respect to A_G . Thus to understand the edge-expansion of HD-walks, it is crucial to understand the structure of their spectra.

6 Approximate Eigendecompositions and Eigenstripping

In this section we prove a general linear algebraic result concerning the spectra of operators that admit an approximate eigendecomposition: their spectra lies tightly concentrated around the decomposition's approximate eigenvalues. Before giving the formal result, we formalize the concept of approximate eigendecompositions.

DEFINITION 6.1. *Let M be an operator over an inner product space V . We call $V = V^1 \oplus \dots \oplus V^k$ a $(\{\lambda_i\}_{i=1}^k, \{c_i\}_{i=1}^k)$ -approximate eigendecomposition if for all i and $v_i \in V^i$, the following holds:*

$$\|Mv_i - \lambda_i v_i\| \leq c_i \|v_i\|.$$

As long as the c_i are sufficiently small, we prove each V^i (loosely) corresponds to an eigenstrip, the span of eigenvectors with eigenvalue closely concentrated around λ_i .

THEOREM 6.1. (EIGENSTRIPPING) *Let M be a self-adjoint operator over an inner product space V , and $V = V^1 \oplus \dots \oplus V^k$ a $(\{\lambda_i\}_{i=1}^k, \{c_i\}_{i=1}^k)$ -approximate eigendecomposition. Let $c_{\max} = \max_i \{c_i\}$, $\lambda_{\text{dif}} = \min_{i,j} \{|\lambda_i - \lambda_j|\}$, and $\lambda_{\text{ratio}} = \frac{\max_i \{|\lambda_i|\}}{\lambda_{\text{dif}}^{1/2}}$. Then as long as c_{\max} is sufficiently small:*

$$c_{\max} \leq \frac{\lambda_{\text{dif}}}{4k},$$

the spectra of M is concentrated around each λ_i :

$$\text{Spec}(M) \subseteq \bigcup_{i=1}^k [\lambda_i - e, \lambda_i + e] = I_{\lambda_i},$$

where $e = O\left(k \cdot \lambda_{\text{ratio}} \cdot c_{\max}^{1/2}\right)$.

This result was recently improved by Zhang [Zha20] to have dependence $e \leq O(\sqrt{k}c_{\max})$, removing the dependence on the approximate eigenvalues altogether. This follows from a simple modification to our proof which we will note below. Finally, it is also worth mentioning that a version of Theorem 6.1 holds with no assumption on c_{\max} , but the assumption substantially simplifies the bounds and is sufficient for our purposes.

Before proving Theorem 6.1, we note a useful property of approximate eigendecompositions of self-adjoint operators: they are approximately orthogonal.

LEMMA 6.1. *Let M be a self-adjoint operator over an inner product space V . Further, let $V = V^1 \oplus \dots \oplus V^k$ be a $(\{\lambda_i\}_{i=1}^k, \{c_i\}_{i=1}^k)$ -approximate eigen-decomposition. Then for $i \neq j$, V^i and V^j are nearly orthogonal. That is, for any $v_i \in V^i$ and $v_j \in V^j$:*

$$|\langle v_i, v_j \rangle| \leq \frac{c_i + c_j}{|\lambda_i - \lambda_j|} \|v_i\| \|v_j\|.$$

Proof. This follows from the fact that M is self-adjoint, and V^i and V^j are approximate eigenspaces. In particular, notice that for any $v_i \in V^i$ and $v_j \in V^j$ we can bound the interval in which $\langle Mv_i, v_j \rangle = \langle v_i, Mv_j \rangle$ lies by Cauchy-Schwarz:

$$\langle Mv_i, v_j \rangle \in \lambda_i \langle v_i, v_j \rangle \pm c_i \|v_i\| \|v_j\|$$

and

$$\langle v_i, Mv_j \rangle \in \lambda_j \langle v_i, v_j \rangle \pm c_j \|v_i\| \|v_j\|.$$

Since these terms are equal, the right-hand intervals must overlap. As a result we get:

$$|(\lambda_i - \lambda_j) \langle v_i, v_j \rangle| \leq (c_i + c_j) \|v_i\| \|v_j\|,$$

as desired. \square

Using Lemma 6.1, we can modify [KO20, Theorem 5.9] to prove Theorem 6.1. Given an eigenvalue μ of M , the idea is to find a probability distribution over $[k]$ for which the expectation of $|\mu - \lambda_i|$ is small, where $i \in [k]$ is sampled from the aforementioned distribution.

Proof. The proof follows mostly along the lines of [KO20, Theorem 5.9], modifying where necessary due to lack of orthogonality. Let ϕ be an eigenvector of M with eigenvalue μ . Our goal is to prove the existence of some λ_i such that $|\mu - \lambda_i|$ is small. To do this, we appeal to an averaging argument. In particular, denoting the component of ϕ in V^i by ϕ_i , we bound the expectation of $|\mu - \lambda_i|^2$ over a distribution P_ϕ given by the (normalized) squared norms $\|\phi_i\|^2$:

$$(6.4) \quad \mathbb{E}_{i \sim P_\phi} [|\mu - \lambda_i|^2] = \frac{1}{\sum_{j=1}^k \|\phi_j\|^2} \sum_{i=1}^k |\mu - \lambda_i|^2 \|\phi_i\|^2.$$

If we can upper bound this expectation by some value c , then by averaging there must exist λ_i such that $|\mu - \lambda_i| \leq \sqrt{c}$, and thus the spectra of M must lie in strips $\lambda_i \pm \sqrt{c}$. To upper bound Equation (6.4), consider the result of pushing the outer summation inside the norm:

$$(6.5) \quad \sum_{i=1}^k |\mu - \lambda_i|^2 \|\phi_i\|^2 = \left\| \sum_{i=1}^k (\mu - \lambda_i) \phi_i \right\|^2 - \sum_{1 \leq i \neq j \leq k} (\mu - \lambda_i)(\mu - \lambda_j) \langle \phi_i, \phi_j \rangle.$$

We will separately bound the two resulting terms, the former by the fact that the ϕ_i are approximate eigenvectors, and the latter by their approximate orthogonality. We start with the former, which follows by a simple application of Cauchy-Schwarz:

$$\begin{aligned} \left\| \sum_{i=1}^k (\mu - \lambda_i) \phi_i \right\|^2 &= \left\| \mu \phi - \sum_{i=1}^k \lambda_i \phi_i \right\|^2 \\ &= \left\| M\phi - \sum_{i=1}^k \lambda_i \phi_i \right\|^2 \\ &= \left\| \sum_{i=1}^k (M\phi_i - \lambda_i \phi_i) \right\|^2 \\ &\leq k \sum_{i=1}^k \| (M\phi_i - \lambda_i \phi_i) \|^2 \\ &\leq k c_{\max}^2 \sum_{i=1}^k \|\phi_i\|^2. \end{aligned}$$

The latter takes a bit more effort. Let λ_{\max} be $\max_i \{|\lambda_i|\}$, then by Lemma 6.1 we have:

$$\begin{aligned} \left| \sum_{1 \leq i \neq j \leq k} (\mu - \lambda_i)(\mu - \lambda_j) \langle \phi_i, \phi_j \rangle \right| &\leq \sum_{1 \leq i \neq j \leq k} |\mu - \lambda_i| |\mu - \lambda_j| \frac{c_i + c_j}{|\lambda_i - \lambda_j|} \|\phi_i\| \|\phi_j\| \\ &\leq 2c_{\max} \lambda_{\text{dif}}^{-1} (\lambda_{\max} + \|M\|)^2 \left(\sum_{i=1}^k \|\phi_i\| \right)^2 \\ &\leq 2k c_{\max} \lambda_{\text{dif}}^{-1} (\lambda_{\max} + \|M\|)^2 \sum_{i=1}^k \|\phi_i\|^2 \end{aligned}$$

Since we'd like our bound to depend only on λ_i and c_i , we must further bound $\|M\|$ which will follow similarly from approximate orthogonality. Let v be a unit eigenvector with eigenvalue $\|M\|$ and v_i be v 's component on V^i ,

then we have:

$$\begin{aligned}
\|M\| &= \|Mv\| \\
&= \left\| \sum_{i=1}^k Mv_i - \lambda_i v_i + \lambda_i v_i \right\| \\
&\leq \sum_{i=1}^k (\lambda_i + c_i) \|v_i\| \\
&\leq (\lambda_{\max} + c_{\max}) \sum_{i=1}^k \|v_i\| \\
&\leq (\lambda_{\max} + c_{\max}) \sqrt{k \sum_{i=1}^k \|v_i\|^2} \\
&\leq (\lambda_{\max} + c_{\max}) \sqrt{2k}.
\end{aligned}$$

where the last step follows from Lemma 6.1 and our assumption on c_{\max} :

$$\begin{aligned}
\sum_{i=1}^k \|v_i\|^2 &= \|v\|^2 + \sum_{1 \leq i \neq j \leq k} \langle v_i, v_j \rangle \\
&\leq 1 + \frac{2c_{\max}}{\lambda_{\text{dif}}} \sum_{1 \leq i \neq j \leq k} \|v_i\| \|v_j\| \\
&\leq 1 + \frac{2c_{\max}}{\lambda_{\text{dif}}} \left(\sum_{i=1}^k \|v_i\| \right)^2 \\
&\leq 1 + \frac{2kc_{\max}}{\lambda_{\text{dif}}} \sum_{i=1}^k \|v_i\|^2 \\
&\leq 1 + \frac{1}{2} \sum_{i=1}^k \|v_i\|^2.
\end{aligned}$$

Together, these bounds imply the existence of some $\lambda_{i'}$ such that:

$$|\mu - \lambda_{i'}| \leq \sqrt{kc_{\max} \left(c_{\max} + 2\lambda_{\text{dif}}^{-1} \left(\lambda_{\max} + (\lambda_{\max} + c_{\max}) \sqrt{2k} \right)^2 \right)},$$

which implies the desired result when accounting for our assumption on c_{\max} . \square

Zhang's [Zha20] improvement to our proof came from the observation that the analysis of the latter term can be simplified by a recursive strategy. In particular, this term can instead be upper bounded by $\frac{2c_{\max}k}{\lambda_{\text{dif}}} \sum_{i=1}^k |\mu - \lambda_i|^2 \|\phi_i\|^2$.

With the appropriate assumption on c_{\max} , plugging this back into Equation (6.5) gives the desired result.

In either case, notice that if c_{\max} is sufficiently small, the intervals I_{λ_i} are disjoint. As a result, each V^i corresponds to an *eigenstrip* W^i :

$$W^i = \text{Span} \{ \phi : M\phi = \mu\phi, \mu \in I_{\lambda_i} \}.$$

The approximate eigenspaces V^i are closely related to the resulting eigenstrips. Indeed, it is possible to show that most of the weight of a function in V^i must lie on W^i , though we will not need this result in what follows. Previous works [KO20, KS20] make stronger claims for the specific case of the HD-Level-Set Decomposition, most notably that V^i and W^i are in fact equivalent on sufficiently strong two-sided local-spectral expanders. Unfortunately, these results are based off of [KO20, Theorem 5.10], whose proof has a non-trivial error we discuss

further in Appendix B. Indeed, were their proof correct, it would imply (due to the generality of their argument) that $V^i = W^i$ for any approximate eigendecomposition. However, it is easy to see this cannot be the case by considering a diagonal 2×2 matrix with an approximate eigendecomposition given by a slight rotation of the standard basis vectors in \mathbb{R}^2 .

7 The Spectra of HD-walks

We now show that the HD-Level-Set Decomposition is an approximate eigendecomposition for any HD-Walk. Combined with Theorem 6.1, this proves Theorem 2.1 that the spectrum of any k -dimensional HD walk is tightly concentrated in $k+1$ eigenstrips. As a result, we give explicit bounds on the spectra of HD-walks, paying special attention to the canonical and partial-swap walks. Finally, we show that the approximate eigenvalues (and thus the values in their corresponding eigenstrips) of the HD-Level-Set Decomposition decrease monotonically for a broad class of HD-Walks we call *complete* walks which, to our knowledge, encompass all walks used in the literature. As we will see in the following section, such decay is crucial for understanding edge expansion.

To start, we recall the definition of pure and HD-walks along with introducing some useful notation.

DEFINITION 7.1. (k -DIMENSIONAL PURE WALK) *Given a weighted, simplicial complex (X, Π) , a k -dimensional pure walk $Y : C_k \rightarrow C_k$ on (X, Π) is a composition:*

$$Y = Z_{2h(Y)} \circ \cdots \circ Z_1,$$

where each Z_i is a copy of D or U , and $h(Y)$ is the height of the walk, measuring the total number of down (or up) operators.

DEFINITION 7.2. (k -DIMENSIONAL HD-WALK) *Given a weighted, simplicial complex (X, Π) , a k -dimensional HD-walk on (X, Π) is an affine combination of pure walks*

$$M = \sum_{Y \in \mathcal{Y}} \alpha_Y Y$$

which is self-adjoint and gives a valid walk on (X, Π) (i.e. has non-negative transition probabilities). We say the height of M , $h(M)$, is the maximal height of any Y with a non-zero coefficient, and say the weight of M , $w(M)$, is the one norm of the α_Y (namely, $w(M) = \sum |\alpha_Y|$).

Our proofs in this section rely mainly on a useful observation of [DDFH18], who show that the up and down operators on two-sided γ -local-spectral expanders satisfy the following relation:

$$(7.6) \quad \left\| D_{i+1}U_i - \frac{1}{i+1}I - \frac{i}{i+1}U_{i-1}D_i \right\| \leq \gamma.$$

This fact leads to a particularly useful structural lemma showing the effect of flipping D through multiple U operators.

LEMMA 7.1. (CLAIM 8.8 [DDFH18]) *Let (X, Π) be a d -dimensional γ -local-spectral expander. Then for all $j < k < d$:*

$$\left\| D_{k+1}U_{k-j}^{k+1} - \frac{j+1}{k+1}U_{k-j}^k - \frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} \right\| \leq \frac{(j+1)(2k-j+2)}{2(k+1)}\gamma.$$

One crucial application of Lemma 7.1 lies in understanding the relation between $\|f_i\|$, and $\|g_i\|$, where $f_i = U_i^k g_i$.

LEMMA 7.2. (LEMMAS 8.10, 8.13, THEOREM 4.6 [DDFH18]) *Let (X, Π) be a d -dimensional γ -local-spectral expander with $\gamma < 1/d$, $f \in C_k$ a function with HD-Level-Set Decomposition $f_0 + \dots + f_k$. Then for all $0 \leq \ell \leq k \leq d$:*

$$\|f_\ell\|^2 = \frac{1}{\binom{k}{\ell}} (1 \pm c_1(k, \ell)\gamma) \|g_\ell\|^2,$$

where $c_1(k, \ell) = O(k^2 \binom{k}{\ell})$.

In Appendix A, we prove a stronger version of both Lemma 7.1 and Lemma 7.2 for $\gamma \leq 2^{-\Omega(k)}$ where the dependence on the first order term γ is polynomial rather than exponential in k . However, since this only provides a substantial improvement for a small range of γ , we use the simpler versions from [DDFH18] throughout the body of the paper. Using Lemma 7.1 and Lemma 7.2, an inductive argument shows that the HD-Level-Set Decomposition is an approximate eigendecomposition. We show this first for the basic case of a pure walk, and then note that the general result follows immediately from the triangle inequality.

PROPOSITION 7.1. *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $Y : C_k \rightarrow C_k$ a pure walk:*

$$Y = Z_{2h(Y)} \circ \cdots \circ Z_1.$$

Let $i_1 \leq \dots \leq i_{h(Y)}$ denote the $h(Y)$ indices at which Z_i is a down operator. Then for all $0 \leq \ell \leq k$, $f \in V_k^\ell$:

$$\left\| Yf - \prod_{s=1}^{h(Y)} \left(1 - \frac{\ell}{\max\{\ell, i_s - 2s + k + 1\}} \right) f \right\| \leq O \left(\gamma h(Y)(k + h(Y)) \binom{k}{\ell} \|f\| \right).$$

Proof. We prove a slightly stronger statement to simplify the induction. For $b > 0$, let $Y_j^b : C_\ell \rightarrow C_{\ell+b}$ denote an unbalanced walk with j down operators, and $j+b$ up operators. If Y_j^b has down operators in positions $i_1 \leq \dots \leq i_j$ and $g_\ell \in H^\ell$, we claim:

$$(7.7) \quad \left\| Y_j^b g_\ell - \prod_{s=1}^j \left(1 - \frac{\ell}{\max\{\ell, i_s - 2s + \ell + 1\}} \right) U_\ell^{b+\ell} g_\ell \right\| \leq \gamma j(b+j) \|g_\ell\|.$$

Notice that since $f \in V_k^\ell$ may be written as $U_\ell^k g_\ell$ for $g_\ell \in H^\ell$, then we may write Yf as $Y_{h(Y)}^{k-\ell} g_\ell$ where $Y_{h(Y)}^{k-\ell}$ has down operators in positions $i_1 + k - \ell \leq \dots \leq i_j + k - \ell$. Combining Equation (7.7) with Lemma 7.2 then implies the result.

We prove Equation (7.7) by induction. The base case $j = 0$ is trivial. Assume the inductive hypothesis holds for all Y_i^b , $i < j$. Notice first that if $i_1 = 1$, we are done since $g_\ell \in H^\ell$, and

$$\prod_{s=1}^j \left(1 - \frac{\ell}{\max\{\ell, i_s - 2s + \ell + 1\}} \right) Y_0^b g_\ell = 0,$$

as $i_s - 2s + \ell + 1 = \ell$ for $s = 1$. Otherwise, it must be the case that one or more copies of the up operator appear before the first down operator, and we may therefore apply Lemma 7.1 to get:

$$Y_j^b g_\ell = \left(\frac{i_1 - 1}{i_1 + \ell - 1} \right) Y_{j-1}^b g_\ell + \Gamma g_\ell,$$

where we can (loosely) bound the spectral norm of Γ by

$$\|\Gamma\| \leq (b+j)\gamma$$

since at worst the first down operator D passes through $b+j$ up operators. By the form of Lemma 7.1, Y_{j-1}^b has down operators at indices $i_2 - 2 \leq \dots \leq i_j - 2$. Then by the fact that $i_1 + \ell - 1 > \ell$ and the inductive hypothesis:

$$\begin{aligned} Y_j^b g_\ell &= \left(\frac{i_1 - 1}{\max\{\ell, i_1 + \ell - 1\}} \prod_{s=1}^{j-1} \frac{i_{s+1} - 2s - 1}{\max\{\ell, i_{s+1} - 2s + \ell - 1\}} \right) Y_0^b g_\ell + \frac{i_1 - 1}{i_1 + \ell - 1} h + \Gamma g_\ell \\ &= \left(\frac{i_1 - 1}{\max\{\ell, i_1 + \ell - 1\}} \prod_{s=2}^j \frac{i_s - 2s + 1}{\max\{\ell, i_s - 2s + \ell + 1\}} \right) Y_0^b g_\ell + \frac{i_1 - 1}{i_1 + \ell - 1} h + \Gamma g_\ell \\ &= \left(\prod_{s=1}^j \frac{i_s - 2s + 1}{\max\{\ell, i_s - 2s + \ell + 1\}} \right) Y_0^b g_\ell + \frac{i_1 - 1}{i_1 + \ell - 1} h + \Gamma g_\ell, \end{aligned}$$

where $\|h\| \leq \gamma(j-1)(b+j-1) \|g_\ell\|$ and we have used the (vacuous) fact that $\max\{\ell, i_1 + \ell - 1\} = i_1 + \ell - 1$. Finally, we can bound the norm of the right-hand error term by:

$$\begin{aligned} \left\| \frac{i_1 - 1}{i_1 + \ell - 1} h + \Gamma g_\ell \right\| &\leq \|h\| + \|\Gamma\| \|g_\ell\| \\ &\leq (j-1)(b+j-1) \|g_\ell\| + (b+j) \|g_\ell\| \\ &\leq j(b+j) \|g_\ell\| \end{aligned}$$

as desired. \square

Since HD-walks are simply affine combinations of pure walks, the triangle inequality immediately implies the result carries over to this more general setting.

COROLLARY 7.1. *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $M = \sum_i \alpha_i Y_i$ a k -dimensional HD-walk on (X, Π) . Then for all $0 \leq \ell \leq k$, $f \in V_k^\ell$:*

$$\|Mf - \lambda_\ell(M)f\| \leq O\left(\gamma w(M)h(M)(k + h(M)) \binom{k}{\ell} \|f\|\right),$$

where

$$\lambda_\ell(M) = \sum \alpha_i \lambda_\ell(Y_i)$$

and $\lambda_\ell(Y_i)$ is the approximate eigenvalue of Y_i given in Proposition 7.1.

It is worth noting that the resulting approximate eigenvalues in Corollary 7.1 are exactly the eigenvalues of M when considered on a sequentially differential poset with $\vec{\delta}_i = i/(i+1)$. We discuss this generalization in more depth and give tighter bounds on the approximate spectra in our upcoming companion paper. It should be noted that this result is similar to one appearing in [AJT19], where a weaker notion of approximate eigenspaces based on the quadratic form $\langle f, Mf \rangle$ is analyzed. Plugging Corollary 7.1 into Theorem 6.1, we immediately get that for small enough γ the true spectra of HD-walks lie in strips around each $\lambda_i(M)$, and thus that the approximate eigenvalues of the HD-Level-Set Decomposition and the spectra of HD-walks are essentially interchangeable.

For concreteness, we now turn our attention to computing the approximate eigenvalues (and thereby the true spectra) of the canonical and swap walks.

COROLLARY 7.2. (SPECTRUM OF CANONICAL WALKS) *Let (X, Π) be a d -dimensional γ -local-spectral expander with γ satisfying $\gamma \leq 2^{-\Omega(k+j)}$, $k+j \leq d$, and $f_\ell \in V_k^\ell$. Then:*

$$\left\| N_k^j f_\ell - \frac{\binom{k}{\ell}}{\binom{k+j}{\ell}} f_\ell \right\| \leq c(k, \ell, j) \|f_\ell\|,$$

where $c(k, \ell, j) = O\left(\gamma j(j+k) \binom{k}{\ell}\right)$. Moreover:

$$\text{Spec}(N_k^j) = \{1\} \cup \bigcup_{j=1}^k \left[\frac{\binom{k}{\ell}}{\binom{k+j}{\ell}} \pm 2^{O(j+k)} \sqrt{\gamma} \right].$$

Proof. By Proposition 7.1, N_k^j is an $(\{\lambda_\ell\}_{\ell=0}^k, \{c(k, \ell, j)\}_{\ell=0}^k)$ -approximate eigendecomposition for

$$\lambda_\ell = \prod_{s=1}^j \left(1 - \frac{\ell}{\max\{k-2s+i_s+1, \ell\}} \right),$$

where $i_1 \leq \dots \leq i_s$ denote the indices of down operators. By the definition of N_k^j we have $i_s = j+s$, and therefore

$$\lambda_\ell = \prod_{s=1}^j \left(1 - \frac{\ell}{k-s+j+1} \right) = \frac{\binom{k}{\ell}}{\binom{k+j}{\ell}}$$

as desired. The bounds on $\text{Spec}(N_k^j)$ follow immediately from plugging the above into Theorem 6.1. \square

A priori, it is not obvious how to bound the spectra of the partial-swap walks, or indeed even that they are HD-walks. However, Alev, Jeronimo, and Tulsiani [AJT19] proved that partial-swap walks may be written as an alternating hypergeometric sum of canonical walks.

PROPOSITION 7.2. (COROLLARY 4.13 [AJT19]) *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma < 1/k$. Then for $0 \leq j \leq k$:*

$$S_k^j = \frac{1}{\binom{k}{k-j}} \sum_{i=0}^j (-1)^{j-i} \binom{j}{i} \binom{k+i}{i} N_k^i.$$

As a result, we can use Corollary 7.2 to bound their approximate eigenvalues and true spectrum.

COROLLARY 7.3. *Let X be d -dimensional two-sided γ -local-spectral expander, $\gamma < 2^{-\Omega(k)}$, $k+j \leq d$, and $f_\ell \in V_k^\ell$. Then:*

$$\left\| S_k^j f_\ell - \frac{\binom{k-j}{\ell}}{\binom{k}{\ell}} f_\ell \right\| \leq c(k) \|f_\ell\|,$$

where $c(k) = \gamma 2^{O(k)}$. Moreover,

$$\text{Spec}(S_k^j) = \{1\} \cup \bigcup_{j=1}^k \left[\frac{\binom{k-j}{\ell}}{\binom{k}{\ell}} \pm 2^{O(k)} \sqrt{\gamma} \right].$$

Proof. By Corollary 7.1, $\bigoplus_{\ell=0}^k V_k^\ell$ is a $(\{\lambda_\ell\}_{\ell=0}^k, \{c'(k, \ell, j)\}_{\ell=0}^k)$ -approximate eigendecomposition for S_k^j with

$$\begin{aligned} \lambda_\ell &= \frac{1}{\binom{k}{k-j}} \sum_{i=0}^j (-1)^{j-i} \binom{j}{i} \binom{k+i}{i} \lambda_\ell(N_k^i) \\ &= \frac{1}{\binom{k}{k-j}} \sum_{i=0}^j (-1)^{j-i} \binom{j}{i} \binom{k+i}{i} \frac{\binom{k}{\ell}}{\binom{k+i}{\ell}} \\ &= \frac{1}{\binom{k}{k-j}} \sum_{i=0}^j (-1)^{j-i} \binom{j}{i} \binom{k-\ell+i}{i} \\ &= \frac{\binom{k-\ell}{j}}{\binom{k}{k-j}} \\ &= \frac{\binom{k-j}{\ell}}{\binom{k}{\ell}}, \end{aligned}$$

and

$$c'(k, \ell, j) = \gamma 2^{O(k)}.$$

This latter fact follows from noting that

$$\|\vec{\alpha}\|_1 = \sum_{i=0}^j \binom{j}{i} \binom{k+i}{i} \leq 2^{2j+k}$$

where $\vec{\alpha}$ consists of the hypergeometric coefficients of Proposition 7.2. The bounds on $\text{Spec}(S_k^j)$ then follow from Theorem 6.1. \square

Together, Corollary 7.2 and Corollary 7.3 prove Theorem 2.3 (assuming γ is sufficiently small).

In Theorem 2.1, we mentioned that approximate eigenvalues $\lambda_i(M)$ are monotonically decreasing for any HD-walk. In fact, to prove this we will need to restrict our original definition of HD-walks slightly, requiring that our walks are always well-defined on the complete complex.

DEFINITION 7.3. (COMPLETE HD-WALK) Let (X, Π) be a weighted, pure simplicial complex and $M = \sum_{Y \in \mathcal{Y}} \alpha_Y Y$ an HD-walk on (X, Π) . We call M complete if for all $n \in \mathbb{N}$ there exist $n_0 > n$ and d such that $\sum_{Y \in \mathcal{Y}} \alpha_Y Y$ is also an HD-walk when taken to be over $J(n_0, d)$.

To our knowledge, all walks considered in the literature (pure, canonical, partial-swap) are complete. We can prove that the eigenstrips of complete HD-walks corresponding to the HD-Level-Set Decomposition exhibit eigenvalue decay by noting that the approximate eigenvalues of Corollary 7.1 are independent of the underlying complex.

PROPOSITION 7.3. Let (X, Π) be a two-sided γ -local-spectral expander, $M = \sum_{Y \in \mathcal{Y}} \alpha_Y Y$ a complete HD-walk over (X, Π) , and γ small enough to apply the conditions of Theorem 6.1. Then for all $0 \leq i < j \leq k$,

$$\lambda_i(M) \geq \lambda_j(M)$$

Proof. The proof follows from two observations. First, recall from Corollary 7.1 that $\lambda_i(M)$ is independent of the underlying complex. Second, any HD-walk on the complete complex can be written as a non-negative sum of partial-swap walks, which satisfy the monotonic decrease property. Let $n \in \mathbb{N}$ be any parameter such that applying $\sum_{Y \in \mathcal{Y}} \alpha_Y Y$ to $J(n, d)$ results in a valid walk (i.e. a non-negative matrix). By the symmetry of $J(n, d)$, the transition probabilities of this walk depends only on size of intersection, and it may thus be written as some convex combination of partial-swap walks:

$$M = \sum_{Y \in \mathcal{Y}} \alpha_Y Y = \sum_{i=0}^k \beta_i S_k^i.$$

Since these walks are equivalent over $J(n, d)$, their spectra must match. Then by Theorem 6.1, it must be the case that for every $1 \leq \ell \leq k$ and n sufficiently large, the intersection of $\sum_{Y \in \mathcal{Y}} \alpha_Y \lambda_\ell(Y) \pm O(1/n)$ and $\sum_{Y \in \mathcal{Y}} \beta_i \lambda_\ell(S_k^i) \pm O(1/n)$ is non-empty. Since we may take n arbitrarily large, this implies the two quantities are in fact equivalent. Finally, by Corollary 7.3 $\lambda_\ell(S_k^i)$ decreases monotonically in ℓ for all i , which implies that the $\lambda_i(M) = \sum_{Y \in \mathcal{Y}} \alpha_Y \lambda_i(Y) = \sum_{Y \in \mathcal{Y}} \beta_i \lambda_i(S_k^i)$ decrease monotonically as desired. \square

8 Pseudorandomness and the HD-Level-Set Decomposition

Now that we have examined the spectral structure of the HD-Level-Set Decomposition, we turn to understanding its combinatorial characteristics. In this section, we give a combinatorial characterization how arbitrary functions project onto the HD-Level-Set decomposition, proving in particular a generalization of Theorem 3.2: pseudorandom sets have bounded projection onto corresponding levels of the complex. We discuss two variants of pseudorandomness, an ℓ_2 -variant stating that the variance across links is small, and an ℓ_∞ -variant stating that the maximum (or ℓ_∞ -norm for arbitrary functions) across links is small. We focus mainly on giving an exact analysis for the ℓ_2 -case, as this forms the structural core of our algorithm for unique games in Section 10. We further discuss the implications of our ℓ_2 analysis to the ℓ_∞ -variant by reduction.

As such, we'll start by analyzing the ℓ_2 -variant. First, let's extend our definition of ℓ_2 -pseudorandomness from sets (boolean functions) to arbitrary functions.

DEFINITION 8.1. (ℓ_2 -PSEUDORANDOM FUNCTIONS) A function $f \in C_k$ is $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom if its variance across i -links is small for all $1 \leq i \leq \ell$:

$$\text{Var}(D_i^k f) \leq \varepsilon_i |\mathbb{E}[f]|$$

As mentioned in Section 3, the key to connecting $\text{Var}(D_i^k f)$ and the HD-Level-Set Decomposition is to notice that by the adjointness of D and U , we can reduce the problem to analyzing the spectral structure of HD-walks. The proof then follows immediately from arguments in the previous section.

THEOREM 8.1. *Let (X, Π) be a γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and let $f \in C_k$ have HD-Level-Set Decomposition $f = f_0 + \dots + f_k$. Then for any $\ell \leq k$, the $\text{Var}(D_\ell^k f)$ is controlled by its projection onto $V_k^0 \oplus \dots \oplus V_k^\ell$ in the following sense:*

$$\text{Var}(D_\ell^k f) \in \sum_{j=1}^{\ell} \frac{\binom{\ell}{j}}{\binom{k}{j}} \langle f, f_j \rangle \pm c\gamma \langle f, f \rangle$$

where $c \leq 2^{O(k)}$.

Proof. To start, notice that since $\langle D_\ell^k f, D_\ell^k f \rangle = \langle U_\ell^k D_\ell^k f, f \rangle$ it is enough to analyze the application of the HD-walk $U_\ell^k D_\ell^k$ to f . By Proposition 7.1, we know that each f_j is an approximate eigenvector satisfying:

$$\left\| U_\ell^k D_\ell^k f_j - \frac{\binom{\ell}{j}}{\binom{k}{j}} f_j \right\| \leq c_1 \gamma \|f\|,$$

where $c_1 \leq 2^{O(k)}$. Combining these observations gives:

$$\begin{aligned} \langle D_\ell^k f, D_\ell^k f \rangle &= \langle f, U_\ell^k D_\ell^k f \rangle \\ &= \sum_{j=0}^k \langle f, U_\ell^k D_\ell^k f_j \rangle \\ &\in \sum_{j=0}^{\ell} \frac{\binom{\ell}{j}}{\binom{k}{j}} \langle f, f_j \rangle \pm c\gamma \langle f, f \rangle \end{aligned}$$

where all constants $c \leq 2^{O(k)}$, and noting that $\langle f, f_0 \rangle = \mathbb{E}[f]^2$ completes the result. \square

Theorem 3.2 follows immediately from combining this result with approximate orthogonality of the HD-Level-Set Decomposition.

LEMMA 8.1. (DDFH THEOREM 4.6) *Let (X, Π) be a γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and let $f \in C_k$ have HD-Level-Set Decomposition $f = f_0 + \dots + f_k$. Then for all $i \neq j$:*

$$|\langle f_i, f_j \rangle| \leq c_1 \gamma \langle f, f \rangle$$

where $c_1 \leq 2^{O(k)}$.

COROLLARY 8.1. *Let (X, Π) be a γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and let $f \in C_k$ be an $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom function. Then for any $1 \leq i \leq \ell$:*

$$|\langle f, f_i \rangle| \leq \binom{k}{i} \varepsilon_i |\mathbb{E}[f]| + c\gamma \langle f, f \rangle$$

where $c \leq 2^{O(k)}$.

Proof. By Lemma 8.1, for all j we have $\langle f, f_j \rangle \geq -c\gamma \langle f, f \rangle$ for some $c \leq 2^{O(k)}$. Then by Theorem 8.1, for all $0 \leq i \leq k$ the variance of $D_i^k f$ is lower bounded by the projection onto f_i :

$$\text{Var}(D_i^k f) \geq \frac{1}{\binom{k}{i}} \langle f, f_i \rangle - c_2 \gamma \langle f, f \rangle,$$

where $c_2 \leq 2^{O(k)}$. Finally, since f is $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom, we have by definition that for all $1 \leq i \leq \ell$, $\text{Var}(D_i^k f) \leq \varepsilon_i |\mathbb{E}[f]|$. Therefore isolating $\langle f, f_i \rangle$ gives the desired upper bound:

$$\begin{aligned} \langle f, f_i \rangle &\leq \binom{k}{i} \text{Var}(D_i^k f) + c_3 \gamma \langle f, f \rangle \\ &\leq \binom{k}{i} \varepsilon_i |\mathbb{E}[f]| + c_3 \gamma \langle f, f \rangle \end{aligned}$$

where $c_3 \leq 2^{O(k)}$. Finally, we already noted that $\langle f, f_i \rangle \geq -c\gamma \langle f, f \rangle$ for some $c \leq 2^{O(k)}$, which gives the desired lower bound and completes the proof. \square

It is worth noting that both Theorem 8.1 and Corollary 8.1 are tight. This is obvious for the former which is a near-equality, and the latter is clearly tight for subsets like i -links which project (almost) entirely onto level i (we'll prove this formally in the next section).

We'll also see in Section 9 that Corollary 8.1 leads to a tight ℓ_2 -characterization of edge-expansion stating that any non-expanding set must have high variance across links. Since an ℓ_∞ -variant of this result for the Grassmann graphs was recently crucial for the resolution of the 2-2 Games Conjecture [KMS18], it is natural to discuss what our results imply for this regime. First, let's formalize ℓ_∞ -pseudorandomness for arbitrary functions.

DEFINITION 8.2. (ℓ_∞ -PSEUDORANDOM FUNCTIONS) *A function $f \in C_k$ is $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_∞ -pseudorandom if for all $1 \leq i \leq \ell$ its local expectation is close to its global expectation:*

$$\|D_i^k f - \mathbb{E}[f]\|_\infty \leq \varepsilon_i.$$

We will prove via reduction to the ℓ_2 -variant that a version of Corollary 8.1 holds in this regime as well. We proceed in two steps. First, we show that ℓ_∞ -pseudorandom functions are also ℓ_2 -pseudorandom assuming a weak local-consistency property.

DEFINITION 8.3. *Let (X, Π) be a weighted, pure simplicial complex. We say a function $f \in C_k$ has ℓ -local constant sign if:*

1. $\mathbb{E}[f] \neq 0$,
2. $\forall s \in X(\ell)$ s.t. $\mathbb{E}_{X_s}[f] \neq 0$: $\text{sign}(\mathbb{E}_{X_s}[f]) = \text{sign}(\mathbb{E}[f])$.

Second, we'll reduce to the case of locally constant sign by noting that we can always shift a function to satisfy this property. With these definitions in hand, we can now state the ℓ_∞ -variant of Theorem 3.2.

THEOREM 8.2. *Let (X, Π) be a γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and let $f \in C_k$ have HD-Level-Set Decomposition $f = f_0 + \dots + f_k$. If f is $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_∞ -pseudorandom, then for all $1 \leq i \leq \ell$:*

$$|\langle f, f_i \rangle| \leq \left(\binom{k}{i} + c(k)\gamma \right) \varepsilon_i^2 + c(k)\gamma \|f\|^2,$$

where $c(k) \leq 2^{O(k)}$, and if f has i -local constant sign:

$$|\langle f, f_i \rangle| \leq \binom{k}{i} \varepsilon_i |\mathbb{E}[f]| + c(k)\gamma \|f\|^2.$$

In dealing with expansion, we will mainly be interested in boolean-valued functions, which always have locally-constant sign and satisfy $\langle f, f \rangle = \mathbb{E}[f]$. Thus in the boolean case we have:

$$\langle f, f_i \rangle \leq \left(\binom{k}{i} \varepsilon_i + 2^{O(k)} \gamma \right) \mathbb{E}[f],$$

which is particularly useful since the expansion of f may be written as $1 - \frac{1}{\mathbb{E}[f]} \langle f, Mf \rangle$. In the case f is non-negative, we can also replace the ℓ_∞ norm with maximum in our definition of pseudorandomness, which is important to show non-expanding sets are locally *denser* than expected. Finally, we note that one can improve the dependence on γ by pushing exponential dependence on k to the second order γ^2 term via more careful analysis of error propagation. However since the analysis is complicated and only gives a substantial improvement for a small range of relevant γ , we relegate such discussion to Appendix A.

We now move to the proof of Theorem 8.2, starting with our generic ℓ_∞ to ℓ_2 reduction for functions with locally-constant sign.

LEMMA 8.2. *Let (X, Π) be a weighted, pure simplicial complex, and $f \in C_k$ a $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_∞ -pseudorandom function with i -local constant sign for any $i \leq \ell$. Then f is also $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom*

Proof. For ease of notation, let $\Pi_k(X_s)$ be shorthand for $\sum_{t \in X_s} \Pi_k(t)$ (i.e. the normalization factor for the above restricted expectation). The trick is to notice that since f has locally constant sign, we may rewrite $\|D_i^k f\|_2^2$ as an expectation over a related distribution P_i :

$$\begin{aligned} \frac{1}{\mathbb{E}[f]} \langle D_i^k f, D_i^k f \rangle &= \sum_{s \in X(i)} \Pi_i(s) \left(\frac{1}{\mathbb{E}[f]} \sum_{t \in X_s} \frac{\Pi_k(t)f(t)}{\Pi_k(X_s)} \right) D_i^k f(s) \\ &= \sum_{s \in X(i)} \left(\frac{1}{\mathbb{E}[f]} \sum_{t \in X_s} \frac{\Pi_k(t)f(t)}{\binom{k}{i}} \right) D_i^k f(s) \\ &= \mathbb{E}_{P_i}[D_i^k f], \end{aligned}$$

where we have used the fact that $\Pi_k(X_s) = \binom{k}{i} \Pi_i(s)$ by Equation (5.3). To understand $P_i(s)$ more intuitively, consider the special case when f is non-negative. Here, Π and f induce a distribution P_k over $X(k)$, where

$$P_k(t) = \frac{\Pi_k(t)f(t)}{\mathbb{E}[f]}.$$

P_k then induces the distribution P_i on $X(i)$ via the following process: draw a face $t \in X(k)$ from P_k , and then choose a i -face $s \subset t$ uniformly at random. Replacing the non-negativity of f with the conditions in the theorem statement still leaves $P_i(s)$ a valid distribution, albeit one with a less intuitive description.

The result then follows from an averaging argument:

$$\left| \frac{1}{\mathbb{E}[f]} \text{Var}(D_i^k f) \right| = \left| \mathbb{E}_{P_i}[D_i^k f] - \mathbb{E}[f] \right| \leq \|D_i^k f - \mathbb{E}[f]\|_\infty$$

We note that when $\mathbb{E}[f] > 0$, the ℓ_∞ -norm may be replaced with maximum in the above. \square

To complete the proof of Theorem 8.2, we reduce to Theorem 8.1 by noting that any function may be shifted to have locally constant sign and applying our reduction.

Proof. [Proof of Theorem 8.2] Note that the latter bound for functions with locally constant sign is immediate since, by Lemma 8.2, f is also $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom and therefore satisfies the guarantees of Corollary 8.1.

For the former, assume for simplicity that $\mathbb{E}[f] \geq 0$ (the negative case follows from a similar argument) and consider the shifted function $f' = f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}$. Notice that as long as $\varepsilon_i > 0$, f' has positive expectation over all i -links and non-zero expectation, and further that

$$f' = f'_0 + f_1 + \dots + f_k,$$

where $f'_0 = f_0 + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}$ and $f = \sum f_i$ is the original HD-Level-Set Decomposition of f . Since adding a constant has no effect on the ℓ_∞ -pseudorandomness, f' remains $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_∞ -pseudorandom and, by our ℓ_∞ to ℓ_2 reduction, $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom as well. Applying Corollary 8.1 then gives:

$$\begin{aligned} \langle f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}, f_i \rangle &\leq \binom{k}{i} \varepsilon_i \mathbb{E}[f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}] + c\gamma \langle f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}, f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1} \rangle \\ &\leq \left(\binom{k}{i} + c\gamma \right) \varepsilon_i^2 + c\gamma \langle f, f \rangle \end{aligned}$$

where we have used the fact that f_i is orthogonal to $\mathbb{1}$ for all $i > 0$. We are left to deal with the case that $\varepsilon_i = 0$, which follows from a limiting argument applying the above to any $\varepsilon > 0$. \square

9 Expansion of HD-walks

In this section we characterize the edge expansion of HD-walks, proving Theorem 2.4, Theorem 2.5, and Corollary 2.1. As a reminder, these results focus on two key aspects of edge-expansion: the expansion of links, and the structure of non-expanding sets. We'll start with the former, but first let's recall the definition of edge-expansion (specified to HD-walks for simplicity).

DEFINITION 9.1. (WEIGHTED EDGE EXPANSION) *Given a weighted simplicial complex (X, Π) , a k -dimensional HD-Walk M over (X, Π) , and a subset $S \subset X(k)$, the weighted edge expansion of S is*

$$\phi(S) = \mathbb{E}_{v \sim \Pi_k|_S} [M(v, X(k) \setminus S)],$$

where

$$M(v, X(k) \setminus S) = \sum_{y \in X(k) \setminus S} M(v, y)$$

and $M(v, y)$ is the transition probability from v to y .

We start by proving Theorem 2.4: that the expansion of i -links is (up to $O(\gamma)$ error) exactly controlled by the eigenvalue of the i th eigenstrip.

THEOREM 9.1. (LOCAL EXPANSION VS GLOBAL SPECTRA) *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$, and M a k -dimensional, complete HD-walk with $k < d$. Then for all $0 \leq i \leq k$ and $\tau \in X(i)$:*

$$\phi(X_\tau) \in 1 - \lambda_i(M) \pm c\gamma,$$

where $c \leq w(M)h(M)^2 2^{O(k)}$.

The main idea behind Theorem 9.1 is simply to show that the indicator function of any i -link always lies almost entirely in V_k^i (or equivalently, almost entirely in the i th eigenstrip W_k^i).

LEMMA 9.1. *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$. Then for all $0 \leq i \leq k < d$ and $\tau \in X(i)$, 1_{X_τ} lies almost entirely in V_k^i . That is for all $j \neq i$:*

$$\langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle \leq c\gamma \langle 1_{X_\tau}, 1_{X_\tau} \rangle$$

where $c \leq 2^{O(k)}$.

Proof. It is enough to analyze the expansion of X_τ with respect to N_k^1 since the quantity can both be analyzed directly, and expressed in terms of X_τ 's HD-Level-Set decomposition. For the direct analysis, recall that N_k^1 describes the process of moving from a k -face σ to a $(k+1)$ -face $T = \sigma \cup \{v\}$, then back to a k -face $\sigma' \subset T$. Crucially, the latter step is performed uniformly at random. Applying N_k^1 to any element in X_τ , the probability of returning to X_τ is exactly the probability that we remove an element in $T \setminus \{\tau\}$ in the final step, which gives:

$$\bar{\phi}(1_{X_\tau}) = \frac{k+1-i}{k+1}.$$

On the other hand, we may also expand out $\bar{\phi}(1_{X_\tau})$ in terms of 1_{X_τ} 's HD-Level-Set decomposition. Using the fact that $1_{X_\tau} = \binom{k}{i} U_i^k 1_\tau \in V_k^0 \oplus \dots \oplus V_k^i$, we may write:

$$\begin{aligned} \bar{\phi}(1_{X_\tau}) &= \frac{1}{\langle 1_{X_\tau}, 1_{X_\tau} \rangle} \sum_{j=0}^i \langle 1_{X_\tau}, N_k^1 1_{X_{\tau,j}} \rangle \\ &= \frac{1}{\langle 1_{X_\tau}, 1_{X_\tau} \rangle} \sum_{j=0}^i \frac{k+1-j}{k} \langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle + \frac{1}{\langle 1_{X_\tau}, 1_{X_\tau} \rangle} \sum_{s=0}^i \langle 1_{X_\tau}, \Gamma_s \rangle, \end{aligned}$$

where $\|\Gamma_s\| \leq 2^{O(k)}\gamma\|1_{X_\tau}\|$ by Corollary 7.1 and the fact that $\|1_{X_{\tau,s}}\| \leq (1 + 2^{O(k)}\gamma)\|1_{X_\tau}\|$. Finally, applying Cauchy-Schwarz to the error term gives:

$$(9.8) \quad \bar{\phi}(1_{X_\tau}) \in \frac{1}{\langle 1_{X_\tau}, 1_{X_\tau} \rangle} \sum_{j=0}^i \frac{k+1-j}{k} \langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle \pm c\gamma$$

for $c \leq 2^{O(k)}\gamma$. Recall by approximate orthogonality (Lemma 8.1), the projections $\langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle$ cannot be too negative, that is $\langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle \geq -c\gamma \langle 1_{X_\tau}, 1_{X_\tau} \rangle$ for some $c \leq 2^{O(k)}$. Then if there exists some $j \neq i$ such that $\langle 1_{X_\tau}, 1_{X_{\tau,j}} \rangle > c_2\gamma \langle 1_{X_\tau}, 1_{X_\tau} \rangle$ for large enough $c_2 \leq 2^{O(k)}$, the LHS of Equation (9.8) lies strictly above $\frac{k+1-i}{k+1}$, giving the desired contradiction. \square

Since the HD-Level-Set decomposition of 1_{X_τ} has no dependence on the walk in question, Theorem 9.1 follows almost immediately.

Proof. [Proof of Theorem 9.1] The expansion of X_τ may be written as:

$$\begin{aligned}\phi(X_\tau) &= 1 - \frac{1}{\alpha} \langle 1_{X_\tau}, M 1_{X_\tau} \rangle \\ &= 1 - \frac{1}{\alpha} \sum_{s=0}^i \langle 1_{X_\tau}, M 1_{X_\tau, s} \rangle,\end{aligned}$$

where α is the density of X_τ and $1_{X_\tau, s} \in V_k^s$. By Corollary 7.1 we can simplify the sum up to an error term:

$$\begin{aligned}\phi(X_\tau) &= 1 - \frac{1}{\alpha} \sum_{s=0}^i \langle 1_{X_\tau}, M 1_{X_\tau, s} \rangle \\ &= 1 - \frac{1}{\alpha} \sum_{s=0}^i \lambda_j(M) \langle 1_{X_\tau}, 1_{X_\tau, s} \rangle + \frac{1}{\alpha} \sum_{s=0}^i \langle 1_{X_\tau}, \Gamma_s \rangle,\end{aligned}$$

where by the fact that $\|1_{X_\tau, s}\| \leq (1 + 2^{O(k)}\gamma) \|1_{X_\tau}\|$ we have $\|\Gamma_s\| \leq O\left(w(M)h(M)(h(M) + k)\binom{k}{s}\gamma \|1_{X_\tau}\|\right)$. By Lemma 9.1 for all $s \neq i$ we can absorb $\langle 1_{X_\tau}, 1_{X_\tau, s} \rangle$ into the error term which gives:

$$\phi(X_\tau) \in 1 - \frac{1}{\alpha} \lambda_i(M) \langle 1_{X_\tau}, 1_{X_\tau, i} \rangle \pm c\gamma$$

for $c \leq w(M)h(M)^2 2^{O(k)}$. Finally, by Lemma 9.1 and approximate orthogonality, we know that $\langle 1_{X_\tau}, 1_{X_\tau, i} \rangle$ is very close to α :

$$(1 - c_1\gamma)\alpha \leq \langle 1_{X_\tau}, 1_{X_\tau, i} \rangle \leq (1 + c_2\gamma)\alpha$$

where $c_1, c_2 \leq 2^{O(k)}$. Combining this with the above completes the proof. Note that the proof falls through for $d = k$ since we cannot analyze N_k^1 on such a complex. The upper bound, however, still holds in this case simply by expanding out the inner product and bounding the inner summation using $\lambda_i(M)$. \square

Theorem 9.1 will form one of two core pieces of our algorithm with unique games. The fact that links corresponding to bad eigenvalues have poor expansion will allow us to patch together good local solutions into a global solution without seeing too much interference. Moreover, close connection between local expansion and stripped eigenvalues will result in the performance of our algorithm being tied directly to ST-rank (which we recall here for convenience).

DEFINITION 9.2. (STRIPPED THRESHOLD RANK) *Let (X, Π) be a two-sided γ -local-spectral expander and M a k -dimensional HD-walk with γ small enough that the HD-Level-Set Decomposition has a corresponding decomposition of disjoint eigenstrips $C_k = \bigoplus W_k^i$ ²⁰. The ST-Rank of M with respect to δ is the number of strips containing an eigenvector with eigenvalue at least δ :*

$$R_\delta(M) = |\{W_k^i : \exists f \in V^i, Mf = \lambda f, \lambda > \delta\}|.$$

We often write just R_δ when M is clear from context.

A basic corollary of Theorem 9.1 is that, like the Johnson graphs, links are small, non-expanding sets (at least when $|X(1)| \gg k$ or γ is small). The second core piece of our algorithm for unique games relies on a certain converse to this result: that *all* non-expanding sets are explained by links. As discussed in Section 1 and Section 2, we consider two regimes for this problem. The first, which we call the ℓ_2 -variant, claims that any non-expanding set must have high *variance* over links—this regime is useful for constructing algorithms for unique games, as we’ll show in the next section. The second is the ℓ_∞ -variant, which claims that any non-expanding set must have a high *maximum* over links—this regime is useful in hardness of approximation. We examine both regimes through their converse: that both ℓ_2/ℓ_∞ -pseudorandom sets expand near-perfectly.

²⁰It should be noted that when the HD-Level-Set Decomposition has spaces with the same approximate eigenvalue, their corresponding eigenstrips technically must be merged. However, since this detail has no effect on our arguments, we ignore it in what follows. We now show how to express the expansion of a set $S \subset X(k)$ with respect to an HD-walk M in terms of the pseudorandomness of S and the ST-rank of M .

THEOREM 9.2. *Let (X, Π) be a two-sided γ -local-spectral expander, M a k -dimensional, complete HD-walk, and let γ be small enough that the eigenstrip intervals of Theorem 6.1 are disjoint. For any $\delta > 0$, let $r = R_\delta(M) - 1$. Then the expansion of a set $S \subset X(k)$ of density α is at least:*

$$\phi(S) \geq 1 - \alpha - (1 - \alpha)\delta - c\gamma - \sum_{i=1}^r (\lambda_i(M) - \delta) \binom{k}{i} \varepsilon_i,$$

where $\lambda_i(M)$ is the approximate eigenvalue given by Corollary 7.1, S is either $(\varepsilon_1, \dots, \varepsilon_r)$ - ℓ_2 -pseudorandom or $(\varepsilon_1, \dots, \varepsilon_r)$ - ℓ_∞ -pseudorandom, and $c \leq w(M)h(M)^2 2^{O(k)}$.

Proof. Recall that the expansion of S may be written as:

$$\phi(S) = 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \langle \mathbb{1}_S, M\mathbb{1}_S \rangle.$$

Decomposing $\mathbb{1}_S = \mathbb{1}_{S,0} + \dots + \mathbb{1}_{S,k}$ by the HD-Level-Set Decomposition, we have:

$$\begin{aligned} \phi(S) &= 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \langle \mathbb{1}_S, M\mathbb{1}_{S,i} \rangle \\ &= 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \lambda_i(M) \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle + \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=1}^k \langle \mathbb{1}_S, \Gamma_i \rangle \end{aligned}$$

where by Corollary 7.1 $\|\Gamma_i\| \leq O\left(w(M)h(M)(h(M) + k) \binom{k}{i} \gamma \|\mathbb{1}_{S,i}\|\right)$. Using Cauchy-Schwarz and the fact that $\|\mathbb{1}_{S,i}\| \leq (1 + 2^{O(k)}\gamma) \|\mathbb{1}_S\|$ (this follows from approximate orthogonality, see [DDFH18, Corollary 8.13]) we can simplify this to

$$\phi(S) \geq 1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^k \lambda_i(M) \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle - e\gamma,$$

where $e \leq w(M)h(M)^2 2^{O(k)}$. Since M is a complete walk, we know the $\lambda_i(M)$ decrease monotonically and as long as γ is sufficiently small, correspond to the eigenvalues in strip W^i as well. Thus we may write:

$$\begin{aligned} \phi(S) &\geq 1 - e\gamma - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^r \lambda_i(M) \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle - \frac{\delta}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=r+1}^k \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \\ &= 1 - e_2\gamma - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=1}^r \lambda_i(M) \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle - \delta \left(1 - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=0}^r \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \right) \\ &= 1 - e_2\gamma - \delta - \frac{1}{\mathbb{E}[\mathbb{1}_S]} \sum_{i=1}^r (\lambda_i(M) - \delta) \langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \end{aligned}$$

Finally, recalling that $\mathbb{1}_{S,i} = \mathbb{E}[\mathbb{1}_S] \vec{1}$ and applying Theorem 8.2 gives the ℓ_2 -variant result, and the ℓ_∞ -variant follows immediately from our ℓ_∞ to ℓ_2 reduction (Lemma 8.2). \square

We now turn to the discussion of a surprisingly subtle point: the tightness of Theorem 9.2. There are two main parameters of interest: the pseudorandomness parameter ε , and the level of the HD-walk k . We'll first prove that in both the ℓ_2 and ℓ_∞ regimes, if we fix the dependence on ε to be linear, our bound is exactly tight.

PROPOSITION 9.1. *Let $X = J(n, d)$ be the complete complex, $2k - t \leq d$, $m|n$, and B_m be the set of all k -faces $\binom{[n/m]}{k}$. Then for any t , B_m witnesses the tightness of Theorem 9.2 with respect to S_k^{k-t} as $n, m \rightarrow \infty$.*

Proof. First, note that it is enough to examine only the ℓ_∞ bound, since the function in question has locally-constant sign the tightness of the ℓ_2 bound follows from Lemma 8.2.

By direct computation, it is not hard to show that the expansion of B_m with respect to S_k^{k-t} is:

$$\phi(B_m) = 1 - \frac{\binom{\frac{n}{m}-k}{k-t}}{\binom{n-k}{k-t}} = 1 - \frac{m^t}{m^k} + O_{k,m}(1/n)$$

On the other hand, we can directly compute that B_m is $(\varepsilon_1, \dots, \varepsilon_k)$ -pseudorandom, where:

$$\varepsilon_i \leq \frac{\binom{\frac{n}{m}-i}{k-i}}{\binom{n-i}{k-i}}$$

Since the complete complex is a two-sided $O(1/n)$ -local-spectral expander [DDFH18], for large enough n Theorem 9.2 gives the bound:

$$\begin{aligned} \phi(B_m) &\geq 1 - \sum_{i=0}^t \binom{t}{i} \frac{\binom{\frac{n}{m}-i}{k-i}}{\binom{n-i}{k-i}} - O_{k,m}(1/n) \\ &= 1 - \frac{\binom{n/m}{k}}{\binom{n}{k}} \sum_{i=0}^t \binom{t}{i} \frac{\binom{n}{i}}{\binom{n/m}{i}} - O_{k,m}(1/n) \\ &\geq 1 - m^{-k} \sum_{i=0}^t \binom{t}{i} m^i - O_{k,m}(1/n) \\ &= 1 - \frac{(m+1)^t}{m^k} - O_{k,m}(1/n) \end{aligned}$$

Thus we see that for large n , the bound is tight up to the leading term in m . \square

However, Proposition 9.1 does not preclude a bound with better (or even no) dependence on k . Indeed, proving that such a bound holds in the ℓ_∞ -regime for the Johnson graphs was an important stepping stone in the proof of the 2-2 Games Conjecture [KMMS18]. On the other hand, we can actually prove that a k -independent bound is *impossible* in the ℓ_2 -regime. The intuition behind the difference can be summarized by examining the behavior of these two variants on a link. In the ℓ_∞ regime, links are of course $\Omega(1)$ -pseudorandom. On the other hand, somewhat counter-intuitively, links are actually $O(\binom{k}{i}^{-1})$ -pseudorandom in the ℓ_2 -regime. Since links tend to have poor expansion, the dependence on k in our bound has to make up for the fact that links are $O(\binom{k}{i}^{-1})$ - ℓ_2 -pseudorandom (this also explains the particular dependence on $\binom{k}{i}$).

PROPOSITION 9.2. *For every $\ell \in \mathbb{N}$ and $c_1(\ell), c_2(\ell) > 0$, there exist $n \gg k \gg \ell$ and an $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom subset S of $S_k^{k/2}$ on $J(n, 2k)$ ²¹ of density α satisfying:*

$$\phi(S) < 1 - \alpha - \lambda_{\ell+1}(S_k^{k/2}) - c_1 \varepsilon^{c_2}$$

Proof. The result follows from letting S be any ℓ -link of the complete complex $J(n, 2k)$ for sufficiently large n, k . In particular, notice that the expansion of S with respect to $S_k^{k/2}$ is then:

$$\phi(S) \leq 1 - 2^{-\ell} + O_k\left(\frac{1}{n}\right).$$

On the other hand, note that $\lambda_{\ell+1}(S_k^{k/2}) \in 2^{-\ell-1} \pm O_k(\frac{1}{n})$. Since one can also check that S is $(\varepsilon_1, \dots, \varepsilon_\ell)$ - ℓ_2 -pseudorandom for $\varepsilon_i \leq \binom{k}{i}^{-1} + O_k(\frac{1}{n})$, we for large enough $n \gg k$ that:

$$1 - \alpha - \lambda_{\ell+1}(S_k^{k/2}) - c_1 \varepsilon^{c_2} \geq 1 - 2^{-\ell-1} - O_k\left(\frac{1}{n}\right) - o_k(1) > \phi(S)$$

as desired. \square

²¹Note that this is exactly the Johnson Graph $J(n, k, k/2)$ mentioned in Section 2.

In other words, no k -independent version of Theorem 9.2 exists for the ℓ_2 -variant, as such a result would violate the upper bound in Proposition 9.2. While this doesn't directly rule out a reduction from ℓ_∞ to ℓ_2 that proves a k -independent bound for the former (the reduction itself would need to be k -dependent in this case), it's known to be impossible in our framework which encompasses the Grassmann where such a result is known to be false [DKK⁺18a].

The difference between a k -independent bound and the regime we consider is most stark when examining the contrapositive of Theorem 9.2, which states that non-expanding sets must be concentrated inside links.

COROLLARY 9.1. *Let (X, Π) be a two-sided γ -local-spectral expander, M a k -dimensional, complete HD-walk, and let γ be small enough to satisfy the requirements of Theorem 6.1. Then for any $\delta > 0$, if $S \subset X(k)$ is a set of density α and expansion:*

$$\phi(S) < 1 - \alpha - (1 - \alpha)\delta - c\gamma$$

for $c \leq w(M)h(M)^22^{O(k)}$, then S is non-trivially correlated with an i -link for $1 \leq i \leq R_{\delta/2}$:

$$\exists 1 \leq i \leq R_{\delta/2}, \tau \in X(i) : \mathbb{E}_{X_\tau}[\mathbb{1}_S] \geq \alpha + \frac{\delta}{c_2 R_{\delta/2} \binom{k}{i} \lambda_i(M)},$$

where $c_2 > 2$ is some small absolute constant.

Notice that the excess correlation implied by Corollary 9.1 decays as k grows large (except in the case of very deep walks like $S_k^{k-O(1)}$); this is one of the main obstructions to using results like Corollary 9.1 for hardness of unique games. On the other hand, we now turn our attention to *algorithms* for unique games, where the ℓ_2 -regime gets its chance to shine.

10 Playing Unique Games on HD-Walks

Following the recent algorithmic framework of [BBK⁺21], we show how to translate our spectral machinery and combinatorial characterization of non-expansion into a polynomial time algorithm for unique games over HD-walks whose finer-grained runtime and approximation guarantees depend on ST-rank. Before we dive into the theorems let's start with some notation for this section.

Notation: We briefly comment on the use of $\mathcal{I} = (M, \mathcal{S})$ to denote an affine UG instance over a random-walk M , since unique games are formally defined over graphs, not random walks. In particular, it is well known that every HD-walk M over $X(k)$ uniquely corresponds to an undirected weighted graph $G_M = (X(k), E)$ (see Section C.1). Thus by $\mathcal{I} = (M, \mathcal{S})$, we really mean that the constraints \mathcal{S} are over the edges of G and the value is calculated according to the distribution over edges E . We will use the $\tilde{O}(\cdot)$ notation in this section to hide log factors, that is, $\tilde{O}(f)$ denotes $O(f \log^c f)$ for any constant $c > 0$.

With notation out of the way, let us describe the main theorem:

THEOREM 10.1. *For any $\varepsilon \in (0, .01)$, there exists an algorithm \mathcal{A} with the following guarantee. Let (X, Π) be a d -dimensional two-sided γ -local spectral expander and M be a k -dimensional complete HD-walk over X such that $\gamma \leq w(M)^{-1}2^{-\Omega(k+h(M))}$ and $d > k$. Let $\mathcal{I} = (M, \mathcal{S})$ be an instance of affine unique games over M with value at least $1 - \varepsilon$. Then \mathcal{A} outputs an $\Omega\left(\frac{\varepsilon^3}{\binom{k}{r(2\varepsilon)}^2}\right)$ -satisfying assignment in time $|X(k)|^{\tilde{O}(\binom{k}{r(2\varepsilon)}^{\frac{1}{c}})}$, where $r(\varepsilon) = R_{1-16\varepsilon}(M)$ is the ST-rank of M .*

Since the HDX literature focuses mostly on canonical and partial-swap walks, we also give the specification of Theorem 10.1 to this class for concreteness. Here we see that the finer-grained performance of our algorithm depends on the *depth* of the walk.

COROLLARY 10.1. *For any $\varepsilon \in (0, .01)$, there exists an algorithm \mathcal{A} with the following guarantee. Let (X, Π) be a d -dimensional two-sided γ -local spectral expander and M be a k -dimensional canonical or partial-swap walk of depth $\beta \in [0, 1]$ over X such that $\gamma \leq w(M)^{-1}2^{-\Omega(k+h(M))}$ and $d > k$. Let $\mathcal{I} = (M, \mathcal{S})$ be an instance of affine unique games over M with value at least $1 - \varepsilon$. Then \mathcal{A} outputs an $\Omega\left(\frac{\varepsilon^3}{\binom{k}{c\varepsilon/\beta}^2}\right)$ -satisfying assignment in time $|X(k)|^{\tilde{O}(\binom{k}{c\varepsilon/\beta}^{\frac{1}{c}})}$ for some absolute constant $c > 0$.*

To prove Theorem 10.1, we follow the general SoS algorithmic paradigm introduced by BBKSS for Johnson graphs (partial-swap walks on the complete complex), described in Section 10.2 for completeness. As discussed in Section 3.3 we abstract BBKSS' analysis to rely on two core structural properties true of the graphs underlying HD-walks:

1. There exists a low-degree Sum-of-Squares proof that non-expanding sets have high variance in size across links.
2. For every small enough ε , there exists a parameter $r = r(\varepsilon)$ such that:

(a) The $(r + 1)$ -st largest (distinct) stripped-eigenvalue of G is small:

$$\lambda_r \leq 1 - \Omega(\varepsilon)$$

(b) The expansion of any r -link is small as well:

$$\forall \tau \in X(s) : \Phi(X_\tau) \leq O(\varepsilon).$$

Recall that we have already proved property (1) in Theorem 9.2, albeit not in the low-degree SoS proof system. In Section 10.4 we modify the proof of Theorem 9.2 to show that it has a degree 2 SoS proof. For property (2), it is clear that the existence of $r(\varepsilon)$ is an inherent consequence of Theorem 9.1, and furthermore it is exactly the $(1 - O(\varepsilon))$ ST-Rank of the HD-walk. Analyzing the generalized BBKSS algorithm using these properties combined with a few more technicalities (described in Section 10.4) then gives an efficient algorithm for unique games over HD-walks on two-sided local spectral expanders.

10.1 Background for Unique Games and SoS Proving Theorem 10.1 from the ground up requires substantial background in the SoS framework. However, since we mostly rely on a number of higher level results from [BBK⁺21] for the SoS side of our work, we cover here only background necessary to understand our methods, and refer the reader to the surveys of [BS14a], [FKP19], [RSS18] and additionally Sections 1, 2, and A of [BBK⁺21] for more information.

The Sum of Squares framework is a method for approximating polynomial optimization problems through semi-definite programming relaxations. In particular, given the problem:

$$\text{Maximize } p \in \mathbb{R}[x_1, \dots, x_n] \text{ constraint to } \{q_i = 0\}_{i=1}^m,$$

for $q_i \in \mathbb{R}[x_1, \dots, x_n]$, the Degree- D Sum of Squares semidefinite programming relaxation outputs in time $n^{O(D)}$ a *pseudoexpectation operator* $\tilde{\mathbb{E}} : \text{poly}_{\mathbb{R}}(n, D) \rightarrow \mathbb{R}$ over polynomials in $\mathbb{R}[x_1, \dots, x_n]$ of degree at most D satisfying:

1. Scaling: $\tilde{\mathbb{E}}[1] = 1$
2. Linearity: $\tilde{\mathbb{E}}[af(x) + bg(x)] = a\tilde{\mathbb{E}}[f(x)] + b\tilde{\mathbb{E}}[g(x)]$
3. Non-negativity (for squares): $\tilde{\mathbb{E}}[f(x)^2] \geq 0$
4. Program constraints: $\tilde{\mathbb{E}}[f(x)q_i(x)] = 0$
5. Optimality: $\tilde{\mathbb{E}}[p(x)] \geq \max_x \{p(x) : \{q_i = 0\}_{i=1}^m\}$

Note that the first four properties give the definition of a pseudoexpectation (under constraints $\{q_i = 0\}_{i=1}^m$), whereas the fifth is promised by the SoS relaxation. The pseudoexpectation operator can equivalently be defined as a weighted expectation over a “pseudodistribution” μ satisfying similar properties (and analogous to an actual distribution). We will sometimes use the pseudodistribution view below when convenient and refer to $\tilde{\mathbb{E}}[\cdot]$ (or $\tilde{\mathbb{E}}_\mu[\cdot]$ in this case) as pseudomoments of the pseudodistribution (see [BS14b] for a more detailed exposition).

A Degree- D Sum of Squares proof of a polynomial inequality $f(x) \leq g(x)$ (where f, g are polynomials of degree at most D) is a method for ensuring the inequality continues to hold over any degree- D

pseudoexpectation. In particular, given constraints $\{q_i = 0\}_{i=1}^m$, a degree- D sum of squares proof of $f \leq g$, denoted by:

$$\{q_i = 0\}_{i=1}^m \vdash_D f \leq g,$$

is a certificate of the form $g(x) = f(x) + \sum s(x)^2 + \sum_i t(x)q_i(x)$ where all terms have degree at most D . Notice that properties 2, 3, and 4 then immediately imply $\tilde{\mathbb{E}}[f(x)] \leq \tilde{\mathbb{E}}[g(x)]$.

Conditioning is a standard algorithmic technique in the SoS paradigm used to improve the value of independently sampling a solution from the output of an SoS semidefinite relaxation (see e.g. [BRS11, BBK⁺21]). Given a degree D pseudodistribution μ , and a degree $< D$ sum of squares polynomial $s(x)$, we can define a new pseudodistribution μ' by *conditioning* on $s(x)$ as follows:

$$\tilde{\mathbb{E}}_{\mu'}[f(x)|s(x)] = \frac{\tilde{\mathbb{E}}_{\mu}[f(x)s(x)]}{\tilde{\mathbb{E}}_{\mu}[s(x)]},$$

for all polynomials $f(x)$ of degree $\leq D - \deg(s)$. We have that μ' is a valid pseudodistribution of degree $D - \deg(s)$ that satisfies the axioms satisfied by μ ²². In an algorithmic context, this is often used to restrict the pseudoexpectation to some partial solution.

10.2 The Algorithm We now present our algorithm for solving Unique Games on HD-walks. We follow the overall framework of BBKSS, which is based on the Sum-of-Squares semidefinite programming (SDP) relaxation paradigm and its view as optimizing over pseudoexpectation operators. The unique games problem (Definition 2.4) can be written as a polynomial optimization problem. In particular, given an instance I of unique games with alphabet Σ and constraints \mathcal{S} over $G(V, E)$ (that are of the form $X_u - X_v = s_{uv} \pmod{k}$), consider the following quadratic optimization problem \mathcal{A}_I over variables $\{X_{v,s}\}_{V \times \Sigma}$ that computes $\text{val}(I)$:

$$\begin{aligned} \text{Maximize:} \quad & \mathbb{E}_{(u,v) \sim E} \left[\sum_{s \in \Sigma} X_{u,s} X_{v,\pi_{uv}(s)} \right] \\ \text{Constraint to:} \quad & X_{v,s}^2 = X_{v,s} \quad \forall v \in V, s \in \Sigma \\ & X_{v,a} X_{v,b} = 0 \quad \forall v \in V, a \neq b \in \Sigma \\ & \sum_{s \in \Sigma} X_{v,s} = 1 \quad \forall v \in V \end{aligned}$$

The variables $X_{v,a}$ are 0/1 indicators that vertex $v \in V$ takes the value $a \in \Sigma$, E is a distribution over the edges of the weighted graph G , the constraints are $\pi_{uv}(s) = s - s_{uv} \pmod{k}$, and the objective function maximizes the fraction of constraints satisfied. We will work with the Degree- D Sum of Squares relaxation of this program, which outputs a degree- D pseudodistribution μ and corresponding pseudoexpectation operator $\tilde{\mathbb{E}}_{\mu} : X^{\leq D} \rightarrow \mathbb{R}$, where $X^{\leq D}$ is the set of all monomials in the X variables up to degree D , and $\tilde{\mathbb{E}}_{\mu}$ satisfies the above equality constraints as axioms. The *value* of $\tilde{\mathbb{E}}_{\mu}$, with respect to the instance I is denoted by $\text{val}_{\mu}(I) = \tilde{\mathbb{E}}_{\mu}[\text{val}_I(X)] = \tilde{\mathbb{E}}_{\mu}[\mathbb{E}_{(u,v) \sim E}[\sum_{s \in \Sigma} X_{u,s} X_{v,\pi_{uv}(s)}]]$. This operator is obtained in time $|V|^{O(D)}$ such that $\tilde{\mathbb{E}}_{\mu}[\text{val}_I(X)] \geq \text{val}(I)$.

We now describe the rounding algorithm that takes a pseudodistribution μ over an almost-satisfiable UG instance I and produces a high value assignment.

Condition&Round: We start with a basic sub-routine which will then be used in the final algorithm for HD-walks. This subroutine takes a pseudodistribution μ for an affine unique games instance $(G(V, E), \Pi)$ on alphabet Σ and outputs an assignment $x \in \Sigma^V$ via the following process:

²²Technically only the axioms that were of degree $\leq D - \deg(s)$

1. Sample a vertex $v \in V$ uniformly at random, and condition μ on event $X_{v,0} = 1$ to get the conditioned pseudodistribution $\mu|(X_{v,0} = 1)$.
2. Sample a solution $x \in \Sigma^V$ by independently sampling a label for every vertex $w \neq v$ from its marginal distribution in $\mu|(X_{v,0} = 1)$: $\Pr[x_w = s] = \tilde{\mathbb{E}}_\mu[X_{w,s}|X_{v,0} = 1]$.

Note that $\tilde{\mathbb{E}}_\mu[\cdot|X_{v,0} = 1]$ is by definition conditioning μ on the polynomial $X_{v,0}$, hence $\tilde{\mathbb{E}}_\mu[X_{w,s}|X_{v,0} = 1] = \tilde{\mathbb{E}}_\mu[X_{w,s}X_{v,0}]/\tilde{\mathbb{E}}_\mu[X_{v,0}]$. Following [BBK⁺21], we define the term Condition&Round-value (abbreviated to CR-val) of an instance I with respect to a pseudodistribution μ :

DEFINITION 10.1. (CONDITION&ROUND VALUE) *The CR-Value of the instance I with respect to a pseudodistribution μ is the expected value of the solution output by Condition&Round on instance I and pseudodistribution μ , denoted $CR\text{-Val}_\mu(I)$. We drop the subscript μ when clear from context.*

Before we describe the main algorithm, an iterative framework for applying Condition&Round, we need to introduce a simple operation on pseudodistributions for affine unique games that allows for ease of analysis:

Symmetrization: Symmetrization is an operation on pseudodistributions introduced in [BBK⁺21] to take advantage of the symmetric structure of affine unique games. The idea is to average the pseudoexpectation operator over shifts $s \in \Sigma$. Formally, given a degree D pseudodistribution μ , define the symmetrized pseudodistribution μ_{sym} via its pseudoexpectation operator as follows. For all degree $\leq D$ monomials:

$$\tilde{\mathbb{E}}_{\mu_{sym}}[X_{u_1,a_1} \dots X_{u_t,a_t}] = \frac{1}{|\Sigma|} \sum_{s \in \Sigma} \tilde{\mathbb{E}}_\mu[X_{u_1,a_1+s} \dots X_{u_t,a_t+s}].$$

We will call a pseudodistribution shift-symmetric if it is invariant under this operation. If μ is a degree D pseudodistribution that satisfies the unique games axioms \mathcal{A}_I it is easy to verify that the symmetrized pseudodistribution μ_{sym} is also a valid degree D pseudodistribution satisfying \mathcal{A}_I with $\text{val}(\mu) = \text{val}(\mu_{sym})$. Furthermore symmetrization can be performed in time subquadratic in the description of $\tilde{\mathbb{E}}_\mu$. As a result, we can perform this operation essentially for free inside our algorithm and therefore assume throughout that we are working with a shift-symmetric pseudodistribution.

We are now ready to describe the main algorithm which is called Iterated Condition&Round. The algorithm follows the strategy presented in [BBK⁺21] Algorithm 6.1], differing mainly in that the parameter $r(\varepsilon)$ satisfying their second condition has been replaced with the $(1 - O(\varepsilon))$ -ST-Rank of the underlying constraint graph.

Iterated Condition&Round: The full algorithm builds a solution by iteratively applying Condition&Round to links. Let M be a complete k -dimensional HD-walk over a d -dimensional two-sided γ -local spectral expander (X, Π) and $G_M = (X(k), E)$ be the corresponding undirected weighted graph on vertex set $X(k)$. Let $I = (M, \mathcal{S})$ be an instance of affine unique games over alphabet Σ with $\text{val}(I) \geq 1 - \varepsilon$ and $r = R_{1-16\varepsilon}(M)$. Further, given a subset $H \subset X(k)$, let I_H denote the restriction of the instance I to the subgraph vertex-induced by H . Given a subroutine for finding a link with high CR-value (see Proposition 10.1), the following process returns an $\Omega_{\varepsilon, r, k}(1)$ satisfying assignment.

1. Let $\delta(\varepsilon) := \Omega\left(\frac{\varepsilon}{\binom{k}{r}}\right)$. Solve the Degree- $D = \tilde{O}(1/\delta(\varepsilon))$ SoS SDP relaxation of unique games, and symmetrize the resulting pseudodistribution to get μ_0 . Set $j = 1$.
2. Let $\text{Dif}(j) = \tilde{\mathbb{E}}_{\mu_0}[\text{val}_I(x)] - \tilde{\mathbb{E}}_{\mu_{j-1}}[\text{val}_I(x)]$. While $\text{Dif}(j) \leq \varepsilon$:
 - (a) Find an r -link X_τ such that the CR-Value of $I|_{X_\tau}$ is at least $\delta(\varepsilon + \text{Dif}(j))$ ²³.
 - (b) Let S_j be the subgraph of X_τ induced by the vertices in $X(k)$ which have not yet been assigned a value in any partial assignment f_i , $i \leq j$, and perform Condition&Round on S_j to get partial assignment f_j .

²³When γ, k, d satisfy certain inequalities then such a link is guaranteed to exist by Proposition 10.1. Therefore we can find it by enumerating over all links $\tau \in X(r)$ and computing $\text{CR-val}(X_\tau)$

(c) Create a new pseudodistribution μ_j by making the marginal distribution over assigned vertices uniform and independent of others, i.e. for all degree $\leq D$ monomials let $\tilde{\mathbb{E}}_{\mu_j}$ be:

$$\tilde{\mathbb{E}}_{\mu_j}[X_{h_1, a_1} \dots X_{h_t, a_t} X_{u_1, b_1} \dots X_{u_m, b_m}] = \frac{1}{|\Sigma|^t} \tilde{\mathbb{E}}_{\mu_{j-1}}[X_{u_1, b_1} \dots X_{u_m, b_m}],$$

where $h_i \in S_j$ and $u_i \in X(k) \setminus S_j$. Increment $j \leftarrow j + 1$.

It is worth noting that the Condition&Round subroutine, and thus the entire Iterated Condition&Round algorithm, can be derandomized by standard techniques like the method of conditional expectations [BBK⁺21].

10.3 Analysis of Algorithm 10.2 BBKSS' analysis of Iterated Condition&Round algorithm relies heavily on analysing a quantity called the Approximate shift-partition potential defined therein. For completeness we include the definitions of the potential functions used in Appendix C.2. For the purpose of this section, the potential $\Phi_{\beta, \nu}^I(X)$ can be thought of as a low-degree polynomial (with degree $= \tilde{O}(1/\nu)$) in the variables of \mathcal{A}_I ($X = (\dots, X_{v, a}, \dots)$ where $v \in V$ and $a \in \Sigma$). Given a pseudodistribution μ for instance I , the pseudoexpectation of the potential will be denoted by $\Phi_{\beta, \nu}^I(\mu) = \tilde{\mathbb{E}}_{\mu}[\Phi_{\beta, \nu}^I(X)]$. β, ν are some parameters in $[0, 1]$ that control the degree of $\Phi_{\beta, \nu}^I(X)$ and the value of the assignment we finally obtain via Condition&Round.

The framework developed by BBKSS for analyzing Iterated Condition&Round can be described as follows. First, they prove that the Condition&Round subroutine, when run on a pseudodistribution μ for a unique games instance I , returns a high value assignment whenever $\Phi_{\beta, \nu}^I(\mu)$ (Definition C.1) is high:

THEOREM 10.2. (BBKSS THEOREM 3.3) *Let $I = (G, \mathcal{S})$ be an affine instance of Unique Games over graph $G = (V, E)$ and the alphabet Σ . Let $\beta, \nu \in [0, 1]$ and μ be a degree- $\tilde{O}(1/\nu)$ shift-symmetric pseudodistribution satisfying the unique games axioms \mathcal{A}_I (Section 10.2). If $\Phi_{\beta, \nu}^I(\mu) \geq \delta$, then on input μ , the Condition&Round Algorithm runs in time $\text{poly}(|V(G)|)$ and returns an assignment of expected value at least $(\delta - \nu)(\beta - \nu)$ for I .*

Using this rounding algorithm as a subroutine the analysis of Iterated Condition&Round in BBKSS proceeds in the following way:

1. For every $I = (G, \Pi)$, where G is a Johnson graph, using the structure theorem for non-expanding sets of G and the properties of the potential function, there exists a link X_{τ} of G such that the potential induced on C is large: $\Phi_{\beta, \nu}^{I|X_{\tau}}(\mu) \geq \text{poly}(\delta)$, where $I|_{X_{\tau}}$ denotes the UG instance induced on the subgraph X_{τ} . Using the rounding theorem 10.2 we get that the Condition&Round value on this link is high.
2. Iterative Condition&Round: Using iterations one can patch together solutions on different links, and since each link is a non-expanding set we get a good solution for the whole graph.

We show that this framework extends to all UG instances on HD-walks. We use the analysis of Condition&Round (Theorem 10.2) in a blackbox way. We start with the analysis of Point 1, which relies on the technical machinery developed in the previous sections for analysing the non-expanding sets of HD-walks. Using this, we show that there always exists a link with high potential and therefore with high Condition&Round value.

PROPOSITION 10.1. *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -spectral expander with $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ and $d > k$, and I be an affine unique games instance over M with value at least $1 - \eta$ where $1/2^k \leq \eta < .02$. Let $r = r(\eta) = R_{1-16\eta}(M) - 1$, and assume $r \leq k/2$. Then given a degree- $\tilde{O}\left(\frac{1}{\eta} \binom{k}{r}\right)$ pseudodistribution satisfying the axioms \mathcal{A}_I , we can find in time $|X(k)|^4$ an r -link X_{τ} with $\text{CR-Value}(X_{\tau}) \geq \Omega\left(\frac{\eta}{\binom{k}{r}}\right)$.*

The proof of this lemma is similar in nature to the proof of the analogous lemma in BBKSS, albeit with all of our technical machinery is plugged in. In particular we use a sum-of-squares version of Theorem 8.1, the relation between eigenvalues and expansion of links (Theorem 9.1), and a new property analyzing the expansion of vertices within a link (Lemma 10.4). We elaborate on this in Section 10.4.

Given the ability to find a link with high CR-value, we can use BBKSS analysis of Iterated Condition&Round ([BBK⁺21, Lemma 6.12]) as a blackbox to conclude it produces an assignment satisfying a large fraction of the constraints for the *whole graph*.

LEMMA 10.1. (LEMMA 6.12 [BBK⁺21]) *Let $\varepsilon \in (0, .01)$, $\delta : [0, 1] \rightarrow [0, 1]$ be any function, and $\delta_{\min} = \min_{\delta(\eta) \in [\varepsilon, 2\varepsilon]} \delta(\eta)$. Let G be a weighted graph²⁴ and I be any affine unique games instance on G with alphabet size $|\Sigma| \geq \Omega\left(\frac{1}{\delta_{\min}}\right)$ and value at least $1 - \varepsilon$.*

Suppose we have a subroutine \mathcal{A} which, for any $\varepsilon \leq \eta \leq 2\varepsilon$, given as input a shift-symmetric degree- D pseudodistribution μ satisfying \mathcal{A}_I with $\mathbb{E}_\mu[\text{val}_I(x)] \geq 1 - \eta$ returns a vertex-induced subgraph H such that:

1. *The CR-Value of I_H is at least $\delta(\eta)$.*
2. *The expansion of H is $O(\eta)$.*

Then if \mathcal{A} runs in time $T(\mathcal{A})$, Iterated Condition&Round²⁵ outputs a solution for I satisfying an $\Omega(\delta_{\min}^2 \varepsilon)$ -fraction of edges of G in time $|V(G)|(T(\mathcal{A}) + |V(G)|^{O(D)})$.

With these results in hand, the final observation to prove Theorem 10.1 lies in the relation between the local expansion of links and global spectra of eigenstrips we proved in Section 9. Namely, since the r th strip is by definition the last one with eigenvalue worse than $1 - O(\varepsilon)$, the expansion of r -links is at most $O(\varepsilon)$ by Theorem 9.1. We therefore meet the conditions of Lemma 10.1 and can use Iterated Condition&Round to output a good global assignment. We now prove Theorem 10.1 assuming Proposition 10.1 and Lemma 10.1 hold.

Proof. [Proof of Theorem 10.1] To start, we first prove that we may assume without loss of generality both that $\varepsilon \geq 1/2^k$, and that $r(2\varepsilon) \leq k/2$ (we will need these properties to meet the conditions of Proposition 10.1 and Lemma 10.1). This relies on the following claim, the proof of which we defer to Appendix C.4.

CLAIM 10.1. *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -local-spectral expander with $d > k$ and $\gamma \leq w(M)^{-1} 2^{-\Omega(h(M)+k)}$. If the expected laziness of M (i.e. $\mathbb{E}_{\Pi_k}[\mathbf{1}_v^T M \mathbf{1}_v]$) is at least .1, then the spectral gap of M is at least $\Omega(1/k)$ and $\lambda_{k/2}(M) \leq .68$.*

With this in mind, note that we can always assume the expected laziness $\mathbb{E}[\mathbf{1}_v^T M \mathbf{1}_v]$ is at most 1/10. This follows from the fact that the expected laziness of M exactly corresponds to the probability of drawing a self-edge on the corresponding graph G_M (see Appendix C.1), and in an affine unique game, every self edge is either always or never satisfiable. Since our game has value at least .99, in a walk with more than .1 laziness, at least .09 of the weight on self-edges must be satisfiable, and therefore *every* assignment will satisfy our approximation guarantee. Assuming then that $\mathbb{E}[\mathbf{1}_v^T M \mathbf{1}_v] \leq .1$, Claim 10.1 implies that the spectral gap of M is at least $\Omega(1/k)$, in which case standard algorithms for unique games on expander graphs (e.g. [MM10]) give the desired result. Similarly, we have $\lambda_{k/2} \leq .68 \leq 1 - 32\varepsilon$ by assumption, so we may further assume $r(2\varepsilon) \leq k/2$ as desired.

Now that we have $1/2^k \leq \varepsilon < .02$ and $r \leq k/2$, we are in position to solve the degree- D SoS relaxation of the unique games integer program \mathcal{A}_I for $D = \tilde{O}\left(\frac{1}{\varepsilon} \binom{k}{r(2\varepsilon)}\right)$, and apply Proposition 10.1 to build a sub-route that satisfies Lemma 10.1. Namely, we have that for any $\varepsilon \leq \eta \leq 2\varepsilon$ and pseudodistribution of value at least $1 - \eta$, Proposition 10.1 finds a link X_τ with high Condition&Round value:

$$\text{CR-val}(X_\tau) \geq \Omega\left(\frac{\eta}{\binom{k}{r(\eta)}}\right).$$

Further, note that by our assumptions on γ and the fact that M is complete, Proposition 7.3 gives that the eigenvalues corresponding to each eigenstrip strictly decrease, and therefore by definition of ST-Rank that the approximate eigenvalue corresponding to the r th eigenstrip is at least $1 - O(\eta)$. By Theorem 9.1, this implies that X_τ has poor expansion:

$$\Phi(X_\tau) \leq O(\eta) + w(M)h(M)^2 2^{O(k)} \gamma \leq O(\eta)$$

since $\eta \geq \varepsilon \geq w(M)2^{O(h(M)+k)}\gamma$ by assumption. As a result, this sub-routine satisfies the conditions of Lemma 10.1 with δ set to $\delta(\eta) = \Omega\left(\frac{\eta}{\binom{k}{r(\eta)}}\right)$. The only catch is that we need the alphabet Σ to satisfy $|\Sigma| \leq \Omega\left(\frac{1}{\delta_{\min}}\right)$. This

²⁴BBKSS only state the result for regular graphs, but their proof works for weighted graphs too, where the value, expansion etc are measured appropriately according to the edge-weights.

²⁵We note that the version of Iterated Condition&Round analyzed in [BBK⁺21] differs slightly from the one we present. Our version simplifies their algorithm a bit but the analysis is exactly the same.

can be assumed without loss of generality, since a random solution satisfies a $1/|\Sigma|$ fraction of constraints in expectation (and can be easily derandomized), which satisfies our approximation guarantee if $|\Sigma| \leq O\left(\frac{1}{\delta_{\min}}\right)$. As a result, applying Lemma 10.1 outputs a $\Omega(\delta_{\min}^2 \varepsilon)$ -satisfying solution for $\delta_{\min} = \min_{\eta \in [\varepsilon, 2\varepsilon]} \left\{ \frac{\varepsilon}{(r(\eta))} \right\} \geq \Omega\left(\frac{\varepsilon}{(r(2\varepsilon))}\right)$. The running time guarantee follows from noting that the sub-routine promised by Proposition 10.1 runs in time $|V(G)|^4$, so Lemma 10.1 then runs in time $|V(G)|^{O(D)}$ with $|V(G)| = |X(k)|$. Since solving the original Degree- D SoS relaxation also takes time only $|V(G)|^{O(D)}$, we get the desired result. \square

10.4 Proof of Proposition 10.1 Now that we have given the algorithmic background, we prove the main technical lemma behind the analysis. The broad outline of the proof is as follows:

1. We first prove a structure theorem (Theorem 10.3) for non-expanding sets of HD-walks. We show an SoS proof of the fact that every non-expanding set must have large variance of size when restricted to links of the complex.
2. Using the structure theorem, in Proposition 10.2 we show that given a pseudodistribution μ with objective value $1 - \eta$ for unique games over an HD-walk M , one can find a link X_τ with high *global* shift-partition potential (Definition C.2).
3. In the final step (Lemma 10.3), we relate the global shift-partition potential to the shift-partition potential on the subgraph induced by X_τ ²⁶: we show that $\Phi_{\beta, \nu}^{I_\tau}(\mu)$ (Definition C.3) is large. By the rounding theorem (Theorem 10.2), we then conclude that the expected value of the Condition&Round algorithm, when performed on X_τ ($\text{CR-val}_\mu(X_\tau)$) must be high.

This proof structure is similar to the analogous Lemma 6.9 of [BBK⁺21]. Most of our technical work goes into proving points 1 and 3 above, and point 2 turns out to be straightforward given BBKSS.

We now turn to proving our structure theorem in the low-degree SoS proof system. We will reprove the lemmas for expansion for pseudorandom sets in HD-walks proved in Section 8 and Section 9, but this time carefully making sure that they are in the low-degree SoS proof system.

THEOREM 10.3. (SOS STRUCTURE THEOREM FOR HD-WALKS) *Let (X, Π) be a two-sided γ -local-spectral expander and M be a k -dimensional, complete HD-walk on X with $\gamma \leq w(M)^{-1}h(M)^{-2}2^{-\Omega(k)}$. Then for any $f \in C_k$ and any $0 \leq \ell \leq k/2$, the expansion²⁷ of f with respect to M is large:*

$$\vdash_2 \langle f, (I - M)f \rangle \geq (1 - \lambda_{\ell+1}) \left((1 - c_1\gamma)(\mathbb{E}[f] + B(f)) - \binom{k}{\ell} \langle D_k^\ell f, D_k^\ell f \rangle \right),$$

where $B(f) = \mathbb{E}[f^2 - f]$ represents the Booleanity constraints and $c_1 \leq 2^{O(k)}$.

The proof of Theorem 10.3 relies on a number of SoS variants of properties of the HD-Level-Set Decomposition used in the previous sections, namely approximate orthogonality and the relation between $\|f_i\|$ and $\|g_i\|$.

LEMMA 10.2. *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $f \in C_k$. Then for $f_i = U_i^k g_i \in V_k^i$:*

$$\vdash_2 \langle f_i, f_i \rangle \in \left(\frac{1}{\binom{k}{i}} \pm c_1\gamma \right) \langle g_i, g_i \rangle,$$

where $c_1 \leq O(k^2)$. Further, weak variants of approximate orthogonality also have degree 2 proofs:

$$\vdash_2 \langle f_i, f_i \rangle \leq (1 + c_2\gamma) \langle f, f \rangle,$$

and

$$\vdash_2 \langle f, f_i \rangle \geq -c_3\gamma \langle f, f \rangle,$$

where $c_2, c_3 \leq 2^{O(k)}$

²⁶This is just the potential measured on the sub-instance I when restricted to the induced subgraph given by the link X_τ .

²⁷This notion of expansion varies slightly from our previous definition, coinciding (up to normalization) for Boolean functions.

The proof of Lemma 10.2 follows from combining arguments in [DDFH18] with standard SoS tricks. For completeness, we give the proof in Appendix C.3. With Lemma 10.2 in hand, Theorem 10.3 follows similarly to Theorem 8.2 and Theorem 9.2.

Proof. [Proof of Theorem 10.3] As in the proof of Theorem 9.2, it is sufficient to bound the weight of f onto low levels of the HD-Level-Set decomposition to prove that the expansion of f is small. First, note that given the function f , the low-level decomposition functions f_i 's are explicit linear functions of the coordinates of f [DDFH18]. We will thus show that Lemma 8.1 has an SoS proof, that is the following relation between $\|D_k^\ell f\|^2$ and f 's weight holds:

$$(10.9) \quad \vdash_2 \sum_{j=0}^{\ell} \langle f, f_j \rangle \leq \binom{k}{\ell} \langle D_k^\ell f, D_k^\ell f \rangle + 2^{O(k)} \gamma \langle f, f \rangle.$$

Before proving Equation (10.9), we check that if the above inequality holds then the theorem statement follows. To see this, first recall from Section 7 that $Mf_i = \lambda_i f_i + \Gamma g_i$ where Γ is a matrix with spectral norm $\leq w(M)h(M)^2 2^{O(k)} \gamma$ which we call C for convenience. It will be useful to have an SoS upper bound on $\langle f, \Gamma g_i \rangle$, which we will use multiple times throughout this proof. First by an SoS version of Cauchy-Schwarz we get that for all real constants $\zeta > 0$:

$$\vdash_2 \langle f, \Gamma g_i \rangle \leq \frac{\zeta}{2} \langle f, f \rangle + \frac{1}{2\zeta} \langle \Gamma g_i, \Gamma g_i \rangle.$$

Now substituting $\zeta = C$ and simplifying using the spectral norm bounds on Γ and Lemma 10.2 we get:

$$\begin{aligned} \langle f, \Gamma g_i \rangle &\leq \frac{C}{2} \langle f, f \rangle + \frac{1}{2C} \langle \Gamma g_i, \Gamma g_i \rangle \\ &\leq \frac{C}{2} \langle f, f \rangle + \frac{C}{2} \langle g_i, g_i \rangle \\ &\leq \frac{C}{2} \langle f, f \rangle + \frac{C}{2} \left(\binom{k}{i} + O(\gamma) \right) \langle f_i, f_i \rangle, \\ &\leq c_1 \gamma \langle f, f \rangle, \end{aligned}$$

where all inequalities are degree 2 SoS inequalities and $c_1 \leq w(M)h(M)^2 2^{O(k)}$. With this in hand and assuming Equation (10.9), the result follows from expanding out $\langle f, (I - M)f \rangle$ and applying Lemma 10.2:

$$\begin{aligned} \langle f, (I - M)f \rangle &= \langle f, f \rangle - \sum_{i=0}^k \lambda_i \langle f, f_i \rangle - \sum_{i=0}^k \langle f, \Gamma g_i \rangle \\ &\geq (1 - c_1 \gamma) \langle f, f \rangle - \sum_{i=0}^{\ell} \lambda_i \langle f, f_i \rangle - \sum_{i=\ell+1}^k \lambda_i \langle f, f_i \rangle \\ &\geq (1 - c_2 \gamma) \langle f, f \rangle - \sum_{i=0}^{\ell} \langle f, f_i \rangle - \lambda_{\ell+1} \sum_{i=\ell+1}^k \langle f, f_i \rangle \\ &= (1 - c_2 \gamma) \langle f, f \rangle - \sum_{i=0}^{\ell} \langle f, f_i \rangle - \lambda_{\ell+1} \langle f, f \rangle + \lambda_{\ell+1} \sum_{i=0}^{\ell} \langle f, f_i \rangle \\ &= (1 - \lambda_{\ell+1}) \left((1 - c_3 \gamma) \langle f, f \rangle - \sum_{i=0}^{\ell} \langle f, f_i \rangle \right) \\ &\geq (1 - \lambda_{\ell+1}) \left((1 - c_4 \gamma) \langle f, f \rangle - \binom{k}{\ell} \langle D_k^\ell f, D_k^\ell f \rangle \right) \\ &= (1 - \lambda_{\ell+1}) \left((1 - c_4 \gamma) \mathbb{E}[f] + (1 - c_4 \gamma) B(f) - \binom{k}{\ell} \langle D_k^\ell f, D_k^\ell f \rangle \right) \end{aligned}$$

where $c_1, c_2, c_3, c_4 \leq w(M)h(M)^2 2^{O(k)}$, and $B(f) = \mathbb{E}[f^2 - f]$.

It remains to prove Equation 10.9. This follows from a similar modification of Lemma 8.1. Notice that by the adjointness of D and U , it is enough to analyze the walk $U_\ell^k D_\ell^k$:

$$\langle D_\ell^k f, D_\ell^k f \rangle = \langle f, U_\ell^k D_\ell^k f \rangle$$

Using Proposition 7.1 and the assumption $\ell \leq k/2$, we can decompose the righthand side as:

$$\binom{k}{\ell} \langle f, U_\ell^k D_\ell^k f \rangle = \sum_{j=0}^{\ell} \binom{k-j}{\ell-j} \langle f, f_j \rangle + \sum_{j=0}^k \langle f, \Gamma g_j \rangle$$

where $\|\Gamma\| \leq w(M)h(M)^2 2^{O(k)} \gamma := C$. Noting that $\binom{k-j}{\ell-j}$ is at least 1 for $0 \leq j \leq \ell$, we can apply the upper bound on $\langle f, \Gamma g_j \rangle$ proved above and Lemma 10.2 to get:

$$\begin{aligned} \binom{k}{\ell} \langle D_\ell^k f, D_\ell^k f \rangle &\geq \sum_{j=0}^{\ell} \langle f, f_j \rangle - c_5 \gamma \langle f, f \rangle - \sum_{j=0}^k \langle f, \Gamma g_j \rangle \\ &\geq \sum_{j=0}^{\ell} \langle f, f_j \rangle - c_6 \gamma \langle f, f \rangle \end{aligned}$$

where all the constants c_i 's are less than $w(M)h(M)^2 2^{O(k)}$ and all inequalities are degree-2 SoS. \square

It is worth giving a brief qualitative comparison of this result to a similar version for the Johnson graphs in [BBK⁺21]. In particular, Theorem 10.3 not only gives a tighter bound (by a factor of $\exp(r)$), but perhaps more importantly shows how viewing the problem from the framework of high dimensional expansion demystifies the original Fourier analytic proof. This understanding allows us to extend the structural result well beyond the Johnson graphs to all HD-walks. In fact, this result also holds for the more general class of expanding posets [DDFH18] (albeit with different parameters). We leave their discussion to future work.

Given this structure theorem, we use the BBKSS' framework to show that the global potential (Definition C.2) on a link is high. This follows because of the properties of the potential function and the pseudodistribution over $(1 - \eta)$ -satisfying assignments. In general, the potential function $\Phi_{\beta,\nu}^I(\mu)$ turns out to be an average taken over non-expanding partitions of the graph. Using our structure theorem, we get that every non-expanding partition must have a large variance across links, and therefore there must exist a link X_τ such that the partition is large even when restricted to X_τ . As a result, there must be a link X_τ where the partitions are large on average too, which corresponds to a quantity called the global potential restricted to X_τ , denoted by $\Phi_{\beta,\nu}^I(\mu|_\tau)$, being large. The proof is essentially the same as [BBK⁺21, Lemma 6.9], but we include it in Appendix C.5 for completeness.

PROPOSITION 10.2. *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -local-spectral expander satisfying $\gamma w(M)2^{O(h(M)+k)} \leq 1/2^k < \eta < 0.02$ and $d > k$, and I be an affine unique games instance over M with value at least $1 - \eta$. Let $r(\eta) = R_{1-16\eta}(M) - 1$, and assume $r \leq k/2$. Let $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$.*

Then given a degree- $\tilde{O}(1/\nu)$ pseudodistribution μ satisfying the axioms \mathcal{A}_I , there exists an r -link X_τ such that the global potential restricted to X_τ is large:

$$\Phi_{\beta,\nu}^I(\mu|_\tau) \geq \frac{1}{4\binom{k}{r}}.$$

In the next lemma, we will relate the global potential to the potential induced on X_τ . In particular, we'll show that since the global potential on X_τ is high (Proposition 10.2), the potential induced on X_τ , $\Phi_{\beta,\nu}^{I_\tau}(\mu)$, must also be high. Since this is just the usual potential function when applied to the sub instance $I|_\tau$ corresponding to the subgraph induced by X_τ (Definition C.3), we can then apply the rounding theorem of BBKSS Theorem 10.2 to surmise that the CR-value of X_τ is high.

LEMMA 10.3. *Assume the conditions of Proposition 10.2 hold. Let X_τ be an r -link for $r = R_{1-16\eta}(M) - 1$ with high global potential: $\Phi_{\beta,\nu}^I(\mu|_\tau) \geq \frac{1}{4\binom{k}{r}}$, for $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$. Then the potential induced on X_τ is also high:*

$$\Phi_{\eta,\nu}^{I|_\tau}(\mu) \geq \frac{1}{8\binom{k}{r}}.$$

We prove Lemma 10.3 in Appendix C.5, but it's worth pausing to discuss the proof, especially in how it differs from the analogous result in BBKSS [BBK⁺21, Claim 6.11]. Recall that the main idea behind Lemma 10.3 is to relate the global potential on X_τ (bounded in Proposition 10.2) to the potential induced on X_τ itself (which implies high CR-Value by Theorem 10.2). In a bit more detail, the global potential on a link depends on the “value” of the vertices in the link (which measures the value of an assignment at the vertex), but is measured with respect to its neighbors across the *entire graph*. On the other hand, the induced potential on the link depends on the value of vertices when measured only with respect to the edges *inside* the link (see Definitions C.2 C.3 for exact details). It is possible to relate the two potentials by observing that the internal-value of a vertex (value with respect to only the neighbors inside the link) can decrease by at most an additive factor equal to its edge-expansion inside the link (i.e. the fraction of edges incident on the vertex that leave the set). By leveraging machinery developed in Section 7 and some additional properties, we show not only that the expansion of links is small (by Theorem 9.1), but in fact that this holds approximately vertex by vertex: for most vertices in the link, only an $O(\eta)$ -fraction of their edges are outgoing. We prove the following claim formally in Appendix C.5

LEMMA 10.4. *Let M be a k -dimensional HD-walk on d -dimensional two-sided γ -local-spectral expander satisfying $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ and $d > k$. Then for every i -link X_τ , the deviation of the random variable $\phi_{X_\tau}(v)$ ($v \sim X_\tau$) is small:*

$$\mathbb{E}_{v \sim X_\tau} [|\phi_{X_\tau}(v) - \phi(X_\tau)|] \leq \frac{1}{2^{11k}},$$

where $\phi_{X_\tau}(v)$ denotes the fraction of edges incident on $v \in X_\tau$ that leave X_τ .

BBKSS use an exact version of this statement for the Johnson (that expansion holds vertex-by-vertex) to prove their analogous version of Lemma 10.3. We relax the conditions required for their proof and show that the approximate statement state above is sufficient in Appendix C.5.

Finally, we complete the section by using Proposition 10.2 and Lemma 10.3 to prove Proposition 10.1 (and thereby Theorem 10.1 as well). We restate the proposition here for convenience:

PROPOSITION 10.3. (RESTATEMENT OF PROPOSITION 10.1) *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -spectral expander with $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ and $d > k$, and I be an affine unique games instance over M with value at least $1 - \eta$ where $1/2^k \leq \eta < .02$. Let $r = r(\eta) = R_{1-16\eta}(M) - 1$, and assume $r \leq k/2$. Then given a degree- $\tilde{O}\left(\frac{1}{\eta} \binom{k}{r}\right)$ pseudodistribution μ satisfying the axioms \mathcal{A}_I , we can find in time $|X(k)|^4$ an r -link X_τ with $CR\text{-val}_\mu(X_\tau) \geq \Omega\left(\frac{\eta}{\binom{k}{r}}\right)$.*

Proof. Let $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$ as in Proposition 10.2. By Proposition 10.2, there exists an r -link X_τ with high global potential:

$$\Phi_{\beta, \nu}^I(\mu|_\tau) \geq \frac{1}{4\binom{k}{r}}.$$

By Lemma 10.3, this implies that X_τ also has high induced potential:

$$\Phi_{\eta, \nu}^{I|_\tau}(\mu) \geq \frac{1}{8\binom{k}{r}}.$$

As mentioned previously, this is just the usual potential function on the instance induced by X_τ (Definition C.3), so we may apply the rounding theorem of BBKSS Theorem 10.2 to bound X_τ ’s CR-Value. In particular, since we have set ν such that $(\eta - \nu) \geq \Omega(\eta)$ and $(\Phi_{\eta, \nu}^{I|_\tau}(\mu) - \nu) \geq \Omega\left(\frac{1}{\binom{k}{r}}\right)$, Theorem 10.2 implies that rounding μ , which is a degree $\tilde{O}(1/\nu) = \tilde{O}\left(\frac{1}{\eta} \binom{k}{r}\right)$ pseudodistribution as required, would give an assignment with large -value on X_τ :

$$CR\text{-val}_\mu(X_\tau) \geq (\eta - \nu)(\Phi_{\eta, \nu}^{I|_\tau}(\mu) - \nu) \geq \Omega\left(\frac{\eta}{\binom{k}{r}}\right).$$

Now that we have shown that there exists an r -link X_τ with large CR-value, it remains to show that we can efficiently find such a link. We can do a brute-force enumeration over all r -links in time $|X(r)|$, compute every CR-value in time $|X(r)|^3$, and therefore in time $|X(r)|^4$ find a link with large CR-value. Since $|X(k)| \geq |X(r)|$ is a standard consequence of X being a local-spectral expander [DDFH18] the proposition follows. \square

11 Acknowledgements

This work stemmed in part from collaboration at the Simons Institute of Theory of Computing, Berkeley during the 2019 summer cluster: “Error-Correcting Codes and High-Dimensional Expansion”. We thank the Simons institute for their hospitality and the organizers of the cluster for creating such an opportunity. Further, the authors would like to thank Sankeerth Rao and Yotam Dikstein for useful discussions in the early stages of this work, and Ella Sharakanski for discussion regarding HD-walk decompositions. We thank Madhu Sudan and Sam Hopkins for helpful comments on initial drafts of this work. Finally, we additionally owe many thanks to Sam Hopkins for discussions on the Sum of Squares framework, and to Vedat Alev for his insights on the spectral structure of higher order random walks.

References

- [ABS15] Sanjeev Arora, Boaz Barak, and David Steurer, *Subexponential algorithms for unique games and related problems*, Journal of the ACM (JACM) **62** (2015), no. 5, 1–25.
- [AJT19] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani, *Approximating constraint satisfaction problems on high-dimensional expanders*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 180–201.
- [AKK⁺08] Sanjeev Arora, Subhash A Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K Vishnoi, *Unique games on expanding constraint graphs are easy*, Proceedings of the fortieth annual ACM symposium on Theory of computing, 2008, pp. 21–28.
- [AKKT15] Naman Agarwal, Guy Kindler, Alexandra Kolla, and Luca Trevisan, *Unique games on the hypercube*, Chicago J. Theor. Comput. Sci. **2015** (2015).
- [AL20] Vedat Levi Alev and Lap Chi Lau, *Improved analysis of higher order random walks and applications*, arXiv preprint arXiv:2001.02827 (2020).
- [Ale20] Vedat Alev, *Higher order random walks, local spectral expansion, and applications.*, Ph.D. thesis, University of Waterloo, 2020.
- [ALO20] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, arXiv preprint arXiv:2001.00303 (2020).
- [ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant, *Log-concave polynomials ii: high-dimensional walks and an fpras for counting bases of a matroid*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 1–12.
- [BBH⁺12] Boaz Barak, Fernando GSL Brandao, Aram W Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou, *Hypercontractivity, sum-of-squares proofs, and their applications*, Proceedings of the forty-fourth annual ACM symposium on Theory of computing, 2012, pp. 307–326.
- [BBK⁺21] Mitali Bafna, Boaz Barak, Pravesh K. Kothari, Tselil Schramm, and David Steurer, *Playing unique games on certified small-set expanders*, STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021 (Samir Khuller and Virginia Vassilevska Williams, eds.), ACM, 2021, pp. 1629–1642.
- [BCC⁺21] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda, *On mixing of markov chains: Coupling, spectral independence, and entropy factorization*, arXiv preprint arXiv:2103.07459 (2021).

- [BKS18] Boaz Barak, Pravesh K Kothari, and David Steurer, *Small-set expansion in shortcode graph and the 2-to-2 conjecture*, arXiv preprint arXiv:1804.08662 (2018).
- [BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer, *Rounding semidefinite programming hierarchies via global correlation*, 2011 ieee 52nd annual symposium on foundations of computer science, IEEE, 2011, pp. 472–481.
- [BS14a] Boaz Barak and David Steurer, *Sum-of-squares proofs and the quest toward optimal algorithms*, arXiv preprint arXiv:1404.5236 (2014).
- [BS14b] _____, *Sum-of-squares proofs and the quest toward optimal algorithms*, arXiv preprint arXiv:1404.5236 (2014).
- [CGSV21] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda, *Rapid mixing for colorings via spectral independence*, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 1548–1557.
- [CJHJ11] Radha Chitta, Rong Jin, Timothy C Havens, and Anil K Jain, *Approximate kernel k-means: Solution to large scale kernel clustering*, Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, 2011, pp. 895–903.
- [CLV20] Zongchen Chen, Kuikui Liu, and Eric Vigoda, *Rapid mixing of glauber dynamics up to uniqueness via contraction*, arXiv preprint arXiv:2004.09083 (2020).
- [CLV21] _____, *Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1537–1550.
- [DD19] Yotam Dikstein and Irit Dinur, *Agreement testing theorems on layered set systems*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 1495–1524.
- [DDFH18] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha, *Boolean function analysis on high-dimensional expanders*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [DDHRZ20] Yotam Dikstein, Irit Dinur, Prahladh Harsha, and Noga Ron-Zewi, *Locally testable codes via high-dimensional expanders*, arXiv preprint arXiv:2005.01045 (2020).
- [DFHT20] Irit Dinur, Yuval Filmus, Prahladh Harsha, and Madhur Tulsiani, *Explicit sos lower bounds from high-dimensional expanders*, arXiv preprint arXiv:2009.05218 (2020).
- [DHK⁺19] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta Shma, *List decoding with double samplers*, Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2019, pp. 2134–2153.
- [DHK⁺21] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta-Shma, *List-decoding with double samplers*, SIAM Journal on Computing **50** (2021), no. 2, 301–349.
- [DK17] Irit Dinur and Tali Kaufman, *High dimensional expanders imply agreement expanders*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 974–985.
- [DKK⁺18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra, *On non-optimally expanding sets in grassmann graphs*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 940–951.
- [DKK⁺18b] _____, *Towards a proof of the 2-to-1 games conjecture?*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 376–389.

[FGYZ21] Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang, *Rapid mixing from spectral independence beyond the boolean domain*, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 1558–1577.

[FKLM20] Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer, *Hypercontractivity on the symmetric group*, arXiv preprint arXiv:2009.05503 (2020).

[FKP19] Noah Fleming, Pravesh Kothari, and Toniann Pitassi, *Semialgebraic proofs and efficient algorithm design*, Foundations and Trends® in Theoretical Computer Science **14** (2019), no. 1-2, 1–221.

[GMR08] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra, *Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph*, 2008 49th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2008, pp. 573–582.

[GS11] Venkatesan Guruswami and Ali Kemal Sinop, *Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with psd objectives*, 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, IEEE, 2011, pp. 482–491.

[HST20] Samuel B Hopkins, Tselil Schramm, and Luca Trevisan, *Subexponential lps approximate max-cut*, 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2020, pp. 943–953.

[JPV21] Vishesh Jain, Huy Tuan Pham, and Thuy Duong Vuong, *Spectral independence, coupling with the stationary distribution, and the spectral gap of the glauber dynamics*, arXiv preprint arXiv:2105.01201 (2021).

[JQST20] Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani, *Unique decoding of explicit epsilon-balanced codes near the gilbert-varshamov bound*, arXiv preprint arXiv:2011.05500 (2020).

[JST21] Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani, *Near-linear time decoding of ta-shma's codes via splittable regularity*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1527–1536.

[Kho02] Subhash Khot, *On the power of unique 2-prover 1-round games*, Proceedings of the thiry-fourth annual ACM symposium on Theory of computing, 2002, pp. 767–775.

[KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell, *Optimal inapproximability results for max-cut and other 2-variable csp's?*, SIAM Journal on Computing **37** (2007), no. 1, 319–357.

[KLLM19] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer, *Hypercontractivity for global functions and sharp thresholds*, arXiv preprint arXiv:1906.05568 (2019).

[KM16a] Tali Kaufman and David Mass, *High dimensional combinatorial random walks and colorful expansion*, arXiv preprint arXiv:1604.02947 (2016).

[KM16b] Subhash Khot and Dana Moshkovitz, *Candidate hard unique game*, Proceedings of the forty-eighth annual ACM symposium on Theory of Computing, 2016, pp. 63–76.

[KM20] Tali Kaufman and David Mass, *Local-to-global agreement expansion via the variance method*, 11th Innovations in Theoretical Computer Science Conference (ITCS 2020), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

[KMM11] Alexandra Kolla, Konstantin Makarychev, and Yury Makarychev, *How to play unique games against a semi-random adversary: Study of semi-random models of unique games*, IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22–25, 2011, 2011, pp. 443–452.

[KMMS18] Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra, *Small set expansion in the johnson graph.*, Electronic Colloquium on Computational Complexity (ECCC), vol. 25, 2018, p. 78.

[KMS17] Subhash Khot, Dor Minzer, and Muli Safra, *On independent sets, 2-to-2 games, and grassmann graphs*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 576–589.

[KMS18] ———, *Pseudorandom sets in grassmann graph have near-perfect expansion*, 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2018, pp. 592–601.

[KO18] Tali Kaufman and Izhar Oppenheim, *Construction of new local spectral high dimensional expanders*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 773–786.

[KO20] ———, *High order random walks: Beyond spectral gap*, Combinatorica (2020), 1–37.

[Kol11] Alexandra Kolla, *Spectral algorithms for unique games*, computational complexity **20** (2011), no. 2, 177–206.

[KR08] Subhash Khot and Oded Regev, *Vertex cover might be hard to approximate to within $2 - \varepsilon$* , Journal of Computer and System Sciences **74** (2008), no. 3, 335–349.

[KS20] Tali Kaufman and Ella Sharakanski, *Chernoff bound for high-dimensional expanders*, To Appear APPROX/RANDOM 2020, 2020.

[KTW14] Subhash Khot, Madhur Tulsiani, and Pratik Worah, *A characterization of strong approximation resistance*, Proceedings of the forty-sixth annual ACM symposium on Theory of computing, 2014, pp. 634–643.

[KV15] Subhash A Khot and Nisheeth K Vishnoi, *The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into l_1* , Journal of the ACM (JACM) **62** (2015), no. 1, 1–39.

[Liu21] Kuikui Liu, *From coupling to spectral independence and blackbox comparison with the down-up walk*, arXiv preprint arXiv:2103.11609 (2021).

[LSV05] Alexander Lubotzky, Beth Samuels, and Uzi Vishne, *Explicit constructions of ramanujan complexes of type ad*, European Journal of Combinatorics **26** (2005), no. 6, 965–993.

[MM10] Konstantin Makarychev and Yury Makarychev, *How to play unique games on expanders*, International Workshop on Approximation and Online Algorithms, Springer, 2010, pp. 190–200.

[OT13] Shayan Oveis Gharan and Luca Trevisan, *A new regularity lemma and faster approximation algorithms for low threshold rank graphs*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Springer, 2013, pp. 303–316.

[Rag08] Prasad Raghavendra, *Optimal algorithms and inapproximability results for every csp?*, Proceedings of the fortieth annual ACM symposium on Theory of computing, 2008, pp. 245–254.

[RSS18] Prasad Raghavendra, Tselil Schramm, and David Steurer, *High-dimensional estimation via sum-of-squares proofs*, pp. 3389–3423, World Scientific, 2018.

[Zha20] Ruizhe Zhang, Personal communication, December 2020.

A Proof of Lemma 7.1

In this section, we prove a strengthening of the main technical lemma of DDFH Section 8 [DDFH18, Claim 8.8], which allows for better control of error propagation.

LEMMA A.1. (STRENGTHENED CLAIM 8.8 [DDFH18]) *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander. Then for all $j < k < d$:*

$$D_{k+1}U_{k-j}^{k+1} - \frac{j+1}{k+1}U_{k-j}^k - \frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} = \sum_{i=-1}^{j-1} \frac{k-i}{k+1} U_{k-1-i}^k \Gamma_i U_{k-j}^{k-1-i}$$

where $\|\Gamma_i\| \leq \gamma$.

Proof. The proof follows by a simple induction. The base cases, $j = 0$ and $k < d$, follow immediately from Equation (7.6). For the inductive step, consider:

$$\begin{aligned} D^{k+1}U_{k-(j+1)}^{k+1} &= \left(D^{k+1}U_{k-j}^{k+1} - \frac{j+1}{k+1}U_{k-j}^k - \frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} \right) U_{k-j-1} \\ &\quad + \frac{j+1}{k+1}U_{k-j-1}^k + \frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} U_{k-j-1} \end{aligned}$$

By the inductive hypothesis, the first term on the RHS may be written as:

$$\left(D^{k+1}U_{k-j}^{k+1} - \frac{j+1}{k+1}U_{k-j}^k - \frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} \right) U_{k-j-1} = \sum_{i=-1}^{j-1} \frac{k-i}{k+1} U_{k-i-1}^k \Gamma_i U_{k-j-1}^{k-1-i},$$

where $\|\Gamma_i\| \leq \gamma$. For the latter term, consider flipping DU and UD . By Equation (7.6) we have:

$$\frac{k-j}{k+1}U_{k-j-1}^k D_{k-j} U_{k-j-1} = U_{k-j-1}^k \left(\frac{1}{k+1}I + \frac{k-j-1}{k+1}U_{k-j-2} D_{k-j-1} + \frac{k-j}{k+1}\Gamma_j \right),$$

for some Γ_j satisfying $\|\Gamma_j\| \leq \gamma$. Combining these observations yields the desired result:

$$\begin{aligned} &D^{k+1}U_{k-(j+1)}^{k+1} - \frac{(j+1)+1}{k+1}U_{k-j+1}^k + \frac{k-(j+1)}{k+1}U_{k-(j+1)-1}^k D_{k-j+1} \\ &= D^{k+1}U_{k-(j+1)}^{k+1} - \frac{j+1}{k+1}U_{k-j-1}^k - U_{k-j-1}^k \left(\frac{1}{k+1}I - \frac{k-j-1}{k+1}U_{k-j-2} D_{k-j-1} \right) \\ &= \left(\sum_{i=-1}^{j-1} \frac{k-i}{k+1} U_{k-i-1}^k \Gamma_i U_{k-j-1}^{k-1-i} \right) + \frac{k-j}{k+1}U_{k-(j+1)}^k \Gamma_j \\ &= \sum_{i=-1}^j \frac{k-i}{k+1} U_{k-1-i}^k \Gamma_i U_{k-(j+1)}^{k-1-i}. \end{aligned}$$

□

We now show how to use this strengthened result to prove tighter bounds on the quadratic form $\langle f, N_k^j f \rangle$ which implies a stronger version of Lemma 7.2 as an immediate corollary. This improvement mainly matters in the regime where $\gamma \leq 2^{-ck}$ for c a small constant.

PROPOSITION A.1. *Let (X, Π) be a d -dimensional γ -local-spectral expander with γ satisfying $\gamma \leq 2^{-\Omega(k+j)}$, $k+j \leq d$, and $f_\ell \in V_k^\ell$. Then:*

$$\langle f_\ell, N_k^j f_\ell \rangle = \frac{\binom{k}{\ell}}{\binom{k+j}{\ell}} \left(1 \pm \frac{j(j+2k+2\ell+3)}{4} \gamma \pm c_3(k, j, \ell) \gamma^2 \right) \langle f_\ell, f_\ell \rangle$$

where $c_3(k, j, \ell) = O((k+j)^3 \binom{k+j}{\ell})$.

Proof. We proceed by induction on j . We will prove a slightly stronger statement for the base-case $j = 1$:

$$\langle f_\ell, D_{k+1} U_k f_\ell \rangle = \left(\frac{k+1-\ell}{k+1} \pm \frac{(k-\ell+1)(k+\ell+2)}{2(k+1)} \gamma \pm c_2(k, \ell) \gamma^2 \right) \langle f_\ell, f_\ell \rangle,$$

where $c_2(k, \ell) = O(k^3 \binom{k}{\ell})$. Recall that f_ℓ may be expressed as $U_\ell^k g_\ell$, for $g_\ell \in H^\ell$. For notational convenience, we write $f_\ell^i = U_\ell^i g_\ell$. Then we may expand the inner product based on Lemma 7.1, and simplify based on applying the naive bounds on N_k^i given by Corollary 7.1:

$$\begin{aligned} \langle f_\ell, D_{k+1} U_k f_\ell \rangle &= \langle f_\ell, D_{k+1} U_\ell^{k+1} g_\ell \rangle \\ &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle + \sum_{i=-1}^{k-\ell-1} \langle f_\ell, \frac{k-i}{k+1} U_{k-1-i}^k \Gamma_i U_\ell^{k-1-i} g_\ell \rangle \\ &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle + \sum_{i=-1}^{k-\ell-1} \frac{k-i}{k+1} \langle N_{k-i-1}^{i+1} f_\ell^{k-1-i}, \Gamma_i f_\ell^{k-1-i} \rangle \\ &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle + \sum_{i=-1}^{k-\ell-1} \frac{k-i}{k+1} \frac{\binom{k-i-1}{\ell}}{\binom{k}{\ell}} \langle f_\ell^{k-1-i}, \Gamma_i f_\ell^{k-1-i} \rangle + \sum_{i=-1}^{k-\ell-1} \frac{k-i}{k+1} \langle h_i, \Gamma_i f_\ell^{k-1-i} \rangle \end{aligned}$$

where $\|h_i\| \leq \gamma(k-\ell)(i+1) \|g_\ell\|$. We now apply Cauchy-Schwarz, and Lemma 7.2 to collect terms in $\langle f_\ell, f_\ell \rangle$:

$$\begin{aligned} \langle f_\ell, D_{k+1} U_k f_\ell \rangle &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle \pm \gamma \sum_{i=-1}^{k-\ell-1} \frac{k-i}{k+1} \frac{\binom{k-i-1}{\ell}}{\binom{k}{\ell}} \langle f_\ell^{k-1-i}, f_\ell^{k-1-i} \rangle \pm a_1(k, \ell) \gamma^2 \langle g_\ell, g_\ell \rangle \\ &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle \pm \gamma \sum_{i=-1}^{k-j-1} \frac{k-i}{k+1} \frac{1}{\binom{k}{j}} \langle g_\ell, g_\ell \rangle \pm a_2(k, \ell) \gamma^2 \langle g_\ell, g_\ell \rangle \\ &= \frac{k-\ell+1}{k+1} \langle f_\ell, f_\ell \rangle \pm \gamma \sum_{i=-1}^{k-j-1} \frac{k-i}{k+1} \frac{\langle f_\ell, f_\ell \rangle}{(1 - c_1(k, \ell) \gamma)} \pm a_2(k, \ell) \gamma^2 \frac{\langle f_\ell, f_\ell \rangle}{(1 - c_1(k, \ell) \gamma)} \\ &= \frac{k-\ell+1}{k+1} \left(1 \pm \frac{(k+\ell+2)}{2} \gamma \pm a_3(k, \ell) \gamma^2 \right) \langle f_\ell, f_\ell \rangle \end{aligned}$$

where the final step comes from a Taylor expansion assuming γ sufficiently small, and $a_3(k, \ell) = O(k^3 \binom{k}{\ell})$.

The inductive step follows from noting that the canonical walk essentially acts like a product of upper walks from lower levels in the following sense:

$$\begin{aligned} \langle f_\ell, N_k^j f_\ell \rangle &= \langle U_k^{k+j-1} f_\ell, D_{k+j} U_k^{k+j} f_\ell \rangle \\ &= \langle U_k^{k+j-1} f_\ell, N_{k+j-1}^1 (U_k^{k+j-1} f_\ell) \rangle. \end{aligned}$$

Thus by the base-case and inductive hypothesis we get:

$$\begin{aligned} \langle f_\ell, N_k^j f_\ell \rangle &= \langle U_k^{k+j-1} f_\ell, N_{k+j-1}^1 (U_k^{k+j-1} f_\ell) \rangle \\ &= \left(\frac{k+j-\ell}{k+j} \pm \frac{(k+j-\ell)(k+j+\ell+1)}{2(k+j)} \gamma \pm c_2(k+j-1, \ell) \gamma^2 \right) \langle f_\ell, N_k^{j-1} f_\ell \rangle \\ &= \frac{\binom{k}{\ell}}{\binom{k+j}{\ell}} \left(1 \pm \frac{(k+j+\ell+1)}{2} \gamma \pm \frac{k+j}{k+j-\ell} c_2(k+j-1, \ell) \gamma^2 \right) \\ &\quad \cdot \left(1 \pm \frac{(j-1)(j+2k+2\ell+2)}{4} \gamma \pm c_3(k, j-1, \ell) \gamma^2 \right) \langle f_\ell, f_\ell \rangle \\ &= \frac{\binom{k}{\ell}}{\binom{k+j}{\ell}} \left(1 \pm \frac{j(j+2k+2\ell+3)}{4} \gamma \pm c_3(k, j, \ell) \gamma^2 \right) \langle f_\ell, f_\ell \rangle, \end{aligned}$$

□

Notice that this immediately implies a stronger version of Lemma 7.2 since $\langle U_\ell^k g_\ell, U_\ell^k g_\ell \rangle = \langle N_\ell^{k-\ell} g_\ell, g_\ell \rangle$. Finally, we conjecture that a stronger result is true, and the error dependence on γ should in fact be $\exp(-\text{poly}(k)\gamma)$. Proving this would require a more careful and involved analysis of how the error term propagates.

B Orthogonality and the HD-Level-Set Decomposition

In this section we discuss in a bit more depth the error in [KO20, Theorem 5.10], and further show by direct counter-example that its implication [KS20] that the HD-Level-Set is orthogonal does not hold. In [KO20], Kaufman and Oppenheim analyze an approximate eigendecomposition of the upper walk N_k^1 for two-sided local-spectral expanders. They prove a specialized version of Theorem 6.1 for this case, and in particular that for sufficiently strong two-sided local-spectral expanders, the spectra of N_k^1 is divided into strips concentrated around the approximate eigenvalues of their decomposition. They call the span of each strip W^i , and note that the W^i form an orthogonal decomposition of the space. Let V^i be the space in the original approximate eigendecomposition corresponding to strip W^i . Kaufman and Oppenheim claim in [KO20, Theorem 5.10] that the W^i are closely related to the original approximate decomposition in the following sense:

$$\forall \phi \in C_k : \|P_{W^i} \phi\| \leq c \|P_{V^i} \phi\|$$

for some constant $c > 0$, where P_{W^i} and P_{V^i} are projection operators. Unfortunately, this relation cannot hold, as it implies [KS20] that the HD-Level-Set Decomposition is orthogonal for sufficiently strong two-sided local-spectral expanders, which we will show below is false by direct example. In slightly greater detail, the issue in the argument is the following. The authors show that for any $j \neq i$:

$$\|P_{W^j} P_{V^i}\| \leq c',$$

for some small constant c' , and then claim that this fact implies for any $\phi \in C_k$:

$$\|P_{W^j} P_{V^i} \phi\| \leq c' \|P_{V^j} \phi\|.$$

Unfortunately, this is not true—the righthand side should read P_{V^i} rather than P_{V^j} for the relation to hold, but this makes it impossible to compare $P_{W^i} \phi$ solely to $P_{V^i} \phi$.

We now move to showing that for any $\gamma > 0$, there exists a two-sided γ -local-spectral expander such that the HD-Level-Set Decomposition is not orthogonal, which implies [KO20, Theorem 5.10] cannot hold by arguments of [KS20].

PROPOSITION B.1. *For any $\gamma > 0$, there exists a two-sided γ -local-spectral expander such that the HD-Level-Set Decomposition is not orthogonal.*

Proof. Our construction is based off of a slight modification of the complete complex $J(n, 3)$. In particular, we consider the uniform distribution Π over triangles $X = \binom{[n]}{3} \setminus \{(123)\}$. It is not hard to see through direct computation that (X, Π) is a two-sided $O(1/n)$ -local-spectral expander. Recall that the link of 1, $U_1^3 \mathbb{1}_1$, lies in $V_3^0 \oplus V_3^1$. Our goal is to prove the existence of a function $f = Ug \in V_3^2$ such that the inner product:

$$(B.1) \quad \langle U_1^3 \mathbb{1}_1, f \rangle \propto \sum_{(1xy) \in X} g(1x) + g(1y) + g(xy)$$

is non-zero. To do this, we first simplify the above expression assuming $g \in \text{Ker}(D_2)$, which we recall implies the following relations:

$$\forall y \in [n] : \sum_{(xy) \in X(2)} \Pi_2(xy) g(xy) = 0.$$

In particular, summing over all $y \in [n]$ gives

$$\sum_{(xy) \in X(2)} \Pi_2(xy) g(xy) = 0.$$

Notice further that by definition of Π_2 , we have $\Pi_2(12) = \Pi_2(13) = \Pi_2(23) = \frac{n-3}{3(\binom{n}{3}-3)}$, and otherwise $\Pi_2(xy) = \frac{n-2}{3(\binom{n}{3}-3)}$. We then may write:

$$\begin{aligned} \sum_{\substack{(1x) \in X(2): \\ x \notin [3]}} g(1x) &= -\frac{n-3}{n-2} (g(12) + g(13)), \\ \sum_{\substack{(xy) \in X(2): \\ (xy) \notin [3] \times [3]}} g(xy) &= -\frac{n-3}{n-2} (g(12) + g(13) + g(23)). \end{aligned}$$

Plugging this into Equation (B.1), the inner product drastically simplifies to depend only on $g(23)$. To see this, we separate the inner product into two terms and deal with each separately:

$$\sum_{(1xy) \in X} g(1x) + g(1y) + g(xy) = \left(\sum_{(1xy) \in X} g(1x) + g(1y) \right) + \sum_{(1xy) \in X} g(xy)$$

We start with the former. Notice that each face $(1z)$ in this term is counted exactly the number of times it appears in a triangle in X , and further that this is exactly how Π_2 is defined. Thus we have:

$$\left(\sum_{(1xy) \in X} g(1x) + g(1y) \right) \propto \sum_{(1x) \in X(2)} \Pi_2(1x) g(1x) = 0.$$

It is left to analyze the latter term. Since (123) is not in our complex, we may write:

$$\begin{aligned} \sum_{(1xy) \in X} g(xy) &= \left(\sum_{\substack{(xy) \in X(2): \\ (xy) \notin [3] \times [3]}} g(xy) \right) - \left(\sum_{\substack{(1x) \in X(2): \\ x \notin [3]}} g(1x) \right) \\ &= -\frac{n-3}{n-2} (g(12) + g(13) + g(23)) + \frac{n-3}{n-2} (g(12) + g(13)) \\ &= -\frac{n-3}{n-2} g(23). \end{aligned}$$

Thus it remains to show that there exists $g \in \text{Ker}(D_2)$ such that $g(23) \neq 0$. Note that the kernel of D_2 is exactly the space of solutions to the underdetermined linear system of equations given by $D_2 g(i) = 0$ for all $1 \leq i \leq n$. Thus we can check if a solution exists with $g(23) = c$ for $c \neq 0$ by ensuring that this constraint is linearly independent of the $D_2 g(i)$. This can be checked through a direct but tedious computation that we leave to the reader. \square

C Unique Games

C.1 Random-walks and Weighted graphs Unique games are defined on weighted, *undirected* constraint graphs, unlike most of the walks we analyze in the previous section. However, it is not hard to see that every self-adjoint random walk corresponds to some underlying undirected graph. In particular, recall that given a weighted, undirected graph $G(V, E)$, G induces a random walk on vertices where the transition probability from $x \in V$ to $y \in V$ is given by the normalized weights:

$$P(x, y) = \frac{W(s, t)}{\sum_{v \in N(s)} W(s, v)}.$$

In fact, *every* self-adjoint walk can be described as such a process.

LEMMA C.1. *Let M be a k -dimensional HD-walk on a weighted simplicial complex (X, Π) . Define G_M to be the graph whose vertex set is $X(k)$ and whose edge set consists of any pair (s, t) such that $M(s, t) \neq 0$. Further, endow the edges of G_M with weight:*

$$W(s, t) = \Pi_k(s)M(s, t).$$

Then M is the random walk induced by G_M .

Proof. The crucial observation for this proof is the following implication due to M being self adjoint:

$$\forall s, t \in X(k) : \Pi_k(s)M(s, t) = \Pi_k(t)M(t, s).$$

Given this fact, let $P_G(s, t)$ denote the transition probability of the induced walk on G . We have:

$$\begin{aligned} P_G(s, t) &= \frac{\Pi_k(s)M(s, t)}{\sum_{v \in N(s)} \Pi_k(v)M(v, t)} \Pi_k(s)M(s, t) \\ &= \frac{M(s, t)}{\sum_{v \in N(s)} M(v, t)} M(s, t) \\ &= M(s, t) \end{aligned}$$

□

C.2 The Shift-Partition Potential The analysis of the Iterated Condition&Round algorithm relies on analysing a quantity called the Approximate shift-partition potential which was defined in BBKSS. For completeness we include the definitions of the various potential functions used in this section.

For analysing the Condition&Round subroutine, define a low-degree polynomial $\Phi_{\beta, \nu} : \Sigma^V \times \Sigma^V \rightarrow [0, \infty)$ called the “approximate shift partition potential”:

DEFINITION C.1. (APPROXIMATE SHIFT-PARTITION POTENTIAL) *For any $\nu, \beta \in (0, 1)$, and two assignments $X, X' \in \Sigma^V$ define the approximate shift-partition potential to be the quantity*

$$\Phi_{\beta, \nu}(X, X') = \sum_{s \in \Sigma} \mathbb{E}_u [\mathbb{1}[X_u - X'_u = s] \cdot p_{\beta, \nu}(\text{val}_u(X))]^2,$$

for $p_{\beta, \nu}(x)$ the degree- $\tilde{O}(1/\nu)$ polynomial which SoS-certifiably approximates the indicator $\mathbb{1}[x \geq \beta]$ for $x \in [0, 1]$ up to an error of ν . The exact definition of $p_{\beta, \nu}$ and the properties required of it are described in more detail in Section 7 of BBKSS and we omit them here.

The approximate shift-partition potential on a pseudodistribution μ of degree at least $O(\deg(\Phi_{\beta, \nu}))$ is defined as:

$$\Phi_{\beta, \nu}(\mu) = \tilde{\mathbb{E}}_{\mu}[\Phi_{\beta, \nu}(X, X')].$$

For analysing the Iterated Condition&Round algorithm we further need to define the potential when restricted to a subgraph which in our case will be a link of the complex.

DEFINITION C.2. (GLOBAL SHIFT-POTENTIAL RESTRICTED TO LINKS) *Let $I = (M, \Pi)$ be a UG instance on a complex X where M is a complete random walk on $X(k)$ with stationary distribution π_k . For any $\nu, \beta \in (0, 1)$ and a link X_τ of the complex X , define the approximate global shift-partition potential restricted to X_τ to be the quantity:*

$$\Phi_{\beta, \nu}(X, X')|_\tau = \sum_{s \in \Sigma} \mathbb{E}_{u \sim \pi_k|_\tau} [\mathbb{1}[X_u - X'_u = s] \cdot p_{\beta, \nu}(\text{val}_u(X))]^2,$$

where $\text{val}_u(X) = \mathbb{E}_{(u, v) \sim M} [\mathbb{1}[X \text{ satisfies } (u, v)]]$, and as before $p_{\beta, \nu}(x)$ is the degree- $\tilde{O}(1/\nu)$ polynomial that ν -approximates the indicator function $\mathbb{1}[x \geq \beta]$. The global potential restricted to X_τ with respect to a pseudodistribution μ will be denoted by $\Phi_{\beta, \nu}(\mu|_\tau)$:

$$\Phi_{\beta, \nu}(\mu|_\tau) = \tilde{\mathbb{E}}_{\mu}[\Phi_{\beta, \nu}(X, X')|_\tau].$$

Note that the global shift-partition potential measures the size of the global partition inside X_τ , and although the expectation is taken only over the vertices $u \in X_\tau$, $\text{val}_u(X)$ is a function of *all* the edges in M that are incident on u , not just the edges in X_τ . We will also need to consider the potential induced on a link, which is defined as applying the potential function (Definition C.1) to the graph induced by X_τ :

DEFINITION C.3. (INDUCED SHIFT-POTENTIAL ON A LINK) *Let $I = (M, \Pi)$ be a UG instance on a complex X where M is a complete random walk on $X(k)$ with stationary distribution π_k . For any $\nu, \beta \in (0, 1)$ and a link X_τ of the complex X , define the approximate induced shift-partition potential on X_τ to be the quantity:*

$$\Phi_{\beta, \nu}^\tau(X, X') = \sum_{s \in \Sigma} \mathbb{E}_{u \sim \pi_k | \tau} [\mathbb{1}[X_u - X'_u = s] \cdot p_{\beta, \nu}(\text{val}_u^\tau(X))]^2,$$

where $\text{val}_u^\tau(X) = \mathbb{E}_{(u, v) \sim M | \tau} [\mathbb{1}[X \text{ satisfies } (u, v)]]$. The induced potential on X_τ with respect to a pseudodistribution μ will be denoted by $\Phi_{\beta, \nu}^\tau(\mu)$:

$$\Phi_{\beta, \nu}^\tau(\mu) = \tilde{\mathbb{E}}_\mu [\Phi_{\beta, \nu}^\tau(X, X')].$$

Note that in the above definition the value of a vertex $u \in X_\tau$ is measured only with respect to the edges incident on u that lie *inside* the link X_τ .

C.3 Sum of Squares and the HD-Level-Set Decomposition This section is devoted to proving Lemma 10.2 which we separate into two parts. First, we examine the relation between $\|f_i\|$ and $\|g_i\|$.

LEMMA C.2. (RESTATEMENT LEMMA 10.2 (PART 1)) *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander. Then for any $f_i = U_i^k g_i \in V_k^i$:*

$$\vdash_2 \langle f_i, f_i \rangle \in \left(\frac{1}{\binom{k}{i}} \pm \frac{(k-i)(k+1)}{i+2} \gamma \right) \langle g_i, g_i \rangle$$

Proof. It is sufficient to prove the following equality:

$$(C.2) \quad \langle f_i, f_i \rangle = \frac{1}{\binom{k}{i}} \langle g_i, g_i \rangle + \langle g_i, \Gamma g_i \rangle,$$

where $\|\Gamma\| < \frac{(k-i)(k+1)}{i+2} \gamma$. To see why, note that:

$$\langle g_i, \Gamma g_i \rangle = \left\langle g_i, \frac{\Gamma + \Gamma^*}{2} g_i \right\rangle$$

where Γ^* is the adjoint of Γ . Since $\frac{\Gamma + \Gamma^*}{2}$ is self-adjoint and $\|\Gamma^*\| = \|\Gamma\|$, we have both:

1. $\vdash_2 \left\langle g_i, \frac{\Gamma + \Gamma^*}{2} g_i \right\rangle \leq \frac{(k-i)(k+1)}{i+2} \gamma \langle g_i, g_i \rangle$
2. $\vdash_2 \left\langle g_i, \frac{\Gamma + \Gamma^*}{2} g_i \right\rangle \geq -\frac{(k-i)(k+1)}{i+2} \gamma \langle g_i, g_i \rangle$

as desired. To prove Equation (C.2), we induct on k . The base case $k = i$ is trivial. Assume $k > i$, then we may write:

$$\langle f_i, f_i \rangle = \langle U_i^{k-1} g_i, D_k U_i^k g_i \rangle$$

In order to apply the inductive hypothesis, we recall the machinery of [DDFH18, Claim 8.8] for pushing D_k through U_i^k . In particular:

$$\left\| D_k U_i^k - \frac{k-i}{k} U_i^{k-1} - \frac{i}{k} U_{i-1}^{k-1} D_i \right\| \leq (k-i) \gamma.$$

Since g_i lies in the kernel of D_i , combining this with our initial observation gives:

$$\langle U_i^{k-1} g_i, D_k U_i^k g_i \rangle = \frac{k-i}{k} \langle U_i^{k-1} g_i, U_i^{k-1} g_i \rangle + \langle g_i, \Gamma g_i \rangle,$$

where $\|\Gamma\| \leq (k-i)\gamma$. Since $U_i^{k-1} g_i \in V_i^{k-1}$. Applying the inductive hypothesis, we see that:

$$\begin{aligned} \langle f_i, f_i \rangle &= \frac{k-i}{k} \left(\frac{1}{\binom{k-1}{i}} \langle g_i, g_i \rangle + \langle g_i, \Gamma' g_i \rangle \right) + \langle g_i, \Gamma g_i \rangle \\ &= \frac{1}{\binom{k}{i}} \langle g_i, g_i \rangle + \langle g_i, \left(\frac{k-i}{k} \Gamma' + \Gamma \right) g_i \rangle, \end{aligned}$$

where:

$$\left\| \frac{k-i}{k} \Gamma' + \Gamma \right\| \leq \frac{(k-i)(k+1)}{i+2} \gamma$$

by the triangle inequality and inductive hypothesis. \square

Second, we prove that a version of approximate orthogonality of the HD-Level-Set Decomposition has a low-degree SoS proof.

LEMMA C.3. (RESTATEMENT LEMMA 10.2 (PART 2)) *let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $f \in C_k$. Then for any $f_i = U_i^k g_i \in V_k^i$ and $f_j = U_j^k g_j \in V_k^j$ we have:*

$$\vdash_2 \langle f_i, f_i \rangle \leq (1 + c_1 \gamma) \langle f, f \rangle,$$

and

$$\vdash_2 \langle f, f_i \rangle \geq -c_2 \gamma \langle f, f \rangle,$$

where $c_1, c_2 \leq 2^{O(k)}$

Proof. First, note that by Lemma 7.1 and the fact that $g_j \in \text{Ker}(D_j)$, we have that:

$$\langle f_i, f_j \rangle = \langle g_i, \Gamma g_j \rangle$$

where $\|\Gamma\| \leq 2^{O(k)}$. We can bound the latter by an SoS version of Cauchy Schwarz. In particular, we have that:

$$\begin{aligned} \langle f_i, f_j \rangle &= \langle g_i, \Gamma g_j \rangle \\ &\leq \frac{\|\Gamma\|}{2} \langle g_i, g_i \rangle + \frac{1}{2 \|\Gamma\|} \langle \Gamma g_j, \Gamma g_j \rangle \\ &\leq \frac{\|\Gamma\|}{2} \langle g_i, g_i \rangle + \frac{\|\Gamma\|}{2} \langle g_j, g_j \rangle \\ &\leq c\gamma(\langle f_i, f_i \rangle + \langle f_j, f_j \rangle) \end{aligned}$$

where all inequalities are degree 2 SoS and $c \leq 2^{O(k)}$. Further, notice that by applying the same argument to $\langle -f_i, f_j \rangle$ we get that $\langle f_i, f_j \rangle$ is bounded above and below:

$$\vdash_2 -c\gamma(\langle f_i, f_i \rangle + \langle f_j, f_j \rangle) \leq \langle f_i, f_j \rangle \leq c\gamma(\langle f_i, f_i \rangle + \langle f_j, f_j \rangle).$$

We now apply this fact directly to prove the two desired inequalities. First, we have

$$\begin{aligned} \langle f, f \rangle &= \sum_{\ell=0}^k \langle f_\ell, f_\ell \rangle + \sum_{\ell \neq m} \langle f_\ell, f_m \rangle \\ &\geq \sum_{\ell=0}^k \langle f_\ell, f_\ell \rangle - \sum_{\ell=0}^k k c \gamma \langle f_\ell, f_\ell \rangle \\ &= (1 - k c \gamma) \sum_{\ell=0}^k \langle f_\ell, f_\ell \rangle \\ &\geq (1 - c_1 \gamma) \langle f_i, f_i \rangle, \end{aligned}$$

where $c_1 \leq 2^{O(k)}$. For small enough γ , Taylor expanding $(1 - c_1\gamma)^{-1}$ gives the desired result. We can now apply the first inequality to easily prove the second inequality:

$$\begin{aligned}\langle f, f_i \rangle &= \langle f_i, f_i \rangle + \sum_{j \neq i} \langle f_i, f_j \rangle \\ &\geq (1 - ck\gamma) \langle f_i, f_i \rangle - \sum_{j \neq i} c\gamma \langle f_j, f_j \rangle \\ &\geq -ck\gamma(1 + c_2\gamma) \langle f, f \rangle \geq -c_3 \langle f, f \rangle\end{aligned}$$

for $c_2, c_3 \leq 2^{O(k)}$. \square

C.4 Remaining Proofs from Section 10.3

CLAIM C.1. (RESTATEMENT CLAIM 10.1) *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -local-spectral expander with $d > k$ and $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ with stationary distribution π . If the total laziness of M $\mathbb{E}[\mathbb{1}_v^T M \mathbb{1}_v]$ is at least .1, then the spectral gap of M is at least $\Omega(1/k)$ and $\lambda_{k/2}(M) \leq .68$.*

Proof. To see this, note that any walk can be (approximately) decomposed by Lemma 7.1 into an affine combination of pure walks of the form $(U_{k-1}D_k)^i$, that is:

$$M = \sum_{i=0}^{h(M)} \alpha_i (U_{k-1}D_k)^i + \Gamma,$$

where $\|\Gamma\| \leq w(M)2^{O((M))}$ and $(U_{k-1}D_k)^0 = I$. Notice that for any $(U_{k-1}D_k)^i$ for $i > 0$, the probability of returning to any given k -face is at most γ . This follows from the fact that the probability of returning to any particular face in the final up step (given that the previous step is at $\sigma \in X(k-1)$) is at most $|\Pi_{\sigma,1}|_\infty \leq \gamma$. This last inequality follows from (X, Π) being a two-sided γ -local-spectral expander satisfying $d > k$, as the link (X_σ, Π_σ) (at level k) is then a γ -spectral expander, which are easily shown to satisfy this property.

Since the total laziness of M is at most .1, by our assumption on γ the above analysis implies that the coefficient α_0 cannot be too large, say at most $1/5$ (since the remaining parts cannot contribute much to the laziness). We can use this fact to bound the spectral gap of M by the same trick used in Proposition 7.3: transferring M over to the complete complex. Since the approximate eigenvalues of M are independent of the underlying complex and $(UD)^i$ has a vanishing lazy component on $J(n, d)$ as n goes to infinity, it must be the case that in lazy component of M on the Johnson complex is also small. Since M can be written as a convex combination of partial-swap walks, all of which have spectral gap at least $1/k$ (save for the identity which carries at most a $1/5$ of the weight), we get that $1 - \lambda_2(M) \geq \Omega(1/k)$ as desired. The bound on $\lambda_{k/2}$ follows similarly noting that all (non-identity) partial-swap walks satisfy $\lambda_{k/2} \leq 1/2$. \square

C.5 Remaining Proofs from Section 10.4 In this section we will prove the claims that were stated without proof in Section 10.4. We will use the exact definition of the potential functions from Appendix C.2 and also the properties of the polynomials $p_{\beta, \nu}(Y)$ that SoS certifiably ν -approximate the indicator function $\mathbb{1}[Y \geq \beta]$. See Section 7 of BBKSS for a detailed overview about these polynomials and their properties.

PROPOSITION C.1. (RESTATEMENT PROPOSITION 10.2) *Let M be a k -dimensional complete HD-walk on a d -dimensional, two-sided γ -local-spectral expander satisfying $\gamma w(M)2^{O(h(M)+k)} \leq 1/2^k < \eta < 0.02$ and $d > k$, and I be an affine unique games instance over M with value at least $1 - \eta$. Let $r(\eta) = R_{1-16\eta}(M) - 1$, and assume $r \leq k/2$. Let $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$. Then given a degree- $\tilde{O}(1/\nu)$ pseudodistribution μ satisfying the axioms \mathcal{A}_I , there exists an r -link X_τ such that the global potential restricted to X_τ is large:*

$$\Phi_{\beta, \nu}^I(\mu|_\tau) \geq \frac{1}{4\binom{k}{r}}.$$

Proof. [Proof of Proposition 10.2] Before diving into the details of the proof we give a proof overview and some high level intuition. Given two UG assignments X, X' , BBKSS define a partition of the graph G into $|\Sigma|$ parts and

call this the approximate shift-partition of G with respect to X, X' . The partition is given by the functions $f_s^{X, X'}$ for $s \in \Sigma$, where $f_s^{X, X'}(u)$ roughly equals the indicator function of whether u belongs to part s . Concretely, we have that $f_s^{X, X'}(u) = \mathbb{1}[X_u - X'_u = s] \cdot p_{\beta, \nu}(\text{val}_u(X))$, that is, u belongs to part s if $X_u - X'_u = s$ and the value of u measured with respect to assignment X is large. Note that any vertex can only belong to one part s but some vertices may not belong to any part. BBKSS prove that for any instance of affine unique games and any two UG assignments X, X' that have large value, the functions $f_s^{X, X'}$ satisfy three properties: a) they include a large number of vertices in total, b) they form a non-expanding partition of G , that is, the fraction of edges that go across parts is small (roughly at most edges violated by X, X'), and finally c) the functions f_s are approximately Boolean-valued on average. These properties are proved in Section 4 of their paper. They then apply the structure theorem for non-expanding sets of G to conclude that there exists a subgraph H where the size of the partition restricted to H ($\sum_s \mathbb{E}_{u \sim H} [f_s^{X, X'}(u)]^2$) must be large. The crucial observation then is that the pseudoexpectation of the restricted size on H is by definition the global potential on H . Therefore by taking pseudoexpectation over $X, X' \sim \mu$ where μ is a pseudodistribution with high value, they conclude that there exists a subgraph with large potential. We will use the same proof strategy to conclude that the global potential on a link is large, using the SoS structure theorem (Theorem 10.3) along the way.

Let us elaborate on the properties (proved in BBKSS Section 4) of the approximate shift-partition given by the functions $f_s^{X, X'}$. We will drop the superscript X, X' for convenience and throughout the proof the pseudoexpectation is taken with respect to $X, X' \sim \mu$. Since μ is a degree $\tilde{O}(1/\nu)$ pseudodistribution with value at least $1 - \eta$ we have that:

1. The approximate shift-partition includes most of the vertices in pseudoexpectation over μ :

$$\tilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \mathbb{E}_u [f_s] \right] \geq 1 - \frac{\eta}{1 - \beta - \nu} - \nu$$

2. The approximate shift-partition is non-expanding²⁸ in pseudoexpectation over μ :

$$\tilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \langle f_s, (I - A_G) f_s \rangle \right] \leq 2\eta + 2 \frac{\eta}{1 - \beta - \nu} + 2\nu$$

3. The functions f_s are approximately Boolean-valued in pseudoexpectation over μ :

$$\tilde{\mathbb{E}}_{\mu} \left[\sum_s B(f_s) \right] = \tilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \mathbb{E}_u [f_s - f_s^2] \right] \leq \frac{\eta}{1 - \beta - \nu} + \nu$$

Since the functions f_s are non-expanding (property (2) above), this allows us to use our ℓ_2 characterization of non-expansion, to get a lower bound on the size of f_s restricted to r -links. We will now use the key observation that $\mathbb{E}_{X_\tau} [f_s]^2$ is exactly $D_r^k f_s(\tau)^2$. In particular, applying Theorem 10.3 to the function f_s with ℓ set to $r(\eta) = R_{1-16\eta}(M)$, gives:

$$\langle f_s, (I - M) f_s \rangle \geq (1 - \lambda_{r+1}(M)) \left((1 - c_1 \gamma) \mathbb{E}[f_s] - \binom{k}{r} \langle D_r^k f_s, D_r^k f_s \rangle + (1 - c_1 \gamma) B(f_s) \right),$$

where $B(f) = \mathbb{E}_{u \sim X_k} [f(u) - f(u)^2]$. Now using the key observation that $D_r^k f_s$ is exactly $\mathbb{E}_{X_\tau} [f_s]$ and noting that $1 - \lambda_{r+1} \geq 16\eta$ by definition, we can simplify the above to:

$$\frac{1}{16\eta} \langle f_s, (I - M) f_s \rangle \geq \left((1 - c_1 \gamma) \mathbb{E}[f_s] - \binom{k}{r} \mathbb{E}_{r \sim X(r)} [\mathbb{E}_{X_\tau} [f_s]^2] + (1 - c_1 \gamma) B(f_s) \right).$$

²⁸If the functions $f_s(u)$ were exactly Boolean-valued and further exactly partitioned the graph, then the term $\sum_s \langle f_s, (I - A_G) f_s \rangle$ is equal to the fraction of edges in G that have endpoints in different parts. Therefore this expression measures the non-expansion of a partition.

Taking a sum over s and then a pseudoexpectation over μ yields:

$$\begin{aligned} \frac{1}{16\eta} \widetilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \langle f_s, (I - M)f_s \rangle \right] &\geq \widetilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \left((1 - c_1\gamma) \mathbb{E}[f_s] - \binom{k}{r} \mathbb{E}_{\tau \in X(r)} \left[\mathbb{E}_{X_{\tau}} [f_s]^2 \right] + (1 - c_1\gamma) B(f_s) \right) \right] \\ &= (1 - c_1\gamma) \widetilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \mathbb{E}[f_s] \right] - \binom{k}{r} \mathbb{E}_{\tau \in X(r)} \left[\widetilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} \mathbb{E}_{X_{\tau}} [f_s]^2 \right] \right] + (1 - c_1\gamma) \widetilde{\mathbb{E}}_{\mu} \left[\sum_{s \in \Sigma} B(f_s) \right], \end{aligned}$$

where $c_1 \leq 2^{O(k)}$.

We will now use the observation that the global potential restricted to X_{τ} exactly corresponds to the second term above! Recall that the global potential (Definition C.2) over an r -link X_{τ} is given by:

$$\Phi_{\beta,\nu}^I(\mu|_{\tau}) = \widetilde{\mathbb{E}}_{X, X' \sim \mu} \left[\sum_{s \in \Sigma} \mathbb{E}_{u \sim X_{\tau}} [f_s^{X, X'}(u)]^2 \right].$$

For convenience of notation we will drop the superscript I in the potential notation henceforth. Re-arranging gives a lower-bound on the global potential averaged across r -links:

$$(C.3) \quad \mathbb{E}_{\tau \in X(r)} [\Phi_{\beta,\nu}(\mu|_{\tau})] \geq \frac{1}{\binom{k}{r}} \left((1 - c_1\gamma) \widetilde{\mathbb{E}} \left[\sum_s \mathbb{E}[f_s] \right] + (1 - c_1\gamma) \widetilde{\mathbb{E}} \left[\sum_s B(f_s) \right] - \frac{1}{16\eta} \widetilde{\mathbb{E}} \left[\sum_s \langle f_s, (I - M)f_s \rangle \right] \right).$$

Plugging in the properties of the functions f_s discussed at the start of the proof, this gives the following bound on the global potential:

$$\mathbb{E}_{\tau \in X(r)} [\Phi_{\beta,\nu}(\mu|_{\tau})] \geq \frac{1}{\binom{k}{r}} \left(1 - c_1\gamma - 2 \frac{\eta}{1 - \beta - \nu} - 2\nu - \frac{1}{16\eta} \left(2\eta + 2 \frac{\eta}{1 - \beta - \nu} + 2\nu \right) \right).$$

Setting $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$ and recalling our assumptions on γ and η , this may be simplified by direct computation to:

$$\mathbb{E}_{\tau \in X(r)} [\Phi_{\beta,\nu}(\mu|_{\tau})] \geq \frac{1}{4\binom{k}{r}}.$$

As a result, we see by averaging that there must exist some r -link with high potential:

$$(C.4) \quad \exists \tau \in X(r) : \Phi_{\beta,\nu}(\mu|_{\tau}) \geq \frac{1}{4\binom{k}{r}},$$

as desired. \square

We will now prove Lemma 10.3.

LEMMA C.4. (LEMMA 10.3 RESTATED) *Assume the conditions of Proposition 10.2 hold. Let X_{τ} be an r -link for $r = R_{1-16\eta}(M) - 1$ with high global potential: $\Phi_{\beta,\nu}^I(\mu|_{\tau}) \geq \frac{1}{4\binom{k}{r}}$, for $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$. Then the potential induced on X_{τ} is also high:*

$$\Phi_{\eta,\nu}^{I|_{\tau}}(\mu) \geq \frac{1}{8\binom{k}{r}}.$$

As discussed in Section 10, proving Lemma 10.3 requires that a strong expansion property holds for links: expansion must be similar vertex-by-vertex. We now formalize this property and prove it holds for HD-walks.

DEFINITION C.4. *The edge-expansion of a vertex v in a set $S \subseteq X(k)$ with respect to a random walk operator M is given by,*

$$\phi_S(M, v) = 1 - \mathbb{1}_v^T M \mathbb{1}_S.$$

The edge-expansion of S with respect to M is the average edge-expansion taken over vertices in S :

$$\phi(M, S) = \mathbb{E}_{v \sim S} [\phi_S(M, v)],$$

where $v \sim S$ is the distribution π_k conditioned on S . As before, we drop M from the notation when clear from context.

We prove that the expansion of vertices in a link cannot vary much from the link's overall expansion (which is small by Theorem 9.1).

LEMMA C.5. (RESTATEMENT OF LEMMA 10.4) *Let M be a k -dimensional HD-walk on d -dimensional two-sided γ -local-spectral expander satisfying $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ and $d > k$. Then for every i -link X_τ , the deviation of the random variable $\phi_{X_\tau}(v)$ ($v \sim X_\tau$) is small:*

$$\mathbb{E}_{v \sim X_\tau} [|\phi_{X_\tau}(v) - \phi(X_\tau)|] \leq \frac{1}{2^{11k}}.$$

We prove this claim by reducing to simpler HD-walks we can analyze directly. With that in mind, let's first prove that a stronger point-wise result holds for any product of "lower walks" $(U_{k-1}D_k)^i$.

PROPOSITION C.2. *Let M be a k -dimensional HD-walk on d -dimensional two-sided γ -local-spectral expander satisfying $\gamma \leq w(M)^{-1}2^{-\Omega(h(M)+k)}$ and $d > k$. For every i -link X_τ , walk $(U_{k-1}D_k)^t$, and k -face $v \in X_\tau$, the expansion of v in X_τ is almost exactly $\phi(X_\tau)$:*

$$|\phi_{X_\tau}((U_{k-1}D_k)^t, v) - \phi((U_{k-1}D_k)^t, X_\tau)| \leq t\gamma.$$

Proof. It is more convenient to analyze the non-expansion at v , $\bar{\phi}_{X_\tau}((U_{k-1}D_k)^t, v) = 1 - \phi_{X_\tau}((U_{k-1}D_k)^t, v)$, which for brevity we will denote by $\bar{\phi}(v)$. Below we will prove that for every $v \in X_\tau$, $\bar{\phi}(v) \in [c, c + t\gamma]$, for some fixed constant c . This immediately implies the result. Since by definition the non-expansion of X_τ denoted $\bar{\phi}((U_{k-1}D_k)^t, X_\tau)$, is an average over $\bar{\phi}(v)$, we get that:

$$|\phi(v) - \phi((U_{k-1}D_k)^t, X_\tau)| = |\bar{\phi}(v) - \bar{\phi}((U_{k-1}D_k)^t, X_\tau)| \leq t\gamma,$$

as desired.

Therefore let us now prove that $\bar{\phi}(v) \in [c, c + t\gamma]$ for all $v \in X_\tau$. First, notice that the non-expansion at v is lower bounded by the probability that the walk does not remove any element from τ in any down-step. Since the down step is uniformly random, this probability is the same across all $v \in X_\tau$. Denote this probability by c .

Let us now upper bound $\bar{\phi}(v)$. The probability that the walk returns to X_τ is equal to c plus the probability that starting from v , the walk leaves the link X_τ at some intermediate (down) step but ends back in X_τ regardless. We will show that the latter probability is at most $t\gamma$.

Consider the first down-step where the walk leaves X_τ , removing an element $w \in \tau$. To end up in the link of τ , the walk needs to add w back in one of the future up steps. The probability of this occurring at an up step from any $\sigma \in X(k-1)$ is exactly $\Pi_{\sigma,1}(w)$ by definition (where $\Pi_{\sigma,1}(w) = 0$ if $w \notin X_\sigma$). Since (X, Π) is a two-sided γ -local-spectral expander, (X_σ, Π_σ) is a standard γ -spectral expander. It is a standard result that such graphs cannot have any vertex of weight greater than γ , thus $\Pi_{\sigma,1}(w) \leq \gamma$. By a union bound the probability that w is added back in any of the future up steps is therefore at most $t\gamma$. Hence the probability that the walk returns to X_τ given that it exited in an intermediate step is at most $t\gamma$, so we get that $\bar{\phi}(v) \in [c, c + t\gamma]$. \square

It's worth noting that the vertex-by-vertex expansion property actually holds exactly for the canonical walks. However, we can only reduce to canonical walks when the height of M is small, an issue we avoid by analyzing lower walks. We now prove Lemma 10.4 by reducing general M to this case.

Proof. [Proof of Lemma 10.4] By repeated application of Lemma 7.1 any k -dimensional HD-walk M on a two-sided γ -local-spectral expander with $k > 0$ can be written as a linear combination of lower products $(U_{k-1}D_k)^i$ up to $O(\gamma)$ error in the following sense:

$$M = \sum_{i=0}^{h(M)} c_i (U_{k-1}D_k)^i + \Gamma,$$

where $(U_{k-1}D_k)^0$ denotes the identity matrix and $\|\Gamma\| \leq w(M)2^{O(h(M))}\gamma$. For simplicity of notation, we write $C := \|\Gamma\|$, which is in turn $\leq 1/2^{13k}$ for γ chosen to be small enough.

We also know by Lemma C.2 that for all i there exists a constant u_i such that for all vertices $v \in X_\tau$, $\phi_{X_\tau}((U_{k-1}D_k)^i, v) = u_i \pm h(M)\gamma$. Hence we get that,

$$1 - \phi_{X_\tau}(M, v) = \sum_{i=0}^j c_i \mathbb{1}_v^T (U_{k-1} D_k)^i \mathbb{1}_{X_\tau} + \mathbb{1}_v^T \Gamma \mathbb{1}_{X_\tau} = \sum_i c_i u_i + \mathbb{1}_v^T \Gamma \mathbb{1}_{X_\tau} \pm w(M)h(M)\gamma.$$

Let $\sum_i c_i u_i = u$. We will show that in expectation over $v \in X_\tau$, $|1 - \phi_{X_\tau}(M, v) - u| \leq |\mathbb{1}_v^T \Gamma \mathbb{1}_{X_\tau}| + w(M)h(M)\gamma$ is small. Define err as the error vector $\Gamma \mathbb{1}_{X_\tau}$. We will bound the 1-norm of err when restricted to coordinates in X_τ . Define the vector s to be $\text{sign}(err(v))$ for $v \in X_\tau$ and 0 otherwise. First note that,

$$\langle s, err \rangle = \sum_{v \in X_\tau} \Pr_{\pi_k}[v] \cdot |err(v)|.$$

Applying Cauchy-Schwarz on the LHS we get that,

$$\begin{aligned} \sum_{v \in X_\tau} \Pi_k(v) |err(v)| &\leq \|s\| \cdot \|err\| \\ &= \|\mathbb{1}_{X_\tau}\| \cdot \|\Gamma \mathbb{1}_{X_\tau}\| \\ &\leq \|\mathbb{1}_{X_\tau}\| \cdot C \|\mathbb{1}_{X_\tau}\| \\ &= C \langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau} \rangle_{\pi_k}. \end{aligned}$$

So we get that,

$$\frac{\sum_{v \in X_\tau} \Pi_k(v) |err(v)|}{\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau} \rangle_{\pi_k}} = \mathbb{E}_{v \sim X_\tau} [|err(v)|] \leq C.$$

Using the fact that $\mathbb{E}_{v \sim X_\tau} [\phi_{X_\tau}(M, v)] = \phi(M, X_\tau)$, we can conclude with a simple trick:

$$|1 - \phi(M, X_\tau) - u| = |\mathbb{E}_{v \sim X_\tau} [1 - \phi_{X_\tau}(M, v) - u]| \leq \mathbb{E}_{v \sim X_\tau} [|err(v)|] + w(M)h(M)\gamma \leq 2C.$$

Since the average deviation of $(1 - \phi_{M, \tau}(v))$ from u is small, the average deviation from the mean $(1 - \phi_M(X_\tau))$ should also be small by triangle inequality:

$$\mathbb{E}_{v \sim X_\tau} [|\phi_{X_\tau}(M, v) - \phi(M, X_\tau)|] \leq \mathbb{E}_{v \sim X_\tau} [|1 - \phi_{X_\tau}(M, v) - u|] + C \leq 3C \leq \frac{1}{2^{11k}}.$$

□

We are now ready to prove Lemma 10.3/Lemma C.4

Proof. [Proof of Lemma 10.3] We will use Lemma 10.4 to relate the global potential on X_τ to the potential induced on X_τ . As a reminder, we proved that for every i -link X_τ , the deviation of the random variable $\phi_{X_\tau}(M, u)$ ($u \sim X_\tau$) is small:

$$\mathbb{E}_{u \sim X_\tau} [|\phi_{X_\tau}(M, u) - \phi(M, X_\tau)|] \leq \frac{1}{2^{11k}}.$$

Let us partition X_τ into two sets, G and \bar{G} , where G contains all vertices u for which $\phi_{M, X_\tau}(u)$ is close to $\phi_M(X_\tau)$. We will henceforth drop the subscripts M, X_τ . Formally we define:

$$G = \left\{ u \in X_\tau \mid |\phi(u) - \phi(X_\tau)| \leq \frac{1}{2^k} \right\}$$

Now note that by Markov's inequality on Lemma 10.4 we have that $\Pr_{u \sim X_\tau}[u \notin G] := c \leq 1/2^{10k}$. For an assignment X , let $Z_{u,s}$ denote $\mathbb{1}[X_u - X'_u = s]$. Expanding out the definition of the global potential on X_τ

(Definition C.2), for any two assignments X, X' we get:

$$\begin{aligned}
\Phi_{\beta,\nu}(X, X')|_{\tau} &= \sum_s \mathbb{E}_{u \sim X_{\tau}} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))] \\
&= \sum_s ((1-c) \mathbb{E}_{u \sim G} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))] + c \mathbb{E}_{u \sim \bar{G}} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))])^2 \\
&\leq \sum_s (1-c)^2 \mathbb{E}_{u \sim G} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))]^2 + \sum_s c(1-c) \mathbb{E}_{u \sim \bar{G}} [Z_{u,s}] \\
&\quad + \sum_s c^2 \mathbb{E}_{u \sim \bar{G}} [Z_{u,s}]^2 \\
&\leq \sum_s (1-c)^2 \mathbb{E}_{u \sim G} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))]^2 + c(1-c) + c^2,
\end{aligned}$$

where $c \leq 1/2^{10k}$ and we used the fact that $p_{\beta,\nu}(y) \leq 1$, when $y \in [0, 1]$, $Z_{u,s} \in [0, 1]$ and $\sum_s Z_{u,s} = 1$.

To relate the two potentials we will relate the quantities $\text{val}_u(X)$ and $\text{val}_u^{\tau}(X)$ for $u \in G$, where $\text{val}_u(X)$ denotes the fraction of edges incident on u that are satisfied by X , and further $\text{val}_u^{\tau}(X)$ denotes the fraction of edges incident on u that lie inside X_{τ} and are satisfied by X . We will use the fact that vertices in G have small expansion. When $r = R_{1-16\eta}(M)$, by Theorem 9.1 the expansion of the r -link X_{τ} is at most $1 - \lambda_{r+1} + w(M)h(M)^2 2^{O(k)}\gamma$ which is less than 17η because of the way the parameters have been set up. Therefore for all vertices u in G , $\phi(u) \leq 18\eta$ since $\eta > 1/2^k$. This implies that for any assignment $X \in \Sigma^V$ and any $u \in G$:

$$\mathbb{1}[\text{val}_u^{\tau}(X) \geq \beta - 18\eta] \geq \mathbb{1}[\text{val}_u(X) \geq \beta],$$

Furthermore by the properties of the polynomials $p_{\beta,\nu}(Y)$, since $\nu < \eta$,

$$p_{\beta-18\eta,\nu}(\text{val}_u^{\tau}(X)) + \nu \geq p_{\beta,\nu}(\text{val}_u(X)) - \nu.$$

Finally expanding out the definition of the potential induced on X_{τ} (Definition C.3), for any two assignments X, X' we get:

$$\begin{aligned}
\Phi_{\beta-18\eta,\nu}^{\tau}(X, X') &= \sum_{s \in \Sigma} \mathbb{E}_{u \sim X_{\tau}} [Z_{u,s} \cdot p_{\beta-18\eta,\nu}(\text{val}_u^{\tau}(X))]^2 \\
&= \sum_{s \in \Sigma} \mathbb{E}_{u \sim X_{\tau}} [Z_{u,s} \cdot (p_{\beta,\nu}(\text{val}_u(X)) - 2\nu)]^2 \\
&\geq \left(\sum_{s \in \Sigma} \mathbb{E}_{u \sim X_{\tau}} [Z_{u,s} \cdot (p_{\beta,\nu}(\text{val}_u(X)))]^2 \right) - 4\nu \\
&\geq \sum_{s \in \Sigma} (1-c)^2 \mathbb{E}_{u \sim G} [Z_{u,s} p_{\beta,\nu}(\text{val}_u(X))]^2 - 4\nu \\
(C.5) \quad &\geq \Phi_{\beta,\nu}(X, X')|_{\tau} - c(1-c) - c^2 - 4\nu,
\end{aligned}$$

where we again used the facts that $Z_{u,s} \in [0, 1]$, $\sum_s Z_{u,s} = 1$ and $p_{\beta,\nu}(y) \in [0, 1]$ when $y \in [0, 1]$. We also have that $c \leq \frac{1}{2^{10k}} \leq \frac{\eta}{56\binom{k}{r}}$ since $\eta > 1/2^k$ and therefore $c \leq \nu$ implying that: $\Phi_{\beta-18\eta,\nu}^{\tau}(X, X') \geq \Phi_{\beta,\nu}(X, X')|_{\tau} - 6\nu$.

Now note that all the inequalities above are sum-of-squares inequalities of degree at most $2 \deg(p) = \tilde{O}(1/\nu)$. We can therefore relate the two potentials when measured with respect to μ , which is a degree- $\tilde{O}(1/\nu)$ pseudodistribution, by applying the pseudoexpectation operator $\tilde{\mathbb{E}}_{\mu}$ to Equation (C.5) above:

$$\Phi_{\beta-18\eta,\nu}^{\tau}(\mu) = \tilde{\mathbb{E}}_{\mu}[\Phi_{\beta-18\eta,\nu}^{\tau}(X, X')] \geq \tilde{\mathbb{E}}_{\mu}[\Phi_{\beta,\nu}(X, X')|_{\tau}] - 6\nu = \Phi_{\beta,\nu}(\mu|_{\tau}) - 6\nu.$$

By the conditions of the lemma, $\Phi_{\beta,\nu}(\mu|_{\tau}) = 1/4\binom{k}{r}$, $\beta = 19\eta$ and $\nu = \frac{\eta}{56\binom{k}{r}}$ we get $\Phi_{\beta,\nu}^{\tau}(\mu) \geq \frac{1}{8\binom{k}{r}}$. This completes the proof of the lemma. \square