Electronic Colloquium on Computational Complexity, Report No. 170 (2020)

High Dimensional Expanders: Random Walks, Pseudorandomness, and Unique Games

Max Hopkins* Tali Kaufman† Shachar Lovett‡

November 7, 2020

Abstract

Higher order random walks (HD-walks) on high dimensional expanders have played a crucial role in a number of recent breakthroughs in theoretical computer science, perhaps most famously in the recent resolution of the Mihail-Vazirani conjecture (Anari et al. STOC 2019), which focuses on HD-walks on *one-sided* local-spectral expanders. In this work we study the spectral structure of walks on the stronger *two-sided* variant, which capture wide generalizations of important objects like the Johnson and Grassmann graphs. We prove that the spectra of these walks are tightly concentrated in a small number of strips, each of which corresponds combinatorially to a level in the underlying complex. Moreover, the eigenvalues corresponding to these strips decay exponentially with a measure we term the *depth* of the walk.

Using this spectral machinery, we characterize the edge-expansion of small sets based upon the interplay of their local combinatorial structure and the global decay of the walk's eigenvalues across strips. Variants of this result for the special cases of the Johnson and Grassmann graphs were recently crucial both for the resolution of the 2-2 Games Conjecture (Khot et al. FOCS 2018), and for efficient algorithms for affine unique games over the Johnson graphs (Bafna et al. Arxiv 2020). For the complete complex, our characterization admits a low-degree Sum of Squares proof. Building on the work of Bafna et al., we provide the first polynomial time algorithm for affine unique games over the Johnson scheme. The soundness and runtime of our algorithm depend upon the number of strips with large eigenvalues, a measure we call High-Dimensional Threshold Rank that calls back to the seminal work of Barak, Raghavendra, and Steurer (FOCS 2011) on unique games and threshold rank.

*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-1909634.

ISSN 1433-8092

1 Introduction

In recent years, high dimensional expanders have begun to play an increasingly important role in our understanding of a wide range of problems within theoretical computer science, including the resolution of the Mihail-Vazirani conjecture [1] and a host of problems across areas such as approximation algorithms [2, 3], sampling [4, 5], agreement testing [6–8], and error correction [9, 10]. These breakthroughs have mostly been driven by the recent introduction of *higher order random walks* (HD-walks) [6, 11, 12, 3], random processes on high-dimensional objects that generalize the standard vertex-edge-vertex walk on expander graphs—walking for instance from a triangle, to a pyramid, and back to a triangle.

Most applications of these walks (e.g. to matroids [1], Glauber Dynamics of Markov chains [4, 5]) rely on a spectral notion of high dimensional expansion known as *one-sided local-spectral expanders* [13, 14]. In this work, we study HD-walks on the stronger *two-sided local-spectral expanders*, a variant introduced by Dinur and Kaufman [11] in the context of agreement testing. Under this stronger setting, HD-walks still capture important structure, offering a broad generalization of non-negative matrices in the Johnson scheme.[1]

The Johnson scheme, historically studied for its connections to coding theory, and its q-analog the Grassmann scheme have seen a recent resurgence due to their connection with the Unique Games Conjecture (UGC). Indeed, it was a deeper understanding of the spectral and combinatorial structure of the Johnson and Grassmann graphs (bases of their respective schemes) that finally lead to the resolution of the 2-2 Games Conjecture [15], completing a long line of work in this direction [16–20]. Conversely, similar structure recently allowed Bafna, Barak, Khotari, Schramm, and Steurer (BBKSS) [21] to give the first polynomial-time algorithm for unique games on the Johnson graphs, raising an interesting interplay between hardness, algorithms, and spectral structure. This breakthrough line of work, however, suffers from a lack of generality: its analysis is rooted in tailoring old Fourier analytic machinery to the Johnson and Grassmann graphs. This specificity results in an unfortunate technical barrier towards further progress on both hardness and algorithms for unique games, where it seems that a broader structural understanding is necessary. In this work, we argue that this barrier may be broken (at least in part) by viewing the Johnson, Grassmann, and related graphs as part of a larger overarching process, a random walk performed on some underlying high-dimensional object.

We study the interplay of local combinatorial and global spectral structure in HD-walks. We prove that the spectra of such walks are tightly concentrated in a small number of strips, each corresponding to a level in the underlying complex, and moreover that the eigenvalues of these strips decay exponentially with a measure we call *depth*. This allows us to give a tight characterization of the edge-expansion of small sets based upon their local combinatorial structure and the eigenvalue decay across strips, a result which in the special cases of the Johnson and Grassmann graphs has been crucially important both for hardness of [16–20] and algorithms for [21] unique games. In greater detail, we show that sets which are *locally-pseudorandom* at level i of the underlying complex expand near-perfectly as long as the ith eigenstrip has *small* eigenvalues. Conversely, we show that sets which are *locally-structured* at level i expand poorly as long as the ith eigenstrip has *large* eigenvalues.

In the special case of the complete complex, our characterization admits a low-degree Sum of Squares proof. Combined with our local-to-global characterization of the expansion of structured sets along with the recent insights of [21], this allows us to provide the first efficient algorithm for affine unique games whose constraint graphs are HD-walks (or equivalently, lie in the Johnson scheme). The soundness and running time of our algorithm depend on a novel spectral parameter we call *High-Dimensional Threshold Rank* (HD-Threshold Rank), which measures the number of strips with eigenvalues above a certain size; this generalizes standard threshold rank, which Barak, Rhagavendra, and Steurer [22] showed in a seminal work to be intimately tied to unique games. To make these results concrete, we study their specification to standard HD-walks from the literature, and show in particular that walks which reach deep into the underlying complex have constant HD-Threshold Rank. In fact, not only are affine unique games particularly easy over such walks, but they carry a stronger characterization of edge-expansion as well, which at a finer-grain level can be seen to depend on the interplay of pseudorandomness and HD-Threshold Rank.

Finally, we note that all of our results extend to a broader set of objects called expanding posets, a generalization of two-sided local-spectral expanders recently introduced by Dikstein, Dinur, Filmus, and

[1]The 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.

Harsha [23]. Since the Grassmann poset is an expanding poset, our results on the expansion of pseudorandom sets extends to a broad generalization of the Grassmann graphs, providing a close connection to Khot, Minzer, and Safra's proof of 2-2 Games [15] (though the characterization is too weak to recover the result itself). For simplicity, in this work we focus solely on the case of two-sided local-spectral expanders, and give the results for general expanding posets and the Grassmann in an upcoming companion paper.

1.1 Local-Spectral Expanders and Higher Order Random Walks

Before discussing our results in greater detail, we first overview the theory of two-sided local-spectral expanders and higher-order random walks.

1.1.1 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 1.1 (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:*

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

We call the elements of $X(i)$ i-faces.[2] *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:*

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

where $\Pi_d = \Pi$.

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

Definition 1.2 (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. 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 link[3] must be a two-sided spectral expander.

Definition 1.3 (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.*[4]

[2] We 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.

[3] The underlying graph of a simplicial complex X is its 1-skeleton $(X(0), X(1))$.

[4] A 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.

1.1.2 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 $f : X(k) \to \mathbb{R}$ 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 [2]). We call an affine combination of pure walks which start and end on $X(k)$ a *k-dimensional HD-walk.*

Definition 1.4 (HD-walk)**.** *Let (X, Π) be a pure, weighted simplicial complex. Let $\mathcal{Y}$ be a family of pure walks $Y : C_k \to C_k$ 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 remains a valid walk (i.e. has non-negative transition probabilities).

Previous work on HD-walks mainly focuses on two natural classes: canonical walks, and partial-swap walks.

Definition 1.5 (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.

Definition 1.6 (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 [2] showed this is the case by expressing them as an alternating hypergeometric sum of canonical walks.

1.1.3 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 1.7 (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:*

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

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.

1.2 Eigenstripping: The Spectral Structure of HD-Walks

It is well known [23, 2] that HD-walks on two-sided local-spectral expanders admit an *approximate eigendecomposition*—a decomposition in which the i-th subspace consists of near-eigenvectors corresponding to some value λ_i. We prove a general linear-algebraic theorem about operators which admit such a decomposition: their true spectra lies in strips tightly concentrated around each λ_i.

Theorem 1.8 (Informal Theorem 2.2: Approximate Eigendecompositions Imply Eigenstripping)**.** *Let M be a self-adjoint operator over an inner product space V, and $V = V^1 \oplus \ldots \oplus V^k$ 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:

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

where $e = O_{k,\lambda}\left(\sqrt{\max_i\{c_i\}}\right)$.

In other words, any sufficiently strong approximate eigendecomposition corresponds to a collection of non-overlapping *eigenstrips*, the span of eigenvectors corresponding to each interval I_{λ_i}. Such behavior was previously known [12] only for the most basic HD-walk, N_k^1, on very strong two-sided local-spectral expanders.

The eigenstrips promised by Theorem 1.8 form their own decomposition closely related (though not necessarily equivalent) to the original V^i. Since in our case the approximate eigendecomposition of interest is combinatorial in nature, Theorem 2.2 will help us to view expansion through both a combinatorial and spectral lens. Indeed, eigenstripping and approximate eigendecompositions are quite useful for understanding expansion. Given such a decomposition on a two-sided local-spectral expander (X, Π), we can express the expansion of a set $S \subset X(k)$ of density α on an HD-walk M as:

$$\Phi(S) = 1 - \frac{1}{\alpha}\langle \mathbb{1}_S, M\mathbb{1}_S \rangle = 1 - \frac{1}{\alpha}\sum_{i=1}^{k}\langle \mathbb{1}_S, M\mathbb{1}_{S,i}\rangle,$$

where $\mathbb{1}_S = \mathbb{1}_{S,1} + \ldots + \mathbb{1}_{S,k}$ is the indicator for S, $\mathbb{1}_{S,i} \in V^i$, and the inner product is understood to be over Π_k. Since the V^i's are closely related to the spectrum of M, we can further write

$$\Phi(S) \approx 1 - \frac{1}{\alpha}\sum_{i=1}^{k}\lambda_i\langle \mathbb{1}_S, \mathbb{1}_{S,i}\rangle, \tag{2}$$

where λ_i is the center of the i-th eigenstrip and corresponds to the approximate eigenvalue of V^i. The expansion of S then hinges on the interplay of two spectral quantities: how the weight of $\mathbb{1}_S$ is distributed on our decompositions, and how the eigenvalues of the underlying walk decay.

1.3 Pseudorandomness and the HD-Level-Set Decomposition

We begin by analyzing the former, which requires understanding the combinatorial structure of HD-walks. HD-walks on sufficiently strong two-sided local-spectral expanders admit a useful combinatorial decomposition due to [23] we call the *HD-Level-Set Decomposition*, which breaks functions on $X(k)$ down by contribution from each level $X(i)$ of the complex for $0 \leq i \leq k$.

Theorem 1.9 (HD-Level-Set Decomposition, Theorem 8.2 [23])**.** *Let (X, Π) be a d-dimensional two-sided γ-local-spectral expander, $\gamma < \frac{1}{d}$, $0 \leq k \leq d$, and let:*

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

Then:

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

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

Extending the work of [23, 2], we show that this decomposition is in fact an approximate eigendecomposition[5] for all HD-walks, and thus by Theorem 1.8 corresponds to a set of disjoint eigenstrips as long as γ is sufficiently small. By leveraging this structure, we prove that for a very broad class of HD-walks (encompassing pure, canonical, partial-swap walks, and more), the eigenvalues corresponding to the HD-Level-Set Decomposition decrease monotonically (see Proposition 4.10). We will assume this is the case for HD-walks we consider throughout the remainder of this section. As a result, for any $0 \leq \ell \leq k$, we can re-write Equation (2) as the bound:

$$\Phi(S) \gtrsim 1 - \alpha - \frac{1}{\alpha}\sum_{i=1}^{\ell}\lambda_i\langle \mathbb{1}_{S,i}, \mathbb{1}_{S,i}\rangle - \lambda_{\ell+1}, \tag{3}$$

where $\mathbb{1}_{S,i} \in V_k^i$, the level-$i$ subspace of the HD-Level-Set Decomposition. Thus we are particularly interested in how the weight of a subset $S \subseteq X(k)$ is distributed on the first ℓ levels of the decomposition.

Intuitively, the HD-Level-Set Decomposition states that the i-th[6] (approximate) eigenspace consists of functions coming from $X(i)$, which manifest as the sum over links of i-faces. One might suspect that the weight of S on V_k^i is then controlled to some extent by its density across i-links. To formalize this, we borrow a notion of pseudorandomness from [20] which measures the distance between global and local densities within S.

[5] Technically, this may require combining V^i with equivalent approximate eigenvalues.
[6] Here we are considering the constant function to be the 0th space.

Definition 1.10 (Pseudorandom Sets)**.** *Let (X,Π) be a weighted, pure simplicial complex. We call a set $S \subset X(k)$ $(\varepsilon_1,\ldots,\varepsilon_\ell)$-pseudorandom if its local expectations at levels 1 through ℓ do not greatly exceed its global expectation. That is if, for all $1 \leq i \leq \ell$, we have:*

$$\forall s \in X(i) : \underset{X_s}{\mathbb{E}}[\mathbb{1}_S] - \mathbb{E}[\mathbb{1}_S] \leq \varepsilon_i,$$

where expectations are taken with respect to Π_k.

As befits our intuition, we prove that functions which are $(\varepsilon_1,\ldots,\varepsilon_\ell)$-pseudorandom have bounded weight on levels $1 \leq i \leq \ell$ (level 0 is the constant part, and thus always has weight $\mathbb{E}[\mathbb{1}_S]^2$).

Theorem 1.11 (Informal Theorem 5.3: Pseudorandomness Controls Eigenweight)**.** *Let (X,Π) be a two-sided γ-local-spectral expander with γ sufficiently small, and $S \subset X(k)$ an $(\varepsilon_1,\ldots,\varepsilon_\ell)$-pseudorandom set of density α. Let $\mathbb{1}_S = \mathbb{1}_{S,1} + \ldots + \mathbb{1}_{S,k}$ be the HD-Level-Set Decomposition of the indicator of S. Then for all $0 < i \leq \ell$, the weight of S on V_k^i is at most:*

$$\langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \leq (1 + O_k(\sqrt{\gamma})) \binom{k}{i} \varepsilon_i \alpha.$$

If the HD-Level-Set is orthogonal, we achieve linear dependence on γ:

$$\langle \mathbb{1}_S, \mathbb{1}_{S,i} \rangle \leq (1 + O_k(\gamma)) \binom{k}{i} \varepsilon_i \alpha.$$

Using Theorem 1.11, we can re-examine Equation (3), our expression of expansion. If $S \subset X(k)$ is an $(\varepsilon_1,\ldots,\varepsilon_\ell)$-pseudorandom set of density α on a sufficiently strong two-sided local-spectral expander, then we can bound the expansion by about:

$$\Phi(S) \gtrsim 1 - \alpha - \sum_{i=1}^{\ell} \lambda_i \binom{k}{i} \varepsilon_i - \lambda_{\ell+1}. \tag{4}$$

Thus we see that as long as α and $\lambda_{\ell+1}$ are small, pseudorandom sets on HD-walks expand near-perfectly. The final piece of the puzzle lies in understanding how the eigenvalues of the HD-Level-Set Decomposition decay.

Before moving to this, however, it is worth pausing to note that Equation (4) provides a tight characterization of expansion for HD-walks on two-sided local-spectral expanders. In particular, for large enough n we can find a subset of any partial-swap walk on the complete complex (i.e. any Johnson graph) with expansion arbitrarily close to Equation (4) (see Proposition 5.4). Conversely, Equation (4) says nothing about sets which are not sufficiently pseudorandom, i.e. when some $\varepsilon_i \gg \binom{k}{i}^{-1}$. In this case we may hope for a bound with worse (sub-linear) dependence on ε_i, but whose coefficient is independent of k. Such a bound is known for the special case of the Johnson graphs [20], but remains open for general HD-walks.

1.4 High-Dimensional Threshold Rank and Deep HD-Walks

Equation (4) raises an interesting dynamic between pseudorandomness and the eigenvalues of the HD-Level-Set Decomposition. To formalize this notion, we can rephrase the equation as follows: for any constant $\delta > 0$, if M has at most r_δ eigenstrips containing eigenvalues greater than δ then:

$$\Phi(S) \gtrsim 1 - \alpha - \delta - \sum_{i=1}^{r_\delta} \lambda_i \binom{k}{i} \varepsilon_i, \tag{5}$$

where S is $(\varepsilon_1,\ldots,\varepsilon_{r_\delta})$-pseudorandom. The interplay between δ and r_δ, a notion we formalize through the introduction of High-Dimensional Threshold Rank (HD-Threshold Rank), is thus crucial for understanding expansion.

Definition 1.12 (High-Dimensional Threshold Rank)**.** *Let M be a linear operator over a vector space V with decomposition $V = \bigoplus_i V^i$ denoted $\mathscr{D}$, where each V^i is the span of some set of eigenvectors. Given $\delta \in \mathbb{R}$, the High-Dimensional Threshold Rank with respect to δ and $\mathscr{D}$ is:*

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

In this work, $\mathscr{D}$ will always correspond to the eigenstrips given by applying Theorem 1.8 to the HD-Level-Set Decomposition, so we will write either $R_\delta(M)$ or just R_δ when M is clear from context.

It is not hard to see that High-Dimensional Threshold Rank is a direct generalization of standard threshold rank, which measures the total number of eigenvalues (with multiplicity) greater than some parameter δ. This is recovered by letting $\mathscr{D}$ be the standard spectral decomposition. For objects like HD-walks which have relatively few eigenstrips, High-Dimensional Threshold Rank may be substantially smaller than threshold rank. In particular, k-dimensional HD-walks have at most $k+1$ eigenstrips despite having up to $n^{O(k)}$ eigenvalues.

Viewing Equation (5) in terms of HD-Threshold Rank gives the following theorem regarding expansion of small, pseudorandom sets on HD-walks.

Theorem 1.13 (Informal Theorem 5.3: Pseudorandom Sets Expand Near-Perfectly)**.** *Given an HD-walk M on a sufficiently strong two-sided local-spectral expander and small constants $\alpha, \delta, \{\varepsilon_i\}_{i=1}^{R_\delta(M)} > 0$, we have that $(\varepsilon_1, \ldots, \varepsilon_{R_\delta(M)})$-pseudorandom sets of density α expand near-perfectly.*

One of the main utilities of this statement lies in its contrapositive: non-expanding sets in HD-walks are correlated with low-level links, where "low-level" is determined by the HD-Threshold Rank of the walk.

Corollary 1.14 (Informal Corollary 5.5: Non-expansion Implies Link Correlation)**.** *Let (X, Π) be a two-sided γ-local-spectral expander with γ sufficiently small and M a k-dimensional HD-walk on (X, Π). Then if $S \subset X(k)$ is a set of density α and expansion:*

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

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

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

It turns out that this viewpoint is crucial both to understanding hardness [15] and algorithms [21] for Unique Games. In the following section we will provide an algorithm for Unique Games on HD-walks over $J(n, d)$ whose soundness and runtime depend on a similar characterization of non-expansion, and as a result depend on the HD-Threshold Rank of the constraint graph.

Since our results depend crucially on R_δ, it is natural to ask the following: *which HD-walks have small High-Dimensional Threshold Rank?* To answer this, we restrict our focus to the main classes of interest in the literature, the canonical and partial-swap walks. In both cases we see that the High-Dimensional Threshold Rank is controlled by the *depth* of the walk.

Definition 1.15 (Depth)**.** *Depth is a parameter in $[0, 1]$ measuring how far a canonical or partial-swap walk reaches into its underlying complex. We say the depth of N_k^j is $\frac{j}{j+k}$ since it walks through j out of $j+k$ levels of the complex. The depth of S_k^j differs slightly due to its restricted nature, and is defined as $\frac{j}{k}$, the fraction of elements it swaps.*

We prove that the (stripped) eigenvalues of the canonical and partial-swap walks decay exponentially, with base dependent on depth.

Theorem 1.16 (Informal Proposition 4.6, Corollary 4.8: Depth Controlled Exponential Eigenvalue Decay)**.** *Let (X, Π) be a two-sided γ-local-spectral expander with γ sufficiently small. Let M be a canonical or partial-swap walk of depth $0 \leq \beta \leq 1$. Then the eigenvalues corresponding to the eigenstrips of M decay exponentially fast:*

$$\lambda_i^{max} \leq e^{-\beta i},$$

where λ_i^{max} is the maximum eigenvalue of M in the eigenstrip corresponding to V_k^i promised by Theorem 1.8. Similarly, the HD-Threshold Rank of M is at most:

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

We compute tight bounds on the exact spectra of these walks in Section 4, along with the general form of spectra for any HD-walk on a two-sided local-spectral expander.

1.5 Local-to-Global: Link Expansion vs Spectra

We have seen how eigenvalue decay and the structure of the HD-Level-Set Decomposition inform the expansion of locally-pseudorandom sets $S \subset X(k)$, but in the case of local-spectral expanders, we are often particularly interested in understanding locally-structured sets as well, i.e. links. Such understanding is crucial in employing the local-to-global paradigm that underlies almost all work on local-spectral expanders. In our case, we show that the local expansion of links is inversely controlled by the corresponding global eigenvalues in the HD-Level-Set Decomposition.

Theorem 1.17 (Informal Theorem 5.6: Local Expansion vs Global Spectrum)**.** *Let (X, Π) be a two-sided γ-local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$, M a k-dimensional HD-walk on (X, Π), and $\lambda_i(M)$ the approximate eigenvalues corresponding to the HD-Level-Set Decomposition. Then for all $0 \leq i \leq k$ and $\tau \in X(i)$:*

$$\Phi(X_\tau) \leq 1 - \lambda_i(M) + O_k(\gamma).$$

When γ is sufficiently small, Theorem 1.8 implies that the $\lambda_i(M)$ correspond to the true global spectrum of M. Thus in a sense, Theorem 1.17 offers a local-to-global approach for situations where standard expansion-based techinques are obstructed by large global eigenvalues. Since the expansion of corresponding links is small, this allows us to operate independently at a local level without losing too much from ignoring edges between links. In the next section, we will see a specific application of this technique to unique games, where Theorem 1.17 allows us to patch together local solutions on links to give a good global solution.

1.6 Playing Unique Games

1.6.1 Background

Our motivation for studying the spectral structure and non-expansion of 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) [24] nearly 20 years ago.

Definition 1.18 (Unique Games)**.** *An instance I of unique games over alphabet Σ is a weighted graph $G(V, E)$, and set of permutations $\Pi = \{\pi_{(u,v)} \in S_\Sigma\}_{(u,v)\in E}$. The value of I, val(I), is the maximum fraction of satisfied constraints over all possible assignments Σ^V:*

$$\max_{x\in\Sigma^V} \mathop{\mathbb{E}}_{(u,v)\sim E}[x_v = \pi_{uv}(x_u)],$$

where edges are drawn corresponding to their weight. For an individual assignment x, we refer to this expectation as val$_I(x)$.

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 [25], vertex-cover [26], and a host of others [24, 27–31]. In this work, we will consider *affine* unique games, a restriction stipulating that Σ is an additive group, and each permutation $\pi_{(u,v)}$ is an additive shift (i.e. $\pi_{(u,v)}(x) = x - a$ for some $a \in \Sigma$). Since the UGC is equivalent to its restriction on affine instances [27], this is not a significant loss in generality.

1.6.2 Algorithms for Unique Games on HD-walks

Recently, Bafna, Barak, Khotari, Schrammm, and Steurer [21] gave the first polynomial time algorithm for affine unique games over the Johnson graphs. Their method, based upon the Sum of Squares semidefinite programming hierarchy (a method for approximating polynomial optimization problems, see Section 6.1), relies on two core structural properties of the Johnson graphs:

1. There exists a low-degree SoS proof that non-expanding sets are concentrated in links.
2. There exists a parameter $r = r(\varepsilon)$ such that:
   (a) The $(r+1)$-st largest (distinct) eigenvalue is small:
   $$\lambda_r \leq 1 - \Omega(\varepsilon)$$
   (b) The expansion of any s-link for any $s < r$ is small:
   $$\forall \tau \in X(s), s < r : \Phi(X_\tau) \leq O(\varepsilon).$$

This second, somewhat cumbersome parameter is found in [21] by direct computation on the Johnson graphs, and controls both the soundness and runtime of BBKSS' algorithm. However, viewing the Johnson graph as an HD-walk, it becomes clear that this behavior actually stems from the local-to-global structure exhibited by Theorem 1.17. As a result, we may replace BBKSS' second parameter with HD-Threshold Rank and informally extend their algorithm to affine unique games over HD-walks[7] M satisfying

1. There exists a low-degree SoS proof that non-expanding sets in M are concentrated in links.
2. M has small HD-Threshold Rank.

More formally, we focus on the case of HD-walks over the complete complex (equivalently, non-negative matrices in the Johnson scheme) where a slight variant of Theorem 1.11 proves that the first condition holds. As a result, we give an algorithm for affine unique games over such constraint graphs whose soundness and runtime depend on HD-Threshold Rank, calling back in a sense to Barak, Steurer, and Rhagavendra's [22] seminal work giving an algorithm for unique games in terms of the constraint graph's standard threshold rank.

Theorem 1.19 (Informal Theorem 6.1: Playing Unique Games on HD-walks)**.** *Let M be a k-dimensional HD-walk on $J(n,d)$ with $n \gg k$, $\varepsilon \in [0, .01)$, and $r(\varepsilon) = R_{1-16\varepsilon}(M)$. Then if I is an instance of affine unique games over M with alphabet Σ satisfying:*

1. $|\Sigma| \geq \Omega\left(\frac{r(2\varepsilon)\binom{k}{r(2\varepsilon)}}{\varepsilon}\right)$,
2. *val*$(I) \geq 1 - \varepsilon$,

there exists an algorithm which outputs an $\Omega\left(\frac{\varepsilon^3}{r(\varepsilon)^2\binom{k}{r(\varepsilon)}^2}\right)$-satisfying assignment in time $|X(k)|^{poly\left(\binom{k}{r(\varepsilon)}, \frac{1}{\varepsilon}\right)}$.

It is worth noting that the High-Dimensional Threshold Rank of a k-dimensional HD-walk is at most $k+1$. Since k is often considered constant with respect to the number of vertices n, Theorem 1.19 gives a polynomial time algorithm for all constraint graphs in the Johnson scheme. To make concrete the finer-grain complexity of the algorithm, we again focus on the canonical and partial-swap walks, which we note are both bases of the Johnson Scheme. As before, the complexity depends on the depth of the walk.

Corollary 1.20 (Corollary 6.2: Unique Games on Deep Walks)**.** *Let I be an instance of affine unique games on a canonical or partial-swap walk of depth $0 \leq \beta \leq 1$ over $J(n,d)$ satisfying the conditions of Theorem 1.19. Then there exists an algorithm that outputs an $\Omega\left(\frac{\beta^2\varepsilon}{\binom{k}{\lceil 20\varepsilon/\beta\rceil}}\right)$-satisfying assignment in time $|X(k)|^{poly\left(\binom{k}{\lceil \varepsilon/\beta\rceil}, 1/\varepsilon\right)}$.*

[7] In reality we will focus only on HD-walks over $J(n,d)$. Extending beyond this case is certainly possible, but requires the walk to be symmetric (i.e. undirected), and to satisfy certain symmetry properties with respect to links which hold for natural classes like the canonical and partial-swap walks.

Thus we see that at a fine-grain level, deeper walks that reach far into the underlying complex perform better than those that only take a few steps. For $\beta = \Theta(1)$ (corresponding to walking $\Omega(k)$ levels into the complex), our algorithm has soundness inverse polynomial in k, and running time exponential in poly(k). On the other hand, for $\beta \leq O(1/k)$ (walks taking only a constant number of steps into the complex), the soundness of our algorithm is inverse exponential, and the runtime doubly exponential in k.

1.7 Beyond Simplicial Complexes: Connections with the UGC

Finally, we briefly discuss the connection of our results to recent progress on the UGC and argue that our framework opens an avenue for further progress. The resolution of the 2-2 Games Conjecture [15] 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 [23] called expanding posets. This class of objects includes expanding subsets of the Grassmann poset we call q-eposets.

Definition 1.21 (q-eposet)**.** *A d-dimensional, weighted, pure q-simplicial complex*[8] *X is a γ-q-eposet if for all $1 \leq i \leq d-1$:*

$$\left\| D_{i+1}U_i - \frac{q-1}{q^{i+1}-1}I - q\frac{q^i-1}{q^{i+1}-1}U_{i-1}D_i \right\| \leq \gamma$$

The particular choice of parameters for q-eposet stems from analysis of the Grassmann poset—in particular the Grassmann poset is an $O_d(1/q^n)$-q-eposet (see [23] for details). Theorem 1.13 and Corollary 1.14 extend naturally to HD-walks on q-eposets. We state the latter result here since it follows without too much difficulty from the arguments in this paper, but the full details (and further generalizations to expanding posets) will appear in our upcoming companion paper.

Corollary 1.22 (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 $1 \leq i \leq r$:

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

Since the Grassmann graphs are simply partial-swap walks on the Grassmann poset[9], Corollary 1.22 provides a direct connection to the proof of the 2-2 Games Conjecture [15]. However, due to the dependence on k, the result is likely too weak to be used in this context as is—improving the bound to be independent of k remains an important open problem. However, we view our method's generality and simplicity as evidence that a deeper understanding of higher order random walks and the spectral structure of the HD-Level-Set Decomposition may be key to further progress on the UGC.

1.8 Related Work

1.8.1 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 [12] 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 [1], but to a number of further breakthroughs in sampling algorithms via a small but consequential improvement on their bound by Alev and Lau [3]. The spectral structure of N_k^1 on the stronger two-sided local-spectral expanders was further studied by Dikstein, Dinur, Filmus, and Harsha (DDFH) [23] who introduced the HD-Level-Set Decomposition, and Kaufman and Oppenheim [6] who introduced a distinct approximate eigendecomposition

[8] This is the q-analog of a simplicial complex, and can be thought of as the downward closure of a set of d-dimensional subspaces of $\mathbb{F}_q^n$.

[9] Seeing that they are HD-walks is non-trivial, and follows from the q-analog of work in [2].

with the benefit of orthogonality (though this came at the cost of additional combinatorial complexity). In recent work, Kaufman and Sharakanski [32] claim that these two decompositions are equivalent on sufficiently strong two-sided γ-local-spectral expanders, but their proof relies on [12, 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.[10] Finally, Alev, Jeronimo, and Tulsiani 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 4, and generalize to expanding posets in the companion paper. For further information on these prior works and their applications, the interested reader should see [33].

1.8.2 Unique Games

The study of unique games has played a central role in hardness-of-approximation since Khot's [24] introduction of the Unique Games Conjecture. One line of work towards *refuting* the UGC focuses on building efficient algorithms for unique games for certain classes of constraint graphs based off of spectral or spectrally-related properties; these include works employing spectral expansion [34, 35], threshold rank [36, 22, 37, 38], hypercontractivity [39], and small-set-expansion or characterized non-expansion [21]. Our work continues to expand this direction with polynomial-time algorithms for (affine) unique games over HD-walks and the introduction of HD-Threshold Rank. On the other hand, recent work towards *proving* the UGC has focused on characterizing non-expanding sets in structures such as the Grassmann [15–20] and Shortcode [19, 15] 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 which we believe may be key to future progress on the UGC.

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 [2] 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 [40] later used cosystolic expanders (a stronger variant) to build explicit instances of 3-XOR that are hard for SoS. While these works are not directly related to ours since the CSP's they study 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.

2 Approximate Eigendecompositions and Eigenstripping

In this section we prove Theorem 1.8: the spectra of any operator with an approximate eigendecomposition is tightly concentrated around the decomposition's approximate eigenvalues. We begin by formalizing what we mean by an approximate eigendecomposition.

Definition 2.1. *Let M be an operator over an inner product space V. We call $V = V^1 \oplus \ldots \oplus V^k$ 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 2.2 (Eigenstripping)**.** *Let M be a self-adjoint operator over an inner product space V, and $V = V^1 \oplus \ldots \oplus V^k$ a $(\{\lambda_i\}_{i=1}^k, \{c_i\}_{i=1}^k)$-approximate eigendecomposition. Let $c_{\max} = \max_i\{c_i\}$, $\lambda_{dif} = \min_{i,j}\{|\lambda_i - \lambda_j|\}$, and $\lambda_{ratio} = \frac{\max_i\{|\lambda_i|\}}{\lambda_{dif}^{1/2}}$. Then as long as $c_{\max}$ is sufficiently small:*

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

[10] It is worth noting that the main results of [12, 32] are unaffected by this error, as an approximate version of [12, Theorem 5.10] sufficient for these results continues to hold.

the spectra of M is concentrated around each λ_i:

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

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

It is worth noting that a version of Theorem 2.2 holds with no assumption on $c_{\max}$, but the assumption substantially simplifies the bounds and is sufficient for our purposes. Before proving Theorem 2.2, we note a useful property of approximate eigendecompositions of self-adjoint operators: they are approximately orthogonal.

Lemma 2.3. *Let M be a self-adjoint operator over an inner product space V. Further, let $V = V^1 \oplus \ldots \oplus V^k$ 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. □

Using Lemma 2.3, we can modify [12, Theorem 5.9] to prove Theorem 2.2. 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 [12, 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 $\|\phi_i\|^2$:

$$\mathop{\mathbb{E}}_{i \sim P_\phi} \left[|\mu - \lambda_i|^2\right] = \frac{1}{\sum\limits_{j=1}^{k} \|\phi_j\|^2} \sum_{i=1}^{k} |\mu - \lambda_i|^2 \|\phi_i\|^2 . \tag{6}$$

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), consider the result of pushing the outer summation inside the norm:

$$\sum_{i=1}^{k} |\mu - \lambda_i|^2 \|\phi_i\|^2 = \left\| \sum_{i=1}^{k} (\mu - \lambda_i)\, \phi_i \right\|^2 - \sum_{1 \leq i \neq j \leq k} (\mu - \lambda_i)(\mu - \lambda_j) \langle \phi_i, \phi_j \rangle . \tag{7}$$

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:

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

The latter takes a bit more effort. Let $\lambda_{\max}$ be $\max_i\{|\lambda_i|\}$, then by Lemma 2.3 we have:

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

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:

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

where the last step follows from Lemma 2.3 and our assumption on $c_{\max}$:

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

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

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

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

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:

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

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 [12, 32] 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 [12, 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 $\mathbb{R}^2$.

3 Pseudorandomness and the HD-Level-Set Decomposition

Now that we have seen how approximate eigendecompositions relate to an operator's spectrum, we take a closer look at the combinatorial structure of the HD-Level-Set Decomposition itself, characterizing how functions project onto each space. In particular, we prove in this section a generalization of Theorem 1.11 to arbitrary functions in C_k. To do this, we will first need to extend our definition of pseudorandomness from sets (boolean functions) to arbitrary functions.

Definition 3.1 (Pseudorandom)**.** *A function $f \in C_k$ is $(\varepsilon_1, \ldots, \varepsilon_\ell)$-pseudorandom if its local expectation is close to its global expectation. That is if, for all $1 \le i \le \ell$, we have:*

$$\forall s \in X(i) : \left| \mathbb{E}_{X_s}[f] - \mathbb{E}[f] \right| \le \varepsilon_i.$$

Our analysis of the projection of such a function onto the HD-Level-Set Decomposition is based upon Garland's method [41], a way of decomposing global information to local information across links. As a result, our (initial) analysis will assume f satisfies a useful local-consistency property.

Definition 3.2. *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.* $\mathbb{E}_{X_s}[f] \neq 0 : \mathit{sign}\left(\mathbb{E}_{X_s}[f]\right) = \mathit{sign}(\mathbb{E}[f])$.

Note that any function may be shifted by some constant to have locally constant sign—this will allow us to generalize our result to all functions. We now state the generalized version of Theorem 1.11, which shows that pseudorandom functions have small projection onto low levels of the HD-Level-Set Decomposition (which, moreover, correspond to the worst eigenvalues for most walks).

Theorem 3.3. *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 $f = f_0 + \ldots + f_k$. If f is $(\varepsilon_1, \ldots, \varepsilon_\ell)$-pseudorandom, then for any $a \in \mathbb{R}^+$ and all $1 \leq i \leq \ell$:*

$$\langle f, f_i \rangle \leq (1 + c(k)\gamma) \frac{a+1}{a} \binom{k}{i} \varepsilon_i^2 + ac(k)\gamma \left\| f \right\|^2,$$

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

$$\langle f, f_i \rangle \leq (1 + c(k)\gamma) \frac{a+1}{a} \binom{k}{i} \varepsilon_i |\mathbb{E}[f]| + ac(k)\gamma \left\| f \right\|^2.$$

Finally, if the HD-Level-Set Decomposition is orthogonal, we can be rid of a. For arbitrary functions we have:

$$\langle f, f_i \rangle \leq (1 + c_1(k,i)\gamma) \binom{k}{i} \varepsilon_i^2,$$

where $c_1(k,i) = O\left(k^2 \binom{k}{i}\right)$, and if f has locally constant sign then:

$$\langle f, f_i \rangle \leq (1 + c_1(k,i)\gamma) \binom{k}{i} \varepsilon_i |\mathbb{E}[f]|.$$

It's worth noting that due the the approximate orthogonality of the HD-Level-Set Decomposition (see Lemma 3.6), this bound holds in absolute value as well since $\langle f, f_i \rangle$ cannot be too negative. It is also worth noting that in the case that the decomposition is orthogonal, we 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.

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, setting $a = 1/\sqrt{\gamma}$ implies Theorem 1.11:

$$\langle f, f_i \rangle \leq \left(1 + 2^{O(k)}\sqrt{\gamma}\right) \binom{k}{i} \varepsilon_i \mathbb{E}[f].$$

This is particularly useful when considering expansion (which we recall may be written as $1 - \frac{1}{\mathbb{E}[f]} \langle f, Mf \rangle$) to cancel the normalization by $\mathbb{E}[f]$–indeed, in Section 5 we will show this gives a tight characterization. Finally, it should be noted that in the case f is non-negative, the absolute value can be removed from the definition of pseudorandomness in this argument, which will later allow us to show that non-expanding sets are locally denser than expected.

We prove Theorem 3.3 through three main steps. First and foremost, we show how weight on V_k^ℓ (the ℓ-th HD-Level-Set Decomposition space) implies an imbalance between local and global expectation at level $X(\ell)$ for any function f with locally constant sign. Second, we analyze how this generalizes to general functions via a simple constant shift. Finally, we simplify the bound by analyzing the relation between $\langle f, f_\ell \rangle$ and $\|g_\ell\|$ based on the machinery of DDFH [23]. We start with the core of our result: any non-trivial projection onto V_k^ℓ implies a disparity between the local and global behaviour of f.

Proposition 3.4. *Let (X, Π) be a two-sided γ-local-spectral expander with $\gamma < 1/k$, and let $f \in C_k$ be a function on k-faces with HD-Level-Set Decomposition $f = f_0 + \ldots + f_k$. Then for all $0 < \ell \leq k$, if f has ℓ-local constant sign, there exists a face $s \in X(\ell)$ such that the difference between local and global expectation is lower bounded by the weight of f's projection onto V_k^ℓ:*

$$\left| \underset{X_s}{\mathbb{E}}[f] - \mathbb{E}[f] \right| \geq \left| \frac{\langle f, f_\ell \rangle^2}{\|g_\ell\|^2 \, \mathbb{E}[f]} \right|,$$

where we recall $g_\ell \in H^\ell$ satisfies $f_\ell = U_\ell^k g_\ell$. If f is non-negative, the inequality holds without absolute value signs.

Proof. We begin by examining the squared norm of E_f^ℓ, the vector of link expectations over $X(\ell)$:

$$\forall s \in X(\ell) : E_f^\ell(s) = \underset{X_s}{\mathbb{E}}[f].$$

Let $\Pi_k(X_s)$ be shorthand for $\sum\limits_{t \in X_s} \Pi_k(t)$ (i.e. the normalization factor for the above restricted expectation). Since f has locally constant sign, we may rewrite the squared norm of E_f^ℓ as an expectation over a related distribution P_ℓ:

$$\frac{1}{\mathbb{E}[f]} \langle E_f^\ell, E_f^\ell \rangle = \sum_{s \in X(\ell)} \Pi_\ell(s) \left(\frac{1}{\mathbb{E}[f]} \sum_{t \in X_s} \frac{\Pi_k(t) f(t)}{\Pi_k(X_s)} \right) E_f^\ell(s) \tag{8}$$

$$= \sum_{s \in X(\ell)} \left(\frac{1}{\mathbb{E}[f]} \sum_{t \in X_s} \frac{\Pi_k(t) f(t)}{\binom{k}{\ell}} \right) E_f^\ell(s) \tag{9}$$

$$= \underset{P_\ell}{\mathbb{E}}[E_f^\ell], \tag{10}$$

where we have used the fact that $\Pi_k(X_s) = \binom{k}{\ell} \Pi_\ell(s)$ by Equation (1). To understand $P_\ell(s)$ 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_ℓ on $X(\ell)$ via the following process: draw a face $t \in X(k)$ from P_k, and then choose a ℓ-face $s \subset t$ uniformly at random. Replacing the non-negativity of f with the conditions in the theorem statement still leaves $P_\ell(s)$ a valid distribution, albeit one with a less intuitive description.

If we can lower bound $\langle E_f^\ell, E_f^\ell \rangle$, Equation (10) then implies a lower bound on $\max\limits_s(|E_f^\ell(s)|)$ by averaging. In particular, it is enough to show that:

$$\langle E_f^\ell, E_f^\ell \rangle \geq \frac{\langle f, f_\ell \rangle^2}{\|g_\ell\|^2} + \mathbb{E}[f]^2. \tag{11}$$

To see why this is sufficient, assume Equation (11) and that $\mathbb{E}[f] > 0$ (the negative case follows similarly). Then we have:

$$\underset{P_\ell}{\mathbb{E}}[E_f^\ell] \geq \frac{\langle f, f_\ell \rangle^2}{\|g_\ell\|^2 \, \mathbb{E}[f]} + \mathbb{E}[f],$$

and further that by averaging there exists an ℓ-face $s \in X(\ell)$ satisfying the desired property:

$$\underset{X_s}{\mathbb{E}}[f] - \mathbb{E}[f] \geq \frac{\langle f, f_\ell \rangle^2}{\|g_\ell\|^2 \, \mathbb{E}[f]}.$$

To show Equation (11), consider splitting the norm into two parts:

$$\langle E_f^\ell, E_f^\ell \rangle = \frac{\left\langle g_\ell, E_f^\ell \right\rangle^2}{\langle g_\ell, g_\ell \rangle} + \mathbb{E}\left[\left(E_f^\ell - \frac{\langle g_\ell, E_f^\ell \rangle}{\langle g_\ell, g_\ell \rangle} g_\ell\right)^2\right] \tag{12}$$

$$\geq \frac{\left\langle g_\ell, E_f^\ell \right\rangle^2}{\langle g_\ell, g_\ell \rangle} + \mathbb{E}\left[E_f^\ell - \frac{\langle g_\ell, E_f^\ell \rangle}{\langle g_\ell, g_\ell \rangle} g_\ell\right]^2 \tag{13}$$

$$= \frac{\left\langle g_\ell, E_f^\ell \right\rangle^2}{\langle g_\ell, g_\ell \rangle} + \mathbb{E}[f]^2, \tag{14}$$

where the last step follows from recalling that g_ℓ is in the kernel of the down operator (and hence $\mathbb{E}[g_\ell] = 0$), and $\mathbb{E}[E_f^\ell] = \mathbb{E}[f]$. To relate this to the projected weight $\langle f, f_\ell \rangle$, we use Garland's method [41] to show that the numerator of the left-hand term is exactly $\langle f, f_\ell \rangle^2$. For any $0 \leq i \leq k$ and $u \in X(i)$, let $y_u \in C_k$ be the indicator for the link X_u: $y_u(t) = \mathbb{1}(t \supset u)$. Then:

$$\begin{aligned}
\langle g_\ell, E_f^\ell \rangle &= \sum_{s \in X(\ell)} \Pi_\ell(s) g_\ell(s) \mathop{\mathbb{E}}_{X_s}[f] \\
&= \sum_{s \in X(\ell)} g_\ell(s) \sum_{t \in X_s} \frac{\Pi_k(t)}{\binom{k}{\ell}} f(t) \\
&= \sum_{s \in X(\ell)} g_\ell(s) \sum_{t \in X(k)} \frac{\Pi_k(t)}{\binom{k}{\ell}} f(t) y_s(t) \\
&= \left\langle \sum_{t \in X(k)} f(t) y_t, \sum_{s \in X(\ell)} g_\ell(s) U_\ell^k y_s \right\rangle \\
&= \langle f, U_\ell^k g_\ell \rangle \\
&= \langle f, f_\ell \rangle
\end{aligned}$$

Plugging this into Equation (14) completes the proof. □

As an immediate corollary, we can upper bound the projection of any function onto the HD-Level-Set Decomposition in terms of its pseudorandomness.

Corollary 3.5. *Let (X, Π) be a γ-local-spectral expander, and let $f \in C_k$ be a function on k-faces with HD-Level-Set Decomposition $f = f_0 + \ldots + f_k$. If f is $(\varepsilon_1, \ldots, \varepsilon_\ell)$-pseudorandom, then for all $1 \leq i \leq \ell$:*

$$\langle f, f_i \rangle^2 \leq \varepsilon_i^2 \left\| g_i \right\|^2.$$

Further, if f has i-local constant sign, then:

$$\langle f, f_i \rangle^2 \leq \varepsilon_i |\mathbb{E}[f]| \left\| g_i \right\|^2.$$

Proof. The latter bound follows immediately from Proposition 3.4. For the former, assume for simplicity that $\mathbb{E}[f] \geq 0$ (the negative case follows from a similar argument). Notice that as long as $\varepsilon_i \neq 0$, $f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}$ has positive expectation over all i links and non-zero expectation, allowing us to apply Proposition 3.4. Note further that in the HD-Level-Set Decomposition, f_0 corresponds to the constant part of f, and we may thus similarly decompose $f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}$ as:

$$f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1} = f_0' + f_1 + \ldots + f_k,$$

where $f_0' = f_0 + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}$. Applying Proposition 3.4 then gives:

$$\langle f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}, f_i \rangle^2 \leq \varepsilon_i \mathbb{E}[f + (\varepsilon_i - \mathbb{E}[f])\mathbb{1}] \left\| g_i \right\|^2.$$

Noting that, for $i > 0$, f_i is orthogonal to $\mathbb{1}$ then gives the desired result. We are left to deal with the case that $\varepsilon_i = 0$, which follows from a limiting argument applying the above to any $\varepsilon > 0$. □

A priori, it is not clear that Corollary 3.5 is particularly useful due to its dependence on $\|g_i\|$. In fact, $\|g_i\|^2$ is closely related to both $\|f_i\|^2$ and $\langle f, f_i\rangle$ on two-sided local-spectral expanders, a fact proved by Dikstein, Dinur, Filmus, and Harsha [23].

Lemma 3.6 (Lemmas 8.10, 8.13, Theorem 4.6 [23])**.** *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 $f_0 + \ldots + f_k$. Then for all $0 \leq \ell \leq k \leq d$:*

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

where $c_1(k,\ell) = O(k^2\binom{k}{\ell})$. Further for all $0 \leq i \neq j \leq k$:

$$\langle f_i, f_j\rangle \leq 2^{O(k)} \|f_i\| \|f_j\| ,$$

and if $\gamma \leq 2^{-\Omega(k)}$, then:

$$\|f_\ell\|^2 \leq \frac{\|f\|^2}{1 - 2^{O(k)}\gamma}.$$

The version of this result appearing in [23] does not have explicit coefficients, but they follow from direct computation. With these in hand, the proof of Theorem 3.3 amounts to a few lines of computation.

Proof of Theorem 3.3. For both the orthogonal and non-orthogonal cases we prove only the result for arbitrary functions. The result for i-local constant sign follows from exactly the same arguments. We start with the case where the HD-Level-Set Decomposition is orthogonal. Here the result is almost immediate from combining Corollary 3.5 and Lemma 3.6:

$$\langle f, f_i\rangle \leq \frac{\varepsilon_i^2 \|g_i\|^2}{\langle f, f_i\rangle} = \frac{\varepsilon_i^2 \|g_i\|^2}{\|f_i\|^2} \leq \frac{\varepsilon_i^2\binom{k}{i}}{(1 - c_1(k,i)\gamma)}$$

where $c_1(k,i) = O(k^2\binom{k}{i})$. Since we assume $\gamma \leq 2^{-\Omega(k)}$, Taylor expanding $\frac{1}{(1-c_1(k,i)\gamma)}$ gives the desired result.

When the HD-Level-Set Decomposition is not orthogonal, we cannot cancel terms on the right and left-hand side, but still have the relation:

$$\begin{aligned}
\langle f, f_i\rangle^2 &\leq \frac{\varepsilon_i^2\binom{k}{i}}{(1 - c_1(k,i)\gamma)} \|f_i\|^2 \\
&= \frac{\varepsilon_i^2\binom{k}{i}}{(1 - c_1(k,i)\gamma)}\langle f, f_i\rangle - \frac{\varepsilon_i^2\binom{k}{i}}{(1 - c_1(k,i)\gamma)} \sum_{j\neq i}\langle f_i, f_j\rangle \\
&\leq (1 + c_2(k,i)\gamma)\varepsilon_i^2\binom{k}{i}\left(\langle f, f_i\rangle + c_3(k,i)\gamma\|f\|^2\right),
\end{aligned}$$

where $c_2(k,i), c_3(k,i) \leq 2^{O(k)}$, and the last step follows from the approximate orthogonality given by Lemma 3.6 and a Taylor expansion. To get the final result, notice that we are done if $\langle f, f_i\rangle \leq a2^{O(k)}\gamma\|f\|^2$. Otherwise, we can write:

$$\langle f, f_i\rangle^2 \leq (1 + c_2(k,i)\gamma)\varepsilon_i^2\binom{k}{i}\left(\langle f, f_i\rangle + \frac{1}{a}\langle f, f_i\rangle\right),$$

which implies the desired result. □

4 The Spectra of HD-walks

We now show that the HD-Level-Set Decomposition is an approximate eigendecomposition for any HD-Walk, and thus by Theorem 2.2 corresponds to a decomposition of the walk's spectrum into tightly concentrated 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 4.1 (k-Dimensional Pure Walk)**.** *Given a weighted, simplicial complex (X, Π), a k-dimensional pure walk $Y : C_k \to C_k$ 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 4.2 (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 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 [23], who show that the up and down operators on two-sided γ-local-spectral expanders satisfy the following relation:

$$\left\| D_{i+1} U_i - \frac{1}{i+1} I - \frac{i}{i+1} U_{i-1} D_i \right\| \leq \gamma. \tag{15}$$

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

Lemma 4.3 (Claim 8.8 [23])**.** *Let (X, Π) be a d-dimensional γ-local-spectral expander. Then for all $j < k < d$:*

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

Indeed Lemma 3.6 is proved through an inductive application of this lemma, which will also be critical for proving that the V_k^i are approximate eigenspaces. In Appendix A, we prove a stronger version of both Lemma 4.3 and Lemma 3.6 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 [23] throughout the body of the paper. Using Lemma 4.3 and Lemma 3.6, 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 4.4. *Let (X, Π) be a two-sided γ-local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $Y : C_k \to C_k$ a pure walk:*

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

Let $i_1 \leq \ldots \leq i_{h(Y)}$ 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 $Y_j^b : C_\ell \to C_{\ell+b}$ denote an unbalanced walk with j down operators, and $j + b$ up operators. If Y_j^b has down operators in positions $i_1 \leq \ldots \leq i_j$ and $g_\ell \in H^\ell$, we claim:

$$\left\| Y_j^b g_\ell - \prod_{s=1}^{j} \left(1 - \frac{\ell}{\max\{\ell, i_s - 2s + \ell + 1\}}\right) U_\ell^{b+\ell} g_\ell \right\| \leq \gamma j(b+j) \|g_\ell\| . \tag{16}$$

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 $i_1 + k - \ell \leq \ldots \leq i_j + k - \ell$. Combining Equation (16) with Lemma 3.6 then implies the result.

We prove Equation (16) 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 4.3 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 4.3, Y_{j-1}^b has down operators at indices $i_2 - 2 \leq \ldots \leq i_j - 2$. 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 4.5. *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 4.4.

It is worth noting that the resulting approximate eigenvalues in Corollary 4.5 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 [2], where a weaker notion of approximate eigenspaces based on the quadratic form $\langle f, Mf\rangle$ is analyzed. Plugging Corollary 4.5 into Theorem 2.2, 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 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.

Proposition 4.6 (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:*

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

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

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

Proof. By Proposition 4.4, N_k^j is an $\left(\{\lambda_\ell\}_{\ell=0}^k, \{c(k,\ell,j)\}_{\ell=0}^k\right)$-approximate eigendecomposition for

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

where $i_1 \leq \ldots \leq i_s$ 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 $\mathrm{Spec}(N_k^j)$ follow immediately from plugging the above into Theorem 2.2. □

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 [2] proved that partial-swap walks may be written as a alternating hypergeometric sum of canonical walks.

Proposition 4.7 (Corollary 4.13 [2])**.** *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 Proposition 4.6 to bound their approximate eigenvalues and true spectrum.

Corollary 4.8. *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:*

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

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

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

Proof. By Corollary 4.5, $\bigoplus_{\ell=0}^{k}V_k^\ell$ is a $\left(\{\lambda_\ell\}_{\ell=0}^{k},\{c'(k,\ell,j)\}_{\ell=0}^{k}\right)$-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 4.7. The bounds on $\mathrm{Spec}(S_k^j)$ then follow from Theorem 2.2. ☐

Together, Proposition 4.6 and Corollary 4.8 prove Theorem 1.16 (assuming γ is sufficiently small).

In the introduction, we mentioned that for a broad class of HD-walks, the approximate eigenvalues of the HD-Level-Set Decomposition exhibit a further property key for understanding expansion: monotonic decay. We now formally define this class and prove the desired property.

Definition 4.9 (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 4.5 are independent of the underlying complex.

Proposition 4.10. *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 2.2. 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 4.5 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 $\sum\limits_{Y \in \mathcal{Y}} \alpha_Y Y$ 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 2.2, it must be the case that for every $1 \leq \ell \leq k$ and n sufficiently large, the intersection of $\sum\limits_{Y \in \mathcal{Y}} \alpha_Y \lambda_\ell(Y) \pm O(1/n)$ and $\sum\limits_{Y \in \mathcal{Y}} \beta_i \lambda_\ell(S_k^i) \pm O(1/n)$ is non-empty. Since we may take n arbitrarily large, this implies the two quantities are in fact equivalent. Finally, by Corollary 4.8 $\lambda_\ell(S_k^i)$ decreases monotonically in ℓ for all i, which implies that the $\lambda_i(M) = \sum \alpha_Y \lambda_i(Y) = \sum \beta \lambda_i(S_k^j)$ decrease monotonically as desired. □

5 Expansion of HD-walks

In this section we prove Theorem 1.13, Corollary 1.14, and Theorem 1.17, capturing the tradeoffs between local structure, HD-Threshold Rank, expansion, and non-expansion. We first recall the definitions of expansion and HD-Threshold Rank from the introduction specified to HD-walks for simplicity.

Definition 5.1 (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*

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

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 will control the expansion of sets in HD-walks partially through the walk's HD-Threshold Rank, a measure of how many eigenstrips contain large eigenvalues.

Definition 5.2 (High-Dimensional 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 $C_k = \bigoplus W_k^i$. The HD-Threshold-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.

It 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 HD-Threshold Rank of M.

Theorem 5.3. *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 2.2 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:*

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

where $\lambda_i(M)$ is the approximate eigenvalue given by Corollary 4.5, S is $(\varepsilon_1, \ldots, \varepsilon_r)$-pseudorandom, and $c \leq w(M)h(M)^2 2^{O(k)}$. Further, if the HD-Level-Set Decomposition is orthogonal, then:

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

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

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

where by Corollary 4.5 $\|\Gamma_i\| \leq O\left(w(M)h(M)(h(M)+k)\binom{k}{i}\gamma \|\mathbb{1}_{S,i}\|\right)$. Using Cauchy-Schwarz and the fact that $\|\mathbb{1}_{S,i}\| \leq (1 + 2^{O(k)}\gamma) \|\mathbb{1}_S\|$ we can simplify this to

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

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:

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

where $e_2 \leq w(M)h(M)^2 2^{O(k)}$ and error from the rightmost term has been absorbed into e_2. Finally, applying the corresponding version of Theorem 3.3 (with $a = O(1/\sqrt{\gamma})$) depending on whether the HD-Level-Set Decomposition is orthogonal and combining $e_2\gamma$ with the resulting error term gives the desired results. $\square$

Theorem 5.3 gives a tight characterization of expansion in the regime of linear dependence on ε_i. In particular, we can find a copy of $J(m,k,t) \subset J(n,k,t)$ whose expansion is arbitrarily close to the bound of Theorem 5.3 by taking $n \gg m$.

Proposition 5.4. *Let $X = J(n,d)$ be the Johnson complex, $2k - t \leq d$, $m|n$, and B_m be the the set of all k-faces $\binom{[n/m]}{k}$. Then for any t, Theorem 5.3 is tight for S_k^{k-t} as $n, m \to \infty$.*

Proof. It is not hard to see by direct computation 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 $(\varepsilon_1, \ldots, \varepsilon_k)$-pseudorandom, where:

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

Since the Johnson Complex is a two-sided $O(1/n)$-local-spectral expander whose HD-Level-Set Decomposition is orthgonal [23], for large enough n Theorem 5.3 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. □

Note that this tightness does not preclude a version of Theorem 5.3 with sub-linear dependence on ε_i and a corresponding coefficient better than $\binom{k}{i}$, or even independent of k. Such a bound is known for the Johnson graphs [20], but is difficult to extend to two-sided local-spectral expanders due to its reliance on symmetry. The difference between the two regimes is perhaps most stark when examining the contrapositive of Theorem 5.3, which states that non-expanding sets must be concentrated inside links.

Corollary 5.5. *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 2.2. Then for any $\delta > 0$, if $S \subset X(k)$ is a set of density α and expansion:*

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

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

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

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

Notice that the excess correlation implied by Corollary 5.5 shrinks with k (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 5.5 for hardness of unique games. Moreover, since we proved in Proposition 5.4 that this is unavoidable in the regime of linear dependence on ε_i, further progress likely requires new techniques beyond analyzing our strategy of analyzing L_2-mass (such as the hypercontractive analysis of [20]). On the other hand, the regime we consider does turn out to useful for considering *algorithms* for unique games.

5.1 Link Expansion: Local vs Global

Before presenting such an algorithm, however, we first discuss what is, in a sense, a converse to the above: a characterization of expansion for *structured* sets (links). Understanding this structure will allow us to solve global problems like unique games by operating at a local scale. In particular, we show that the local expansion of links in HD-walks is inversely related to the walk's global spectra.

Theorem 5.6 (Local Expansion vs Global Spectra)**.** *Let (X, Π) be a two-sided γ-local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$, and M a k-dimensional, complete HD-walk. Then for all $0 \leq i \leq k$ and $\tau \in X(i)$:*

$$\Phi(X_\tau) \leq 1 - \lambda_i(M) + c\gamma,$$

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

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

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

where α is the density of X_τ and $\mathbb{1}_{X_\tau,s} \in V_k^s$. Notice that because X_τ is a link, it comes from level at most i:

$$\mathbb{1}_{X_\tau} = \binom{k}{i} U_i^k \mathbb{1}_\tau \in V_k^0 \oplus \ldots \oplus V_k^i$$

Then similar to Theorem 5.3, we may write:

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

where by Corollary 4.5 $\|\Gamma_s\| \leq O\left(w(M)h(M)(h(M)+k)\binom{k}{s}\gamma\,\|\mathbb{1}_{X_\tau,s}\|\right)$. Since the approximate eigenvalues of M decrease monotonically, we can further write:

$$\begin{aligned}\Phi(X_\tau) &\leq 1 - \frac{\lambda_i(M)}{\alpha}\sum_{s=0}^{i}\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau,s}\rangle + \frac{1}{\alpha}\sum_{s=0}^{i}\langle \mathbb{1}_{X_\tau}, \Gamma_s\rangle, \\ &= 1 - \lambda_i(M) + \frac{1}{\alpha}\sum_{s=0}^{i}\langle \mathbb{1}_{X_\tau}, \Gamma_s\rangle + c\gamma,\end{aligned}$$

where $c \leq w(M)h(M)^2 2^{O(k)}$ and we have dealt with the error term as in the proof of Theorem 5.3. □

Theorem 5.6 will play a crucial role in our algorithm for unique games, allowing us to patch together local solutions over links corresponding to bad eigenvalues of the constraint graph. We believe this paradigm of exploiting local expansion corresponding to large eigenvalues is of independent interest, and may be useful for related problems like agreement testing and PCPs.

6 Playing Unique Games on HD-Walks

Recently, BBKSS [21] proposed a polynomial-time algorithm based on the Sum of Squares (SoS) semidefinite programming hierarchy for affine unique games on the Johnson graphs. Their strategy relies on two core structural properties of the underlying constraint graphs:

1. There exists a low-degree SoS proof that non-expanding sets are concentrated in links.
2. There exists a parameter $r = r(\varepsilon)$ such that:
   (a) The $(r+1)$-st largest (distinct) eigenvalue is small:
   $$\lambda_r \leq 1 - \Omega(\varepsilon)$$
   (b) The expansion of any s-link with $s < r$ is small:
   $$\forall \tau \in X(s), s < r : \Phi(X_\tau) \leq O(\varepsilon).$$

This second parameter, found in [21] by direct computation on the Johnson graphs, determines both the soundness and runtime of their algorithm. However, in the context of our framework it is not hard to see that the parameter's existence is an inherent consequence of Theorem 5.6, and furthermore that it is exactly the HD-Threshold-Rank of the underlying constraint graph. Combined with an SoS variant of Theorem 5.3 for the complete complex, this allows us to build an algorithm for affine unique games over any HD-walk on the complete complex with soundness and runtime dependent on HD-Threshold Rank.

Theorem 6.1. *Let M be a k-dimensional HD-walk on $X = J(n,d)$, $n \geq 2^{\Omega(k)}$, $\varepsilon \in [0, .01)$, and $r(\varepsilon) = R_{1-16\varepsilon}(M)$. Then given an instance of affine unique games over M with alphabet Σ such that:*

1. $|\Sigma| \geq \Omega\left(\frac{r(2\varepsilon)\binom{k}{r(2\varepsilon)}}{\varepsilon}\right)$,
2. *val*$(I) \geq 1 - \varepsilon$,

there exists an algorithm outputting an $\Omega\left(\frac{\varepsilon^3}{r(2\varepsilon)^2\binom{k}{r(2\varepsilon)}^2}\right)$*-satisfying assignment in time* $|X(k)|^{poly\left(\binom{k}{r(\varepsilon)}, \frac{1}{\varepsilon}\right)}$.

For concreteness, we examine the specification of Theorem 6.1 to standard HD-walks, the canonical and partial-swap walks. While for fixed k Theorem 6.1 provides a polynomial time algorithm for affine unique games for all HD-walks, we see its fine-grain performance depends on the depth of the walk.

Corollary 6.2. *Let $X = J(n,d)$, $n > 2^{\Omega(k)}$, $\varepsilon \in [0, .01)$. Then there exists a universal constant $c > 0$ such that if I is an instance of affine unique games on a k-dimensional canonical or partial-swap walk of depth $0 \leq \beta \leq 1$ with alphabet size at least* $|\Sigma| \geq \Omega\left(\frac{\binom{k}{c\frac{\varepsilon}{\beta}}}{\beta}\right)$ *and value at least $1 - \varepsilon$, there exists an algorithm outputting an* $\Omega\left(\frac{\beta^2\varepsilon}{\binom{k}{c\frac{\varepsilon}{\beta}}^2}\right)$*-satisfying assignment in time* $|X(k)|^{poly\left(\binom{k}{\frac{\varepsilon}{\beta}}, \frac{1}{\varepsilon}\right)}$.

Theorem 6.1 relies on the Sum of Squares semidefinite programming hierarchy and its relation to Unique Games, which we now overview before giving the proof.

6.1 Background: Sum of Squares and Unique Games

Proving Theorem 6.1 from the ground up requires substantial and non-trivial background in the SoS framework. However, since we mostly rely on a number of higher level results from [21] for the SoS side of our work, we cover here only background necessary to understand our methods, and refer the reader to Sections 1, 2, and A of [21] for additional information.

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

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

for $q_i \in \mathbb{R}[x_1, \ldots, x_n]$, the Degree-$D$ Sum of Squares semidefinite programming relaxation outputs in time $n^{O(D)}$ a *pseudo-expectation operator* $\tilde{\mathbb{E}} : X^{\leq D} \to \mathbb{R}$ over monomials in $\mathbb{R}[x_1, \ldots, x_n]$ 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 pseudo-expectation (under constraints $\{q_i = 0\}_{i=1}^m$), whereas the fifth is promised by the SoS relaxation.

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 a degree-D pseudo-expectation. 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 are at most degree-D. Notice that properties 2, 3, and 4 then immediately imply $\tilde{\mathbb{E}}[f(x)] \leq \tilde{\mathbb{E}}[g(x)]$.

Unique games can be written as a polynomial optimization problem. In particular, given an instance I of unique games with alphabet Σ and constraints Π over $G(V, E)$, consider the following quadratic optimization problem over variables $\{X_{v,s}\}_{V \times \Sigma}$ that computes $\mathrm{val}(I)$:

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

Following BBKSS, we will call the constraints of this program $\mathcal{A}_{\mathcal{I}}$. We will work with the Degree-D Sum of Squares relaxation of this program, which outputs a degree-D pseudo-expectation satisfying $\mathcal{A}_{\mathcal{I}}$ in time $|V|^{O(D)}$ such that $\tilde{\mathbb{E}}[\mathrm{val}_I(x)] \geq \mathrm{val}(I)$. Throughout our proof, we will modify this pseudo-expectation in two ways.

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. [22, 21]). Given a pseudo-expectation $\tilde{\mathbb{E}}$ and a sum of squares polynomial $s(x)$, we can define a new pseudo-expectation by *conditioning* on $s(x)$ as follows:

$$\tilde{\mathbb{E}}[f(x)|s(x)] = \frac{\tilde{\mathbb{E}}[f(x)s(x)]}{\tilde{\mathbb{E}}[s(x)]}.$$

In this work, we will only be interested in the case where $s(x)$ is the indicator of some binary variable. In other words, in our context this process can be thought of as conditioning on the value of some $X_{u,a}$.

Symmetrization is an operation on pseudo-expectations introduced in [21] to take advantage of the symmetric structure of affine unique games. The idea is to symmetrize over shifts $s \in \Sigma$. Formally, given a pseudo-expectation $\tilde{\mathbb{E}}$, define the s-shifted pseudo-expectation $\tilde{\mathbb{E}}_s$ to be:

$$\tilde{\mathbb{E}}_s[X_{u_1,a_1} \cdots X_{u_t,a_t}] = \tilde{\mathbb{E}}[X_{u_1,a_1-s} \cdots X_{u_t,a_t-s}].$$

BBKSS then define the symmetrization operator which maps $\tilde{\mathbb{E}}$ to:

$$\tilde{\mathbb{E}}_{\mathrm{sym}} = \frac{1}{|\Sigma|} \sum_{s\in\Sigma} \tilde{\mathbb{E}}_s,$$

and call a pseudo-expectation shift-symmetric if it is invariant under this operation. Let $\tilde{\mathbb{E}}$ be the pseudodistribution output by the degree-D SoS relaxation of unique games. BBKSS note three important properties:

1. $\tilde{\mathbb{E}}_{\mathrm{sym}}$ is a degree-D pseudo-expectation satisfying $\mathcal{A}_{\mathcal{I}}$.

2. Symmetrization can be performed in time subquadratic in the descrption of $\tilde{\mathbb{E}}$.

3. The objective value is invariant under symmetrization: $\tilde{\mathbb{E}}[\text{val}(I)] = \underset{\text{sym}}{\tilde{\mathbb{E}}}[\text{val}(I)]$.

As a result, we may freely assume throughout that we are working with a shift-symmetric pseudo-expectation.

6.1.1 The Algorithm

We have now covered sufficient background to present the algorithm behind Theorem 6.1, Iterated Condition and Round. The algorithm follows the strategy presented in [21, Algorithm 6.1], differing mainly in that the parameter $r(\varepsilon)$ satisfying their second condition has been replaced with HD-Threshold Rank.

Condition and Round: We start with a basic sub-routine common to the SoS literature, which takes a pseudo-expectation $\tilde{\mathbb{E}}$ 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:

1. Sample a vertex $v \in V$ uniformly at random, and condition on event $X_{v,0} = 1$ to get the conditional pseudo-expectation $\tilde{\mathbb{E}}[\cdot | X_{v,0} = 1]$,

2. Sample $x \in \Sigma^V$ by independently sampling each coordinate x_u from the categorical distribution $\Pr[x_u = s] = \tilde{\mathbb{E}}[X_{u,s} | X_{v,0} = 1]$.

Following [21], we call the expected value of the solution output by Condition and Round the *CR-Value* of the instance I, denoted CR-Val(I). It is worth noting that Condition and Round, and thus the entire algorithm, can be de-randomized by standard techniques like the method of conditional expectations [21].

Iterated Condition and Round: The full algorithm builds a solution by iteratively applying Condition and Round to links. Let M be a k-dimensional HD-walk on $X = J(n,d)$, $I = (M,\Pi)$ 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. The following process returns an $\Omega_{\varepsilon,r,k}(1)$ satisfying assignment.

1. Let $\delta(\varepsilon) = \frac{\varepsilon}{r\binom{k}{r}}$. Solve the Degree-$D = \tilde{O}\left(1/\delta(\varepsilon)\right)$ SoS SDP relaxation of unique games, and symmetrize the resulting pseudo-expectation to get $\tilde{\mathbb{E}}_0$. Set $j = 1$.

2. Let $\text{Dif}(j) = \tilde{\mathbb{E}}_0[val_I(x)] - \tilde{\mathbb{E}}_{j-1}[val_I(x)]$. While $\text{Dif}(j) \leq \varepsilon$:

   (a) Find an r'-link X_τ for $r' \leq r$ such that the CR-Value of $I|_{X_\tau}$ is at least $\delta(\varepsilon + \text{Dif}(j))$.

   (b) Let S_j be the subgraph of X_τ induced by the vertices which have not yet been assigned a value in any partial assignment f_i, $i \leq j$, and perform Condition and Round on S_j to get partial assignment f_j.

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

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

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

6.2 Proving Theorem 6.1

The remainder of this section is devoted to proving Theorem 6.1. Our main technical contribution lies in showing how key structural properties of the Johnson graphs exploited by [21] are in fact inherent local-to-global properties of HD-walks. On the algorithmic side, we appeal directly to [21], employing the following lemma which shows how to reduce Iterated Condition and Round to Condition and Round on subgraphs with poor expansion.

Lemma 6.3 (Lemma 6.12 [21]). *Let $\varepsilon_0 \in (0,1)$ be a universal constant, $\varepsilon < \varepsilon_0/2$, $\delta : [0,1] \to [0,1]$ a function, and $\delta_{min} = \min_{\delta(\eta) \in [\varepsilon, 2\varepsilon]}$. Let G be a random walk*[11] *on any graph 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 given as input $\tilde{\mathbb{E}}$, a shift-symmetric degree-D pseudo-expectation satisfying $\mathcal{A}_I$ with $\tilde{\mathbb{E}}[val_I(x)] \geq 1-\eta \geq 1-\varepsilon_0$ 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 and Round[12] *outputs a solution for I satisfying an $\Omega(\delta_{min}^2 \varepsilon)$-fraction of edges of G in time $|V(G)|(T(\mathcal{A}) + |V(G)|^{O(D)})$.*

Our task is thus reduced to efficiently finding a subgraph with poor expansion on which Condition and Round has high expected value. BBKSS prove this fact for the Johnson graph in part by exploiting the two properties we discussed at the beginning of the section, that is:

1. There exists a low-degree SoS proof that non-expanding sets are concentrated in links.
2. There exists a parameter $r = r(\varepsilon)$ such that:
   (a) The $(r+1)$-st largest (distinct) eigenvalue is small:
   $$\lambda_r \leq 1 - \Omega(\varepsilon)$$
   (b) The expansion of any s-link with $s < r$ is small:
   $$\forall \tau \in X(s), s \leq r : \Phi(X_\tau) \leq O(\varepsilon).$$

In essence, we have already shown these properties hold for HD-walks in Section 5. All that remains is to show our methods fit into the SoS framework, and to strengthen the results for the special case of the complete complex. While it may be possible to extend our arguments beyond the complete complex (at least for some subset of HD-walks), it is not clear whether or not this more general case fits into the SoS framework. We discuss this issue in more detail after proving property 1 for the complete complex.

Proposition 6.4. *Let M be a k-dimensional HD-walk on $X = J(n,d)$, $n \geq \Omega(k^3)$, and $\{\lambda_i\}_{i=0}^k$ be eigenvalues corresponding to the HD-Level-Set Decomposition. Finally, let $r \leq \lceil k/2 \rceil$ and λ_b be parameters such that $\forall i \geq r+1$, $\lambda_i \leq \lambda_b$. Then the expansion*[13] *with respect to M of any function $f \in C_k$ which is not correlated with any s-link for $s \leq r$ is large:*

$$\vdash_2 \langle f, (I-M)f \rangle \geq (1-\lambda_b)\left(\mathbb{E}[f] - \frac{\binom{k}{r}\binom{n-r}{k-r}}{\binom{n-2r}{k-r}}\left(\sum_{i=0}^{r} \langle E_f^i, E_f^i \rangle_{X(i)}\right) + B(f)\right),$$

where $B(f) = \mathbb{E}[f^2 - f]$ measures the booleanity of f, and we recall for $\tau \in X(i)$, $E_f^i(\tau) = \mathop{\mathbb{E}}_{X_\tau}[f]$.

Proof. Let $f = f_0 + \ldots + f_k$ give the HD-Level-Set Decomposition of f. Recall from the proof of Theorem 1.11 that:

$$\langle E_f^i, E_f^i \rangle = \frac{\langle f_i, f_i \rangle^2}{\langle g_i, g_i \rangle} + \mathbb{E}\left[\left(E_f^i - \frac{\langle f_i, f_i \rangle}{\langle g_i, g_i \rangle} g_i\right)^2\right], \tag{17}$$

[11] BBKSS only state the result for regular graphs, but their proof only requires that the weight is spread evenly across vertices, a fact which holds for any random walk.

[12] Technically the algorithm is slightly more general, replacing links with any subgraph along with a number of parameters. See [21, Algorithm 6.13].

[13] Here we use a slightly different notion of expansion which will be useful later in the proof.

where E_f^i is the vector of densities of f in i-links, that is for $\tau \in X(i)$, $E_f^i(\tau) = \mathbb{E}_{X_\tau}[f]$. Further, it is well known [23] that for $J(n,d)$ the ratio of $\langle f_i, f_i \rangle$ to $\langle g_i, g_i \rangle$ is in fact a constant:

$$\frac{\langle f_i, f_i \rangle}{\langle g_i, g_i \rangle} = \frac{\binom{n-2i}{k-i}}{\binom{k}{i}\binom{n-i}{k-i}}.$$

Since E_f^i and g_i are both linear in the coefficients of f, Equation (17) then provides the following degree-2 SoS proof:

$$\vdash_2 \langle f_i, f_i \rangle \leq \frac{\binom{k}{i}\binom{n-i}{k-i}}{\binom{n-2i}{k-i}} \langle E_f^i, E_f^i \rangle.$$

In fact, it's worth noting that we can even replace $\langle E_f^i, E_f^i \rangle$ on the RHS with $\mathrm{Var}(E_f^i)$, tightening the bound. However, since the term does not provide much additional advantage in what follows, we drop it for simplicity. Expanding $\langle f, (I-M)f \rangle$, we get the desired bound:

$$\begin{aligned}
\langle f, (I-M)f \rangle &= \langle f, f \rangle - \sum_{i=0}^{k} \lambda_i \langle f_i, f_i \rangle \\
&= \langle f, f \rangle - \sum_{i=0}^{r} \lambda_i \langle f_i, f_i \rangle - \sum_{i=r+1}^{k} \lambda_i \langle f_i, f_i \rangle \\
&\leq \langle f, f \rangle - \sum_{i=0}^{r} \langle f_i, f_i \rangle - \lambda_b \sum_{i=r+1}^{k} \langle f_i, f_i \rangle \\
&= (1-\lambda_b) \left(\mathbb{E}[f] + B(f) - \sum_{i=0}^{r} \langle f_i, f_i \rangle \right) \\
&\leq (1-\lambda_b) \left(\mathbb{E}[f] + B(f) - \frac{\binom{k}{r}\binom{n-r}{k-r}}{\binom{n-2r}{k-r}} \sum_{i=0}^{r} \langle E_f^i, E_f^i \rangle \right)
\end{aligned}$$

where both inequalities are degree-2 SoS (in the coefficients of f), and the last relies further on our assumption that $n \geq \Omega(k^3)$. □

It is worth giving a brief qualitative comparison of this result to a similar version given in [21], who employ the old Fourier analytic techniques used originally by [20]. Proposition 6.4 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 drastically simplifies the proof (which takes up a 10-page appendix in [21]). It is also worth highlighting here that the key barrier to extending Theorem 6.1 to HD-walks beyond the complete complex lies in our use of the fact that $\langle f_i, f_i \rangle / \langle g_i, g_i \rangle$ is a constant. While this term is tightly bounded for two-sided local-spectral expanders, to our knowledge it is not necessarily a constant, and the inequality no longer fits into the SoS framework as a result. On the other hand, $\langle f_i, f_i \rangle / \langle g_i, g_i \rangle$ is always constant for objects called sequentially differential posets (objects introduced by Stanley [42] and discussed in the context of high dimensional expansion by [23]). We will discuss the extension of our results to such objects in our upcoming companion paper.

We now turn our attention to the second structural property necessary for BBKSS' analysis: the existence of a parameter $r(\varepsilon)$ such that $\lambda_r \leq 1 - \Omega(\varepsilon)$, while the expansion of s-links for $s < r$ is at most $O(\varepsilon)$. For strong enough two-sided local-spectral expanders, it is not hard to see that applying Theorem 5.6 implies not only that $r(\varepsilon)$ exists, but further that it is exactly the HD-Threshold Rank of the constraint graph. However, depending on the constraint graph, it is possible that the complete complex is not a sufficiently strong two-sided local-spectral expander for this to be the case, so we must modify our approach. In fact, it is possible to prove a stronger, exact version of Theorem 5.6 by recalling that the HD-Level-Set Decomposition is an exact eigendecomposition for $J(n,d)$ [23], which will give us a sufficiently strong result for any HD-walk.

Proposition 6.5. *Let M be a k-dimensional HD-walk over $X = J(n, d)$, $\{\lambda_i\}_{i=0}^k$ be eigenvalues corresponding to the HD-Level-Set Decomposition*[14]*, and r a parameter such that:*

$$\lambda_0 \geq \ldots \geq \lambda_r.$$

Then for all $0 \leq i \leq r$, the expansion of i-links is inversely related to λ_i:

$$\forall \tau \in X(i), \Phi(X_\tau) \leq 1 - \lambda_i.$$

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

$$\begin{aligned}
\Phi(X_\tau) &= 1 - \frac{1}{\alpha}\langle \mathbb{1}_{X_\tau}, M\mathbb{1}_{X_\tau}\rangle \\
&= 1 - \frac{1}{\alpha}\sum_{j=0}^{k}\langle \mathbb{1}_{X_\tau}, M\mathbb{1}_{X_\tau,j}\rangle \\
&= 1 - \frac{1}{\alpha}\sum_{j=0}^{k}\lambda_i\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau,j}\rangle
\end{aligned}$$

where α is the density of X_τ, and $\mathbb{1}_{X_\tau,j} \in V_j$. Notice that because X_τ is a link, it comes from level at most i:

$$\mathbb{1}_{X_\tau} = \binom{k}{i} U_i^k \mathbb{1}_\tau \in V_k^0 \oplus \ldots \oplus V_k^i.$$

Thus we may write:

$$\begin{aligned}
\Phi(X_\tau) &= 1 - \frac{1}{\alpha}\sum_{j=0}^{i}\lambda_j\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau,j}\rangle \\
&\leq 1 - \lambda_i\frac{1}{\alpha}\sum_{j=0}^{i}\langle \mathbb{1}_{X_\tau}, \mathbb{1}_{X_\tau,j}\rangle \\
&= 1 - \lambda_i
\end{aligned}$$

as desired. □

Then as long as the eigenvalues of the HD-Level-Set decrease, we see that both of our desired local-to-global properties are satisfied. Unfortunately, there is a slight technical issue: the eigenvalues of the HD-Level-Set Decomposition don't necessarily decrease monotonically.[15] For instance, the eigenvalues of the partial-swap walks, a basis for HD-walks over the complete complex, instead decrease until they become $O_k(1/n)$, at which point they begin to alternate in sign. This, however, is not much more than a technical annoyance since it is true that sufficiently *large* eigenvalues of the HD-Level-Set Decomposition decrease monotonically as long as the walk is not too lazy.

Lemma 6.6. *Let M be an HD-walk on $X = J(n, d)$, $n > 2^{\Omega(k)}$, where the probability of staying at a vertex is at most $1/50$. Let $\{\lambda_i\}_{i=0}^k$ be eigenvalues corresponding to the HD-Level-Set Decomposition, and let $r = R_{.68}(M)$. Then the following three properties hold:*

1. *$r \leq \lceil k/2 \rceil$*
2. *$\lambda_0 \geq \ldots \geq \lambda_r$*

[14]For $J(n, d)$, it is well known that each space in the HD-Level-Set Decomposition consists of eigenvectors with a single eigenvalue λ_i, see e.g. [23].

[15]While this may seem to contradict Proposition 4.10, the issue is actually caused by spaces with equivalent approximate eigenvalues.

3. $\forall i \geq r : \lambda_{r-1} \geq \lambda_i$*, and* $\lambda_i < 0.68$*.*

Proof. Any HD-walk on $J(n,d)$ may be expressed as a convex combination of partial-swap walks:

$$M = \sum_{i=0}^{k} \alpha_i S_k^i,$$

where we are further guaranteed that the lazy term $\alpha_0 \leq 1/50$. Let λ_j^i denote the eigenvalue of V_k^i with respect to S_k^i. Since the partial-swap walks share eigenspaces, we have:

$$\lambda_j = \sum_{i=0}^{k} \alpha_i \lambda_j^i.$$

Since $J(n,2k)$ is a two-sided $O_k(1/n)$-local-spectral expander, Corollary 4.8 implies that

$$\lambda_j^i = \frac{\binom{k-i}{j}}{\binom{k}{j}} \pm \frac{2^{O(k)}}{n}.$$

Using this fact, we start by proving the second and third properties (technically the third follows immediately from the second, but we prove them jointly). Note that since these properties imply that eigenvalues decrease until they are bounded from above by .68, we can then directly check the first property (that r is at most $\lceil k/2 \rceil$) by computing the worst-case value of $\lambda_{\lceil k/2 \rceil}$. To prove the latter properties, notice that it is sufficient to show the following for all j:

1. If $\lambda_j \geq 0.675$, $\forall i > j$, $\lambda_i < \lambda_j$
2. $\forall i > j, \lambda_i - \lambda_j < \frac{2^{O(k)}}{n} < .005$

If these conditions hold, the eigenvalues of M must decrease until they reach one of value less than .675, at which point the remaining values are upper bounded by .68.

The second of these facts is trivial, following directly from the form of the λ_j and letting n be sufficiently large. To prove the first, notice that if $\lambda_j \geq 0.675$, M must have non-trivial weight α_i for some $1 \leq i \leq k-j$. In particular, for large enough n we have that:

$$\exists 1 \leq i \leq k-j : \alpha_i > \Omega(1/k).$$

Further, for such an i, the j-th eigenvalue of S_k^i is substantially larger than all subsequent eigenvalues:

$$\forall t > j : \lambda_j^i - \lambda_t^i > 2^{-O(k)}.$$

We can now bound the difference of λ_j and subsequent eigenvalues. For all $t > j$ we have:

$$\begin{aligned}
\lambda_j - \lambda_t &= \sum_{m=0}^{k} \alpha_m \left(\lambda_j^m - \lambda_t^m\right) \\
&\geq 2^{-O(k)} + \sum_{m=0, m \neq i}^{k} \alpha_m \left(\lambda_j^m - \lambda_t^m\right) \\
&\geq 2^{-O(k)} - k\frac{2^{O(k)}}{n} \\
&\geq 0,
\end{aligned}$$

as desired. It is left to show that for any M satisfying $\alpha_0 \leq 1/50, \lambda_{\lceil k/2 \rceil} < .68$. Recall that $\lambda_{\lceil k/2 \rceil} = \sum \alpha_i \lambda_{\lceil k/2 \rceil}^i$. For large enough n, this is maximized by letting $\alpha_0 = 1/50$, and $\alpha_1 = 49/50$, so it is sufficient to check that $1/50 + 49/50\lambda_{\lceil k/2 \rceil}^1 < .68$. This can be checked directly by noting that $\lambda_{\lceil k/2 \rceil}^1 \leq \frac{k - \lceil k/2 \rceil}{k} + O_k(1/n)$. □

Since affine unique games are trivial[16] over walks with a constant lazy component, we can assume without loss of generality that the lazy component of all walks we consider is at most $1/50$. We are finally ready to show how the discussed properties can be leveraged to find a link with high CR-Value. The argument, which follows BBKSS [21, Lemma 6.9], centers around applying Proposition 6.4 to a potential function whose value lower bounds the CR-value of the instance. The exact form of the potential is not particularly important, but we need it to satisfy a few properties in order to successfully apply Proposition 6.4.

Theorem 6.7 (BBKSS [21] Sections 4,6)**.** *Let $\varepsilon \in (0, 1/17)$ and $\beta, \nu \in (0, 1)$ be parameters such that $\beta \geq 17\varepsilon$, and $\nu < \varepsilon$. Let $I = (G(V, E), \Pi)$ be an instance of affine unique games, and $\tilde{\mathbb{E}}$ a Degree-$D = \tilde{O}(1/\nu)$ shift-symmetric pseudo-expectation satisfying the corresponding constraints $\mathcal{A}_I$ with value at least $1 - \varepsilon$.*[17] *Finally, let H be a vertex-induced subgraph of G such that for all $v \in H$, the weight of edges leaving v is equivalent and at most 16ε. Then there exists a potential function $\Phi^{I,H}_{\beta,\nu}$ in the variables of $\tilde{\mathbb{E}}$ such that:*

1. *If $\tilde{\mathbb{E}}\left[\Phi^{I,H}_{\beta,\nu}\right] \geq \delta$, then the CR-value of I_H is at least $(\delta - 3\nu)(\beta - 16\varepsilon - \nu)$*
2. *There exists a family of degree $\leq D$ polynomials $\mathcal{F}_{\beta,\nu}$ independent of H such that:*

$$\Phi^{I,H}_{\beta,\nu} = \sum_{f \in \mathcal{F}_{\beta,\nu}} \mathop{\mathbb{E}}_{H}[f]^2$$

and the following properties hold:

   (a) The pseudo-expectation of $\sum \mathbb{E}[f]$ is large:

$$\tilde{\mathbb{E}}\left[\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathbb{E}[f]\right] \geq 1 - \frac{\varepsilon}{1 - \beta - \nu} - \nu$$

   (b) The pseudo-expectation of the expansion of $\sum f$ is small:

$$\tilde{\mathbb{E}}\left[\sum_{f \in \mathcal{F}_{\beta,\nu}} \langle f, (I - A_G) f\rangle\right] \leq 2\varepsilon + 2\frac{\varepsilon}{1 - \beta - \nu} + 2\nu$$

   (c) The pseudo-expectation of $\sum B(f)$ is small:

$$\tilde{\mathbb{E}}\left[\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathbb{E}[f - f^2]\right] \leq \frac{\varepsilon}{1 - \beta - \nu} + \nu$$

As written, Theorem 6.7 does not appear in [21], but follows directly from a number of claims and propositions throughout Sections 4 and 6 of their work. The key to finding a link with high CR-value is then to notice the connection between the form of the potential in Theorem 6.7 and Proposition 6.4's characterization of expansion.

Proposition 6.8. *Let M be a k-dimensional HD-walk on $X = J(n, d)$, $n \geq 2^{\Omega(k)}$, $\varepsilon \in [0, .02)$, $r = r(\varepsilon) = R_{1-16\varepsilon}(M) - 1$, and I be an affine unique games instance over M with value at least $1 - \varepsilon$. Then given a degree-$\tilde{O}\left(\frac{1}{\varepsilon} r \binom{k}{r}\right)$ pseudo-expectation satisfying the axioms $\mathcal{A}_{\mathcal{I}}$, we can find in time $n^{O(r)}$ an s-link X_τ for $0 \leq s \leq r$ with CR-Value $\Omega\left(\frac{\varepsilon}{r\binom{k}{r}}\right)$.*

[16] Self-edge constraints are either always satisfied or unsatisfiable, so for a lazy enough walk with high value any assignment to the nodes will satisfy sufficient constraints.

[17] Technically this pseudo-expectation is actually split over two "independent samples" of the variables corresponding to the instance I. Since this does not particularly matter for our arguments, we ignore this detail. See [21, Definition 2.2] for more information.

Proof. As long as the conditions of Theorem 6.7 are satisfied, it is sufficient to find a link X_τ with high potential. To find such a link, recall that the potential promised by Theorem 6.7 may be written as:

$$\Phi^{I,H}_{\beta,\nu} = \sum_{f \in \mathcal{F}} \mathop{\mathbb{E}}_{H}[f]^2$$

for some family of functions $\mathcal{F}$. This is closely related to the characterization of expansion in Proposition 6.4, where the term $\langle E^i_f, E^i_f \rangle$ may be re-written as:

$$\langle E^i_f, E^i_f \rangle = \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\mathop{\mathbb{E}}_{X_\tau} [f]^2 \right].$$

Using this connection, we can apply Proposition 6.4 across all functions in the family to lower bound the expected potential across small links in our complex. Using properties (a), (b), and (c) of Theorem 6.7 to lower bound this expectation, we complete the proof by an averaging argument.

More formally, let $\mathcal{F}_{\beta,\nu}$ be the family of functions corresponding to the potentials promised by Theorem 6.7. Notice that by our restriction on ε and Lemma 6.6, $r = R_{1-16\varepsilon}(M) - 1$ and $\lambda_b = 1 - 16\varepsilon$ satisfy the conditions of Proposition 6.4. Applying this to each $f \in \mathcal{F}_{\beta,\nu}$ yields:

$$\begin{aligned}
\vdash_2 \frac{1}{16\varepsilon} \sum_{f \in \mathcal{F}_{\beta,\nu}} \langle f, (I-M)f \rangle &\geq \sum_{f \in \mathcal{F}_{\beta,\nu}} \left(\mathbb{E}[f] - \frac{\binom{k}{r}\binom{n-r}{k-r}}{\binom{n-2r}{k-r}} \left(\sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\mathop{\mathbb{E}}_{X_\tau} [f]^2 \right] \right) + B(f) \right) \\
&= \left(\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathbb{E}[f] \right) - \frac{\binom{k}{r}\binom{n-r}{k-r}}{\binom{n-2r}{k-r}} \left(\sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathop{\mathbb{E}}_{X_\tau} [f]^2 \right] \right) + \sum_{f \in \mathcal{F}_{\beta,\nu}} B(f) \\
&= \left(\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathbb{E}[f] \right) - \frac{\binom{k}{r}\binom{n-r}{k-r}}{\binom{n-2r}{k-r}} \left(\sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\Phi^{I,X_\tau}_{\beta,\nu} \right] \right) + \sum_{f \in \mathcal{F}_{\beta,\nu}} B(f),
\end{aligned}$$

which in turn gives the aforementioned lower-bound on the potential of low-level links:

$$\vdash_2 \sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\Phi^{I,X_\tau}_{\beta,\nu} \right] \geq \frac{\binom{n-2r}{k-r}}{\binom{k}{r}\binom{n-r}{k-r}} \left(\left(\sum_{f \in \mathcal{F}_{\beta,\nu}} \mathbb{E}[f] \right) + \left(\sum_{f \in \mathcal{F}_{\beta,\nu}} B(f) \right) - \left(\sum_{f \in \mathcal{F}_{\beta,\nu}} \frac{\langle f, (I-M)f \rangle}{16\varepsilon} \right) \right). \quad (18)$$

Since r is the first index such that $\lambda_{r+1} < 1 - 16\varepsilon$, it follows from Lemma 6.6 and our restrictions on ε that $\lambda_r \geq 1 - 16\varepsilon$, and further by Proposition 6.5 that for any s-link X_τ, $0 \leq s \leq r$, $\Phi(X_\tau) \leq 16\varepsilon$. Since this additionally holds vertex-by-vertex by symmetry, all links in Equation (18) satisfy the requirements of Theorem 6.7, and taking the pseudo-expectation implies:

$$\sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\tilde{\mathbb{E}} \left[\Phi^{I,X_\tau}_{\beta,\nu} \right] \right] \geq \frac{\binom{n-2r}{k-r}}{\binom{k}{r}\binom{n-r}{k-r}} \left(1 - 2\frac{\varepsilon}{1-\beta-\nu} - 2\nu - \frac{1}{16\varepsilon} \left(2\varepsilon + 2\frac{\varepsilon}{1-\beta-\nu} + 2\nu \right) \right).$$

by setting $\beta = 17\varepsilon$ and $\nu = \frac{\varepsilon}{32r\binom{k}{r}}$ and recalling our assumptions on n and ε, this may be further simplified to:

$$\sum_{i=0}^{r} \mathop{\mathbb{E}}_{\tau \in X(i)} \left[\tilde{\mathbb{E}} \left[\Phi^{I,X_\tau}_{\beta,\nu} \right] \right] \geq \frac{1}{8\binom{k}{r}}.$$

As a result, we see by averaging that there must exist some s-link X_τ for $s \leq r$ with high potential:

$$\exists X_\tau, |\tau| \leq r : \tilde{\mathbb{E}} \left[\Phi^{I,X_\tau}_{\beta,\nu} \right] \geq \frac{1}{8(r+1)\binom{k}{r}}. \quad (19)$$

By Theorem 6.7, the CR-Value of $I|_{X_\tau}$ is then at least

$$\text{CR-Val}(I|_{X_\tau}) \geq \left(\frac{1}{8(r+1)\binom{k}{r}} - \frac{3\varepsilon}{32r\binom{k}{r}} \right) \left(\varepsilon - \frac{\varepsilon}{32r\binom{k}{r}} \right) \geq \Omega \left(\frac{\varepsilon}{r\binom{k}{r}} \right)$$

as desired. $\square$

The proof of Theorem 6.1 follows almost immediately from Lemma 6.3, Proposition 6.5, and Proposition 6.8.

Proof. We begin by setting the parameters for Lemma 6.3. Let the function δ be given by:

$$\delta(\eta) = \frac{1}{50r(\eta)\binom{k}{r(\eta)}}.$$

Then by Proposition 6.8 and Proposition 6.5, there exists a sub-routine which finds an s-link X_τ, $0 \leq s \leq r(\varepsilon)$, in time $n^{O(r(\varepsilon))}$ such that:

1. The CR-Value of I_{X_τ} is at least $\delta(\varepsilon)$.
2. The expansion of X_τ is small: $\Phi(X_\tau) < 16\varepsilon$.

Thus we are in position to apply Lemma 6.3, which completes the proof. □

7 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. Finally, we 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

[1] 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. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–12, 2019.

[2] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani. Approximating constraint satisfaction problems on high-dimensional expanders. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 180–201. IEEE, 2019.

[3] Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. *arXiv preprint arXiv:2001.02827*, 2020.

[4] 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.

[5] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of glauber dynamics up to uniqueness via contraction. *arXiv preprint arXiv:2004.09083*, 2020.

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

[7] Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1495–1524. IEEE, 2019.

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

[9] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta Shma. List decoding with double samplers. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2134–2153. SIAM, 2019.

[10] Yotam Dikstein, Irit Dinur, Prahladh Harsha, and Noga Ron-Zewi. Locally testable codes via high-dimensional expanders. *arXiv preprint arXiv:2005.01045*, 2020.

[11] Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 974–985. IEEE, 2017.

[12] Tali Kaufman and Izhar Oppenheim. High order random walks: Beyond spectral gap. *Combinatorica*, pages 1–37, 2020.

[13] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for ramanujan complexes and topological expanders. *Geometric and Functional Analysis*, 26(1):250–287, 2016.

[14] Izhar Oppenheim. Local spectral expansion approach to high dimensional expanders part i: Descent of spectral gaps. *Discrete & Computational Geometry*, 59(2):293–330, 2018.

[15] Khot Subhash, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 592–601. IEEE, 2018.

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

[17] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 376–389, 2018.

[18] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 940–951, 2018.

[19] 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.

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

[21] Mitali Bafna, Boaz Barak, Pravesh Kothari, Tselil Schramm, and David Steurer. Playing unique games on certified small-set expanders. *arXiv preprint arXiv:2006.09969*, 2020.

[22] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *2011 ieee 52nd annual symposium on foundations of computer science*, pages 472–481. IEEE, 2011.

[23] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha. Boolean function analysis on high-dimensional expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

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

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

[26] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2- ε. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.

[27] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007.

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

[29] Radha Chitta, Rong Jin, Timothy C Havens, and Anil K Jain. Approximate kernel k-means: Solution to large scale kernel clustering. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 895–903, 2011.

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

[31] Subhash A Khot and Nisheeth K Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into l1. *Journal of the ACM (JACM)*, 62(1):1–39, 2015.

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

[33] Vedat Alev. *Higher Order Random Walks, Local Spectral Expansion, and Applications.* PhD thesis, University of Waterloo, 2020.

[34] Sanjeev Arora, Subhash A Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K Vishnoi. Unique games on expanding constraint graphs are easy. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 21–28, 2008.

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

[36] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *Journal of the ACM (JACM)*, 62(5):1–25, 2015.

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

[38] Alexandra Kolla. Spectral algorithms for unique games. *computational complexity*, 20(2):177–206, 2011.

[39] Boaz Barak, Fernando GSL Brandao, Aram W Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 307–326, 2012.

[40] Irit Dinur, Yuval Filmus, Prahladh Harsha, and Madhur Tulsiani. Explicit sos lower bounds from high-dimensional expanders. *arXiv preprint arXiv:2009.05218*, 2020.

[41] Howard Garland. p-adic curvature and the cohomology of discrete subgroups of p-adic groups. *Annals of Mathematics*, pages 375–423, 1973.

[42] Richard P Stanley. Differential posets. *Journal of the American Mathematical Society*, 1(4):919–961, 1988.

A Proof of Lemma 4.3

In this section, we prove a strengthening of the main technical lemma of DDFH Section 8 [23, Claim 8.8], which allows for better control of error propagation.

Lemma A.1 (Strengthened Claim 8.8 [23]). *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 (15). For the inductive step, consider:

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

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

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

where $\|\Gamma_i\| \leq \gamma$. For the latter term, consider flipping DU and UD. By Equation (15) 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 3.6 as an immediate corollary. This improvement mainly matters in the regime where $\gamma \leq 2^{-ck}$ for c a small constant.

Proposition A.2. *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 4.3, and simplify based on applying the naive bounds on N_k^i given by Corollary 4.5:

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

where $\|h_i\| \le \gamma(k-\ell)(i+1)\,\|g_\ell\|$. We now apply Cauchy-Schwarz, and Lemma 3.6 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 3.6, 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 propogates.

B Orthogonality and the HD-Level-Set Decomposition

In this section we discuss in a bit more depth the error in [12, Theorem 5.10], and further show by direct counter-example that its implication [32] that the HD-Level-Set is orthogonal does not hold. In [12], Kaufman and Oppenheim analyze an approximate eidgendecomposition of the upper walk N_k^1 for two-sided local-spectral expanders. They prove a specialized version of Theorem 2.2 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 corresonding to strip W^i. Kaufman and Oppenheim claim in [12, 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 [32] 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 [12, Theorem 5.10] cannot hold by arguments of [32].

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:

$$\langle U_1^3\mathbb{1}_1, f\rangle \propto \sum_{(1xy)\in X} g(1x) + g(1y) + g(xy) \tag{20}$$

is non-zero. To do this, we first simplify the above expression assuming $g \in \mathrm{Ker}(D_2)$, 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:

$$\sum_{\substack{(1x)\in X(2):\\ x\notin[3]}} g(1x) = -\frac{n-3}{n-2}\left(g(12)+g(13)\right),$$

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

Plugging this into Equation (20), 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:

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

Thus it remains to show that there exists $g \in \mathrm{Ker}(D_2)$ 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 $D_2g(i) = 0$ 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 $D_2g(i)$. This can be checked through a direct but tedious computation that we leave to the reader. $\square$

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il