



Hypercontractivity on High Dimensional Expanders: a Local-to-Global Approach for Higher Moments

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

November 24, 2021

Abstract

Hypercontractivity is one of the most powerful tools in Boolean function analysis. Originally studied over the discrete hypercube, recent years have seen increasing interest in extensions to settings like the p -biased cube, slice, or Grassmannian, where variants of hypercontractivity have found a number of breakthrough applications including the resolution of Khot's 2-2 Games Conjecture (Khot, Minzer, Safra FOCS 2018). In this work, we develop a new theory of hypercontractivity on high dimensional expanders (HDX), an important class of expanding complexes that has recently seen similarly impressive applications in both coding theory and approximate sampling. Our results lead to a new understanding of the structure of Boolean functions on HDX, including a tight analog of the KKL Theorem and a new characterization of non-expanding sets.

Unlike previous settings satisfying hypercontractivity, HDX can be asymmetric, sparse, and very far from products, which makes the application of traditional proof techniques challenging. We handle these barriers with the introduction of two new tools of independent interest: a new explicit combinatorial Fourier basis for HDX that behaves well under restriction, and a new local-to-global method for analyzing higher moments. Interestingly, unlike analogous second moment methods that apply equally across all types of expanding complexes, our tools rely inherently on simplicial structure. This suggests a new distinction among high dimensional expanders based upon their behavior beyond the second moment.

*Department of Computer Science, Harvard University, MA 02138. Email: mitalibafna@g.harvard.edu.

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

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

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

1 Introduction

Introduced over 50 years ago today, *hypercontractivity* remains one of the most powerful tools in the analysis of boolean functions. Originally used to prove numerous landmark results on the discrete hypercube such as the KKL Theorem [KKL88] and Majority is Stablest [MOO05], the study of hypercontractivity has since seen a resurgence on extended domains such as the p -biased cube [KLLM19], slice [KMMS18], and Grassmannian [KMS18]. Fascinatingly, these regimes all share a common thread: while hypercontractivity doesn't hold in general, it is satisfied for certain classes of *pseudorandom functions*. This recently discovered phenomenon has led to a slew of breakthroughs, most famously including the resolution of Khot's 2-2 Games Conjecture [KMS18]. Unfortunately, the scope of these results is currently restricted, as all known proof techniques rely on product structure or other strong symmetries, and no unifying theory is known to exist.

In this work we take the first substantive step towards solving this issue with the introduction of a new theory of hypercontractivity for the general class of *high dimensional expanders* (HDX). HDX are a family of expanding complexes that have seen an explosion of work in recent years, leading to major breakthroughs across a number of areas including (among others) the recent construction of c3-LTCs [DEL⁺21], and efficient approximate sampling for many important systems (e.g. for matroid bases [ALOV19], independent sets [ALO20], Ising models [AJK⁺21], and more). Our results lead to a new understanding of the structure of boolean functions on HDX, including a tight analog of the KKL Theorem, and a characterization of non-expanding sets similar to that used in the proof of 2-2 Games [KMS18]. Proving such results previously seemed out of reach since HDX are very far from products, asymmetric, and can be quite sparse. To handle these challenges, we introduce a new set of tools including a new explicit Fourier decomposition and a local-to-global method for analyzing higher order moments. Interestingly, unlike previous ℓ_2 -based techniques which apply equally across all types of expanding complexes, our methods rely crucially on the underlying HDX structure being *simplicial*. This suggests a new stratification of spectral HDX based upon their behavior *beyond the second moment*.

1.1 Contributions

Before jumping into a more detailed breakdown of our results, we start by giving an informal overview of our main contributions within the broader context of classical Fourier analysis and the theory of high dimensional expanders.

Classical Fourier Analysis: Classical Fourier Analysis on the discrete hypercube focuses on analyzing functions $f : \{0, 1\}^n \rightarrow \mathbb{R}$ through their *Fourier Expansion*, a decomposition that breaks f into a series of orthogonal “level functions,” each corresponding to the projection of f onto a certain eigenspace of the (noisy) hypercube graph.¹ At a basic level, a function's Fourier decomposition gives a nice method for understanding its *second moment*, since orthogonality allows one to move between this and the standard basis freely (a result usually known as *Parseval's Theorem*). On the other hand, in computer science, we are usually interested in analyzing the special class of *boolean functions* $f : \{0, 1\}^n \rightarrow \{0, 1\}$. These functions exhibit rich structure that Parseval's Theorem isn't equipped to capture—to understand them, we usually need to look beyond the second moment.

Hypercontractivity, introduced in 1970 by Bonami [Bon70] (and later independently by Beckner [Bec75] and Gross [Gro75]), is exactly the tool for the job. In its simplest form, hypercontractivity boils down to the statement that the fourth moment of low levels of the Fourier decomposition should behave nicely. Namely that the i th level of a boolean function f , denoted f_i , should satisfy:

$$\|f_i\|_4 \leq 2^{O(i)} \|f_i\|_2. \tag{1}$$

¹More generally, these are the eigenspaces of the Hamming scheme.

This deceptively simple observation, known in the above form as “Bonami’s Lemma” [Bon70], led to many landmark results including the KKL Theorem [KKL88], noise-sensitivity of sparse functions [KKL88], Friedgut’s Junta Theorem [Fri98], and Majority is Stablest [MOO05]. What’s more, hypercontractivity (and its resulting applications) actually extend beyond the hypercube. After KKL’s seminal work, many authors studied extensions and applications of hypercontractivity [BKK⁺92, Tal94, FK96, Fri98], but it wasn’t until recently that tight analogs of Equation (1) were developed for general product spaces [KLLM19] as well as for other structured domains such as the symmetric group [FKLM20] and Grassmannian [KMS18]. These extended domains differ from the hypercube in that they are only hypercontractive for special classes of pseudorandom functions, but are nevertheless responsible for an impressive set of applications including analogs of classical results, a variety of new sharp threshold theorems [KLLM19, LM19, KLLM21], and perhaps most famously the proof of the 2-2 Games Conjecture [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMMS18, KMS18]. Unfortunately, despite the stark similarities between these settings, no unified theory explaining the phenomenon exists. Further, all known techniques rely heavily on product structure or other strong forms of symmetry, which makes it difficult to approach the problem in more general settings.

Fourier Analysis on HDX: High dimensional expanders (HDX) are a class of robustly connected complexes that have seen an incredible amount of development and application throughout theoretical computer science in the past few years, most famously in coding theory [DEL⁺21, EKZ20, JST21, KO21, KT21b, DDHRZ20, JQST20, DHK⁺19] and approximate sampling [ALOV19, AL20, ALO20, CLV20, CLV21, CGŠV21, FGYZ21, JPV21, Liu21, BCC⁺21], but also in agreement testing [DK17, DD19, KM20], CSP-approximation [AJT19, BHKL20], and (implicitly) hardness of approximation [KMMS18, KMS18]. In this work, we study a central notion of high dimensional expansion called *two-sided local-spectral expansion*, originally developed by Dinur and Kaufman [DK17] to build sparse agreement testers. For simplicity, we’ll often refer to these objects just as *local-spectral expanders*, but the reader should be aware we always refer to the two-sided variant, not the weaker one-sided variant commonly used in approximate sampling.

Interestingly, local-spectral expanders are actually known to admit a (nascent) theory of Fourier analysis [DDFH18, KO20]. Initial works in this area have focused on the development and application of Fourier Decompositions and Parseval’s Theorem, and while the existing theory does have a few interesting applications (e.g. an FKN theorem for HDX [FKN02, DDFH18], efficient CSP-approximation [AJT19, BHKL20]), it is subject to the same limitations as original second moment methods on the hypercube: they simply don’t capture the richer structure of boolean functions. Let’s consider a concrete and important example: the expansion of pseudorandom sets (an analog of “sparse functions are noise-sensitive” on the hypercube).² Traditionally proved via hypercontractivity, a variant of this result on the Grassmannian recently led to the resolution of the 2-2 Games Conjecture [KMS18]. On the other hand, Bafna, Hopkins, Kaufman, and Lovett [BHKL20] showed that second moment methods cannot recover such a result. While they are able to recover some sort of characterization with these techniques, it necessarily decays as the dimension grows to infinity, becoming trivial in the regime useful for hardness of approximation—if we want to do better, it appears we need a theory of hypercontractivity.

This is easier said than done: local-spectral expanders look nothing like any object previously known to satisfy hypercontractivity. They can be sparse, asymmetric, and very far from products. Moreover, there are no known techniques for analyzing local-spectral expanders beyond the second moment.³ Even DDFH and KO’s Fourier decompositions are intrinsically tied to second moment methods, since they are defined by linear algebraic manipulation of the standard inner product. Surprisingly, it turns out that these barriers are not inherent, and can be removed with the introduction of just two new tools: a *combinatorial* Fourier

²The connection lies in the fact that the noise-sensitivity result can equivalently be phrased as saying that small sets on the noisy hypercube are expanding.

³We note that recent works in the sampling literature have considered entropic notions of high dimensional expansion, but the underlying assumptions are much stronger than local-spectral expansion.

decomposition for HDX, and a new local-to-global method to replace reliance on product structure in the analysis of higher moments.

Our new decomposition is the natural analog of the standard Fourier decomposition on product spaces (often called the “orthogonal” or “Efron-Stein” decomposition). It is equivalent to old decompositions in an ℓ_2 -sense (and therefore shares all relevant ℓ_2 -based properties), but comes with a number of additional benefits: it has simple explicit and recursive forms, and it behaves nicely under restriction. This allows us to bring to bear much of the power of more traditional Fourier-analytic machinery, which often relies on these same properties. Historically, however, applying this machinery in a useful fashion has also required the underlying object to be a product, or to satisfy some other strong symmetry. Our second key observation is that while individual variables in a local-spectral expander may be highly correlated, they look independent *on average*. More concretely, this means that in the analysis of *expectations* (such as a higher moment), we are free to treat the underlying variables as independent even if they actually exhibit a very high level of correlation.

Hypercontractivity on HDX: Leveraging these tools, we build a theory of hypercontractivity on HDX. Concretely, we prove that Equation (1) holds on local-spectral expanders for an appropriate notion of pseudorandom functions—ones that are not concentrated in any local restrictions on the complex.⁴ Combined with BHKL’s recent spectral analysis of *higher order random walks* (which, for the moment, we’ll think of as analogs of the noisy hypercube graphs or Hamming scheme), this leads to the resolution of a number of open questions in boolean function analysis. To start, we provide a tight characterization of (edge) expansion on higher order random walks, which, unlike previous methods [BHKL20], *does not decay with dimension*. This matches the version of the result on the Grassmannian which led to the resolution of the 2-2 Games Conjecture [KMS18], and opens yet another avenue towards the use of HDX in hardness of approximation. We also introduce natural analogs⁵ of two classic Fourier-analytic notions: *influence* and the *noise operator*. Combining these with the above recovers tight variants of both the KKL Theorem and noise-sensitivity of sparse (or in this case pseudorandom) functions.

Beyond these concrete applications, hypercontractivity on HDX also has interesting implications in the broader context of discrete Fourier analysis and high dimensional expansion. For the former, our result gives the first general class of hypercontractive objects beyond products, and combined with bounded degree constructions [LSV05, KO18], the first example of hypercontractivity over any sparse object at all.⁶ For the latter, our result suggests a new stratification among notions of local-spectral expansion. This requires some additional explanation. While local-spectral expanders were originally introduced only over simplicial complexes, they were quickly extended to more general settings such as the Grassmannian, or even to general ranked posets [DDFH18]. While these classes of local-spectral expanders are essentially equivalent in an ℓ_2 sense [DDFH18, BHKL20, KT21a], our analysis of the fourth moment crucially relies on simplicial structure. We conjecture that this is an inherent rather than technical barrier: only special classes of underlying objects (e.g. Grassmannian, simplicial complexes) satisfy hypercontractivity, and thereby lead to the strongest known form of spectral high dimensional expanders.

⁴In the high dimensional expansion literature, these restrictions are known as *links*.

⁵When applied to the embedding of the hypercube into a simplicial complex, these definitions return the standard notions.

⁶Formally, the notion of “sparseness” here requires a bit of explanation. Previous settings of hypercontractivity all have natural representations as dense complexes. For instance, the hypercube complex has $2n$ vertices and 2^n top-level faces, and even weakly hypercontractive sparsifications like the short code [BGH⁺12] have quasi-polynomially many faces.

2 Background

Before stating our results more formally, we give a quick overview of the theory of local-spectral expanders and higher order random walks. Local-spectral expansion is a robust notion of connectivity on weighted hypergraphs introduced by Dinur and Kaufman [DK17] in the context of agreement testing. As is standard in the area, we will view d -uniform hypergraphs $H \subseteq \binom{[n]}{d}$ as (pure) **simplicial complexes**:

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

where $X(d) = H$, $X(i) \subseteq \binom{[n]}{i}$ is given by downward closure, and $X(0) = \emptyset$. We note that this notation is off by one from much of the HDX literature which considers $X(i) \subseteq \binom{[n]}{i+1}$. This notation is standard in the topological literature (where an i -simplex indeed as $i + 1$ points), but is less natural for our purely combinatorial work.

Most recent work on high dimensional expansion is based on the *local-to-global paradigm*, in which local properties of a complex are lifted to a desired global property (e.g. mixing or agreement testing). The main local structure of interest are called **links**. For every “ i -face” $\tau \in X(i)$, the link of τ is the subcomplex obtained by restriction to faces including τ :

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

A simplicial complex is said to be a γ -**local-spectral expander** if (the graph underlying) every link is a γ -spectral expander.⁷

Higher order random walks are an analog of the standard walk on expander graphs that moves between two vertices via an edge. Kaufman and Mass [KM16] observed that this process can be applied at any level of a simplicial complex: one could move between edges via a triangle, or triangles via a pyramid. Formally, these walks are defined as a composition of **averaging operators**, objects that have become ubiquitous tools in the study of high dimensional expanders. Denote the space of functions $\{f : X(k) \rightarrow \mathbb{R}\}$ as C_k . For a function $f \in C_k$, the (level k) **Up** and **Down operators** lift and lower f to level $k + 1$ and $k - 1$ respectively by averaging:

$$U_k f(\tau) = \mathbb{E}_{\sigma \subset \tau} [f(\sigma)],$$

$$D_k f(\tau) = \mathbb{E}_{\sigma \supset \tau} [f(\sigma)].$$

It will often be useful to compose the down or up operators multiple times to move between levels k and i , we denote this by $D_i^k = D_i \circ \dots \circ D_k$ and $U_i^k = U_k \circ \dots \circ U_i$. Informally, HD-walks are simply affine combinations of composed averaging operators. For instance, the basic composition $N_k^i = U_k^{k+i} D_k^{k+i}$, called a **canonical walk**, is the random process which moves between two k -faces via a shared $(k + i)$ -face.

3 Results

We now move to an informal description of our results. We view our work as having three main contributions. First, we introduce and develop a new theory of Fourier analysis on high dimensional expanders. This includes a new explicit Fourier decomposition, as well as a number of natural generalizations of Fourier-analytic ideas such as influence and the noise operator to simplicial complexes. Second, we prove that our Fourier-analytic decomposition satisfies a hypercontractive inequality for the special subclass of *pseudorandom functions*, and use this fact to characterize the small set expansion of HD-walks and give a version of Bourgain’s Theorem (an analog of KKL on product spaces) on HDX. Finally, en route to our hypercontractivity theorem, we introduce a new method of localization on high dimensional expanders of independent interest that enables local-to-global analysis of higher order moments.

⁷A graph is a γ -spectral expander if the second largest eigenvalue of its adjacency matrix is at most γ in absolute value.

3.1 The Bottom-Up Decomposition

We start with a discussion of our new explicit Fourier-analytic decomposition. All previously known Fourier bases on local-spectral expanders [KO20, DDFH18] are linear algebraic in nature, and have no known closed form. While these decompositions certainly have their place and are sufficient for a number of interesting applications [DDFH18, AJT19, BHKL20], they often fall short when finer-grained calculation is required. To alleviate this issue, we introduce a new combinatorial decomposition on simplicial complexes that is an analog of the classic orthogonal (sometimes called Efron-Stein) decomposition on product spaces. For simplicity, we'll start by introducing the decomposition in a simple recursive form, and give the explicit form later in this section.

Definition 3.1 (Bottom-Up (Recursive Form)). *Let X be a d -dimensional pure simplicial complex and $f \in C_k$ any function. Recursively define the i th level function(s) to be:*

$$g_{\uparrow i} = D_i^k f - \sum_{j=0}^{i-1} \binom{i}{j} U_j^i g_{\uparrow j}, \quad f_{\uparrow i} = \binom{k}{i} U_i^k g_{\uparrow i}$$

We call $f = \sum_{i=0}^k f_{\uparrow i}$ the *Bottom-Up Decomposition*.

Strictly speaking, while the consideration of this basis is new over general simplicial complexes, it was first studied for the special case of the complete complex by [KMMS18]. There, the authors took advantage of the complex's near-product structure to show that the decomposition gives an (approximate) Fourier basis close to the eigendecomposition of f with respect to the well-studied Johnson graphs. We prove that the assumption of near-product structure is actually unnecessarily strong—it is enough for the underlying complex to be sufficiently expanding.

Theorem 3.2 (Bottom-Up Properties (Informal Lemma 7.6+Theorem 7.8)). *Let X be a two-sided γ -local-spectral expander, and M an HD-walk. Then for any $f \in C_k$, and $0 \leq i < j \leq k$:*

1. $\langle f_{\uparrow i}, f_{\uparrow j} \rangle \approx 0$
2. $\|f\|_2^2 \approx \sum_{i=0}^k \|f_{\uparrow i}\|_2^2$
3. $\exists \lambda_i$ s.t. $M f_{\uparrow i} \approx \lambda_i f_{\uparrow i}$

Theorem 3.2 is also very similar to an analogous result for the HD-Level-Set Decomposition in [DDFH18, Theorem 1.3]. We will cover their definition in greater detail in Section 7. For the moment, it suffices to note that their decomposition also breaks f into $k + 1$ Fourier levels, which we similarly denote by $f = \sum_{i=0}^k f_{\downarrow i}$. It turns out that the similarities between the HD-Level-Set and Bottom-Up Decompositions are no accident—the two decompositions are actually nearly equivalent.

Theorem 3.3 (Bottom-Up Approximates HD-Level-Set (Theorem 7.8)). *Let X be a two-sided γ -local-spectral expander and $f \in C_k$. Then the Bottom-Up and HD-Level-Set Decomposition are close in ℓ_2 -norm:*

$$\|f_{\uparrow i} - f_{\downarrow i}\|_2^2 \leq 2^{O(k)} \gamma \|f\|_2^2.$$

Similarly,

$$|\langle f_{\uparrow i}, f_{\uparrow i} \rangle - \langle f_{\downarrow i}, f_{\downarrow i} \rangle| \leq 2^{O(k)} \gamma \|f\|_2^2.$$

The main advantage of the Bottom-Up Decomposition then lies in its simple recursive structure and explicit form, which we now describe.

Proposition 3.4 (Bottom-Up Explicit Form (Theorem 7.2)). *Let X be a d -dimensional pure simplicial complex, $f \in C_k$ any function, and $g_{\uparrow i}$ be as given in Definition 3.1. Then:*

$$g_{\uparrow i} = \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} U_j^i D_j^k f,$$

or equivalently $\forall \tau \in X(i)$:

$$g_{\uparrow i}(\tau) = \sum_{\sigma \subseteq \tau} (-1)^{|\tau \setminus \sigma|} \mathbb{E}_{X_\sigma} [f]$$

In Section 7, we will see how these properties are useful for analyzing finer-grained structures like restriction that are often key to classical Fourier-analytic arguments. It is unknown how to analyze such properties for prior linear algebraic decompositions, and determining whether the latter share similar structure at this level remains an interesting open problem.

3.2 Hypercontractivity

Now that we have introduced our relevant Fourier-analytic decomposition, we turn our attention to the study of hypercontractivity. Hypercontractivity is one of the most powerful tools in boolean function analysis and is crucial to proving many of area's key results (e.g. KKL [KKL88], FKN [FKN02], Majority is Stablest [MOO05], sharp threshold theorems [FB99], etc.). Informally, hypercontractivity can be thought of as a niceness condition on “low-degree” functions. We’ll start by considering a simple variant often called the Bonami or Bonami-Beckner lemma, [Bon70] which states that a “degree- i ” function p should satisfy:

$$\|p\|_4 \leq 2^{O(i)} \|p\|_2.$$

Classically, we might think of p as being a degree- i polynomial, corresponding to the i th Fourier level of a boolean function. The corresponding statement in our context is therefore that the i th level of the Bottom-Up Decomposition should satisfy an analogous inequality:

$$\|f_{\uparrow i}\|_4 \leq 2^{O(i)} \|f_{\uparrow i}\|_2. \quad (2)$$

Unfortunately, it is well known that Equation (2) cannot hold in our setting, even over the complete complex. However, it is possible that the inequality could hold for *natural subclasses* of functions. Indeed, such a phenomenon is known to occur on general product distributions [KLLM19], where **pseudorandom** functions satisfy a form of Equation (2).

Definition 3.5 (Pseudorandomness). *Let X be a simplicial complex and $f \in C_k$. We say f is (ε, i) -pseudorandom if it is sparse in every i -link in the following two senses:*

1. For all $\tau \in X(i)$:

$$\left| \mathbb{E}_{X_\tau} [f] \right| \leq \varepsilon \|f\|_\infty$$

2. For all $\tau \in X(i)$:

$$\langle f|_\tau, f|_\tau \rangle \leq \varepsilon \|f\|_\infty^2$$

In applications, we will often only care about non-negative functions, in which case the second condition can be removed completely (as it is implied by the first). We note that functions satisfying Definition 3.5 are also sometimes called *global* since they are not concentrated in any local structure [KLLM19, LM19]. We call them pseudorandom in keeping with prior literature on the Johnson and Grassmann graphs [KMMS18, KMS18], and because they cannot be distinguished from an $(\varepsilon$ -sparse) random function by examining density inside links. Finally, note that Definition 3.5 requires f to be sparse. We conjecture that our results should hold in the dense regime as well, and discuss this further in Section 4.

Hypercontractivity for restricted subclasses is still a very powerful tool. KLLM’s result, for instance, led to the resolution of Majority is Stablest in the p -biased setting [LM19], and the resolution of several conjectures in extremal combinatorics as well [KLLM21]. Unfortunately, their technique cannot be used to analyze simplicial complexes. Unlike the complete complex (which is essentially a product space), bounded degree HDX are necessarily quite far from products. Nevertheless, we prove that a variant of Equation (2) continues to hold.

Theorem 3.6 (Hypercontractivity on HDX (Informal Theorem 8.1)). *Let X be a sufficiently strong two-sided γ -local-spectral expander and $f \in C_k$ an (ε, i) -pseudorandom function. Then the following hypercontractive inequality holds:*

$$\mathbb{E}[f_{\uparrow i}^4] \leq 2^{O(i)} \varepsilon \mathbb{E}[f_{\uparrow i}^2] \|f\|_{\infty}^2 + c_k \gamma^{1/2} \varepsilon \|f\|_2^2 \|f\|_{\infty}^2,$$

where $c_k \leq \min\{2^{O(k)}, k^{O(i)}\}$.

Our overall framework for proving Theorem 3.6 roughly follows KMMS’ strategy for the complete complex. However, even with analogous results for the Bottom-Up Decomposition in hand, most of their techniques fail dramatically in our setting due to local-spectral expanders’ distinct lack of product structure. In fact, Theorem 3.6 gives the first general class of hypercontractive objects beyond product spaces, and combined with known bounded degree constructions of local-spectral expanders [LSV05, KO18], the first example over any *sparse* domain at all. In Section 3.4, we’ll discuss how we tackle these traditionally hard-to-handle structures with the introduction of a new notion of average-case independence that relates closely to local-spectral expansion. Our method actually allows for analysis well beyond the 4th moment, and can also be used to extend Theorem 3.6 to 2-to- $2q$ hypercontractivity (where the 4-norm is replaced by a higher $2q$ -norm). We focus on the 2-to-4 case in this work for simplicity.

Before moving on to applications of Theorem 3.6, it is worth discussing another typical form of hypercontractivity and how it translates to the setting of simplicial complexes. Hypercontractivity is frequently expressed in terms of an object called the **noise operator**. On the hypercube, the noise operator T_{ρ} acts as an averaging process on boolean strings which replaces each coordinate with a random bit with probability $1 - \rho$. In this context, hypercontractivity states that T_{ρ} should act as a *smoothing operator* in the following sense:

$$\|T_{\rho}f\|_4 \leq \|f\|_2 \tag{3}$$

for some constant ρ . While it is not immediately obvious how to define T_{ρ} over a simplicial complex, one can imagine a similar process where each vertex in a k -face is removed with probability $1 - \rho$, and is then re-randomized over relevant k -faces. We formalize this procedure in terms of the averaging operators.

Definition 3.7 (Noise Operator). *Let X be a d -dimensional pure simplicial complex. The noise operator $T_{\rho}^k(X) : C_k \rightarrow C_k$ at level $k \leq d$ of the complex is:*

$$T_{\rho}^k(X) = \sum_{i=0}^k \binom{k}{i} (1 - \rho)^i \rho^{k-i} U_{k-i}^k D_{k-i}^k.$$

We write just T_{ρ} when clear from context.

When applied to the hypercube complex,⁸ this natural analog returns exactly the standard boolean noise operator T_ρ . Combining standard arguments with the spectral properties of the Bottom-Up Decomposition, we can also prove a variant of Equation (3) for pseudorandom functions on HDX. To state this result, it will be useful to have a notion of degree: as on the hypercube, we say the degree of a function f is the largest i such that $f_{\uparrow i}$ is non-zero.

Corollary 3.8 (Informal Proposition 10.15). *Let X be a sufficiently strong two-sided γ -local-spectral expander and $f \in C_k$ a degree i , (δ, i) -pseudorandom function for $\delta \leq \varepsilon \left\| \frac{f}{\|f\|_\infty} \right\|_2^2$. Then for some constant $\rho = \Theta(1)$:*

$$\|T_\rho f\|_4 \leq \varepsilon^{1/4} \|f\|_2.$$

3.3 Applications

A classical application of hypercontractivity is to give what is known as a “level- i inequality” that bounds low-level weight of a boolean function. We can use Theorem 3.6 to give an analog on HDX for pseudorandom functions.

Theorem 3.9 (Level- i inequality (Informal Theorem 9.4)). *Let X be a two-sided γ -local-spectral expander with γ sufficiently small and $f \in C_k$ an (ε, i) -pseudorandom boolean function of density α . Then the weight on $f_{\uparrow i}$ is bounded by:*

$$\langle f_{\uparrow i}, f_{\uparrow i} \rangle \leq 2^{O(i)} \varepsilon^{1/3} \alpha.$$

Level- i inequalities have a plethora of applications in boolean Fourier analysis. We’ll look at the analog of two classical applications: one to small-set expansion, and the other to the structure of functions with low influence. Starting with the former, let’s recall the basic definition of edge-expansion.

Definition 3.10. *Let M be a walk on the k th level of a simplicial complex X . The (edge) expansion of a subset $S \subseteq X(k)$ is the average probability of leaving S in a single step of the walk:*

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

where $M(v, X(k) \setminus S)$ is the probability the walk leaves S starting from v .

Informally, a walk is called a *small-set expander* if all small subsets expand. Traditionally, the level- i inequality on the discrete hypercube is used to show that the noisy hypercube graph is a small-set expander. The analogous result on simplicial complexes, however, isn’t true: HD-walks (which generalize graphs like the noisy hypercube) have well-known examples of small non-expanding sets: *links* [KMMS18, BHKL20]. Using Theorem 3.9, we can prove a converse to this result: *any non-expanding set must be concentrated in a link.*

Theorem 3.11 (Characterizing non-expansion on HD-walks (Informal Theorem 9.3)). *For every $0 < \delta < 1$, there exists some $\varepsilon > 0$ and $r \in \mathbb{N}$ such that for all large enough k the following holds. For any HD-walk⁹ on a sufficiently strong two-sided local-spectral expander X and any subset $S \subseteq X(k)$, if S has expansion at most $\Phi(S) \leq \delta$, then S is concentrated in a low-level link:*

$$\exists i \leq r, s \in X(i) : \frac{|X_s \cap S|}{|X_s|} \geq \varepsilon$$

⁸The hypercube complex has vertex set $[n] \times \{0, 1\}$, where the first entry stands for a coordinate and the second entry a value. The top level $X(n)$ consists of all binary strings and is exactly the hypercube.

⁹Formally, this statement only holds for HD-walks such as $N_k^{\Theta(k)}$ which exhibit sufficiently fast eigenvalue decay. We give a more general formulation in the main body that holds for all HD-walks (see Theorem 9.3).

Expansion is also closely related to a well-studied Fourier-analytic quantity called **total influence**. On the boolean hypercube, the total influence of a function measures its total variability across each coordinate:

$$I[f] = \sum_{i=1}^n \Pr_{x \sim \{0,1\}^n} [f(x) \neq f(x \oplus e_i)]$$

where e_i is the i th standard basis vector. One of the most celebrated results in the analysis of boolean functions is the KKL Theorem [KKL88], which states that any function with low total influence must have an influential coordinate. In domains beyond the hypercube (such as product spaces), total influence is usually instead written equivalently as:

$$I[f] = \langle f, Lf \rangle$$

where L is the (un-normalized) **Laplacian operator** (see Section 10 for more details). While the KKL Theorem does not hold over product spaces (even say the p -biased hypercube), a useful analog known as “Bourgain’s Sharp Threshold Theorem” [FB99, Appendix] does. Bourgain’s Theorem states that if a boolean function has small total influence, there must exist a link (on the hypercube a subcube) in which the function is much denser than expected.

We prove an analogous result for HDX. The Laplacian formulation of total influence has a natural generalization on simplicial complexes:

$$I_X[f] = \langle f, k(I - U_{k-1}D_k)f \rangle$$

that returns the standard definition over the hypercube complex (see Section 10). Using Theorem 3.9, we prove that any function with low total influence must be concentrated in a link.

Theorem 3.12 (Bourgain’s Theorem for HDX (Informal Theorem 10.5)). *Let X be a sufficiently strong two-sided γ -local-spectral expander, and $f \in C_k$ a boolean function. Then for any $0 \leq K \leq k$, if $I[f] \leq K\text{Var}(f)$, there exists $i \leq K$ and an i -face τ such that the link of τ is dense:*

$$\mathbb{E}_{X_\tau}[f] \geq 2^{-O(K)}.$$

Note that Theorem 10.5 is actually a bit weaker than Bourgain’s Theorem in the sense that it only promises a link that is much denser than average when the function f is sparse. We conjecture that this result should hold in the dense regime as well (see Section 4 for details). On the other hand, unlike Bourgain’s Theorem (which has a density increase of $2^{-O(K^2)}$ rather than $2^{-O(K)}$ for general functions), our result is tight.¹⁰

Proposition 3.13 (Bourgain’s Theorem Lower Bound (Informal Proposition 10.6)). *Let $c \geq 1$ be any constant and $K > 1$ an integer. For all $K \ll k \ll n$, there exists a Boolean function $f \in C_k$ on the k -dimensional complete complex on n vertices satisfying:*

1. *The influence of f is small:*

$$I[f] \leq K\text{Var}(f).$$

2. *For every $i \leq cK$, all i -links are sparse:*

$$\forall i \leq cK, \tau \in X(i) : \mathbb{E}_{X_\tau}[f] \leq 2^{-\Omega(K)}.$$

¹⁰A similar tight version of Bourgain’s Theorem for sparse functions on the p -biased cube was proved by [KLLM19].

3.4 Localization (Average Independence)

Our hypercontractive inequality is derived from a new method of localization on high dimensional expanders of independent interest. Localization itself is of course not new—indeed such techniques have recently become synonymous with HDX. However, most prior work in the literature focuses on the localization of *second moments*, whereas hypercontractivity requires the analysis of *higher moments*. Traditionally, analysis beyond the second moment is difficult on HDX due to an inherent lack of product structure. We show that this can often be circumvented by a new method of decorrelating variables.

Theorem 3.14. *Let X be a d -dimensional two-sided γ -local-spectral expander and $f \in C_k$. Then for any $j \leq d - k$ and $\tau \in X(j)$, the global and localized expectation of f over X_τ differ by an operator with small spectral norm:*

$$\mathbb{E}_{X_\tau(k)}[f] - \mathbb{E}_{X(k)}[f] = \Gamma f(\tau)$$

where $\Gamma : C_k \rightarrow C_j$ satisfies $\|\Gamma\| \leq O_{k,j}(\gamma)$.

We emphasize that the first expectation in this definition is given by *localizing* rather than *restricting* f . In other words we are averaging over k -faces in the link X_τ (which are $(k + j)$ -faces in the original complex) rather than over k -faces in the original complex X that contain τ . This latter notion of restriction is also very important in analysis of HDX. We discuss further in Section 6.

Theorem 3.14 should really be thought of as saying that, on average, f can be decorrelated from “irrelevant” j -faces that don’t appear in the input. This is particularly useful when analyzing objects like HDX with high correlation. To understand the technique a bit more concretely, let’s look at a basic example application.

Let X be a γ -local-spectral expander. We will often be interested in analyzing certain expected products on X . For instructive purposes, let’s take a look at an example of such a product with just two instances of some $g \in C_2$:

$$\mathbb{E}_{a \sim X(1)} \mathbb{E}_{b \sim X_a(1)} \mathbb{E}_{c \sim X_{ab}(1)} [g(a, b)g(a, c)] = \mathbb{E}_{a \sim X(1)} \mathbb{E}_{b \sim X_a(1)} [g(a, b) \mathbb{E}_{c \sim X_{ab}(1)} [g(a, c)]] . \quad (4)$$

Notice that if we were working over a product space, the distribution of $c \sim X_{ab}(1)$ would be the same as the distribution of $c \sim X_a(1)$. This allows us to significantly simplify the above:

$$\begin{aligned} \mathbb{E}_{a \sim X(1)} \mathbb{E}_{b \sim X_a(1)} [g(a, b) \mathbb{E}_{c \sim X_{ab}(1)} [g(a, c)]] &= \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b)] \mathbb{E}_{c \sim X_a(1)} [g(a, c)] \right] \\ &= \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b)]^2 \right] . \end{aligned}$$

On the other hand in an HDX (especially one of bounded degree), this could be far from true since b and c can be highly correlated. Theorem 3.14 provides a simple technique for circumventing this issue. Let $g|_a$ be the restriction of g to a , that is $g|_a(b) = g(a, b)$. Theorem 3.14 promises that

$$\mathbb{E}_{c \sim X_{ab}(1)} [g(a, c)] = \mathbb{E}_{c \sim X_a(1)} [g(a, c)] + \Gamma g|_a(b),$$

where $\|\Gamma\| \leq O(\gamma)$. This allows us to recover the same form as above up to $O(\gamma)$ error:

$$\begin{aligned} \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b)] \mathbb{E}_{c \sim X_{ab}(1)} [g(a, c)] \right] &= \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b)]^2 \right] + \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b) \cdot \Gamma g(b)] \right] \\ &\leq \mathbb{E}_{a \sim X(1)} \left[\mathbb{E}_{b \sim X_a(1)} [g(a, b)]^2 \right] + O_g(\gamma), \end{aligned}$$

where we have ignored some terms in g for simplicity and the last step follows from an application of Cauchy-Schwarz and the spectral norm (see Section 8 for details).

We emphasize that while Equation (4) in particular could also have been analyzed through a more direct application of the swap walk, *these standard techniques fail miserably when additional copies of g are added*. Since there will generally be j copies of g in analysis of the j th moment, this means that traditional techniques cannot go beyond the second moment. On the other hand, our technique is applied individually to each copy of g , so it is essentially irrelevant how many times it appears in the product.

4 Discussion

Before getting into the details and formalization of the above, we take a moment to give a more careful treatment of some interesting open problems and related work.

4.1 Open Problems

Hypercontractivity, both on the cube and on extended domains, has led to an astounding number of applications since its introduction some 50 years ago. We recover just a small sample of these classical applications in our work, and believe the theory will give rise to further results in the analysis of boolean functions. However, rather than surveying a list of classical results one might wish to extend (we refer the reader to O’Donnell’s excellent book [O’D14] for this), we’ll instead focus on three open problems we feel are most directly raised by our work.

Perhaps the most obvious direction left open is to extend hypercontractivity to the *dense regime*. While our definition of pseudorandomness implicitly assumes the underlying function is sparse, we conjecture that all of our results should hold under a weaker notion of pseudorandomness that drops this assumption.

Definition 4.1 (Pseudorandomness (Dense Regime)). *Let X be a simplicial complex and $f \in C_k$ a boolean function. We say f is (ε, i) -pseudorandom if its local and global average are close on every i -link:*

$$\forall \tau \in X(i) : \left| \mathbb{E}_{X_\tau}[f] - \mathbb{E}[f] \right| \leq \varepsilon.$$

While the stronger notion we use in this work is certainly sufficient for some applications (e.g. characterizing expansion, noise-sensitivity) and is line with previous work [KMMS18, KMS18, KLLM19], it does seem to fall short in other areas. A good example of this is our variant of Bourgain’s Theorem. While our version only promises the existence of a dense link, the original result on product spaces actually promises a link with *higher than average density* (albeit by a factor of $2^{-O(K^2)}$ instead of $2^{-O(K)}$), which could be recovered by proving hypercontractivity for the above definition. More generally, proving hypercontractivity for this dense variant opens the door to a broader spectrum of applications than the sparse regime alone can handle.

The second problem we’d like to discuss is more focused on the theory of high dimensional expanders itself. As mentioned in the introduction, local-spectral expansion can be extended well beyond simplicial complexes to many natural poset structures including the Grassmann poset [DDFH18, KT21a], where hypercontractivity was crucial to resolving the 2-2 Games Conjecture [KMS18]. The spectral and ℓ_2 -structure of these expanding posets (eposets) is well understood [DDFH18, AJT19, BHKL20], and essentially has no dependence on the underlying poset structure.¹¹ In stark contrast, our results break down over general eposets at several key points. In fact, it seems likely that the Bottom-Up Decomposition is not even a Fourier basis (fails to satisfy Theorem 3.2) over general eposets, since the proof relies heavily on simplicial structure (see Lemma 7.5). On the other hand, variants of hypercontractivity are known for some special eposets

¹¹Different poset parameters result in different eigenvalues, but the structure is otherwise the same.

such as the Grassmann poset. The key difference in these cases is that the definition of pseudorandomness necessarily changes. This raises a natural question: do all eposets satisfy hypercontractivity for some notion of pseudorandomness, or are structures like the Grassman poset and simplicial complexes “special”? We conjecture that the latter is the case, and that these objects represent a new, stronger class of spectral high dimensional expanders.

Our third proposed problem is not raised quite as directly by this work, but is hard to ignore in light of recent breakthroughs in approximate sampling via HDX [ALOV19, AL20, ALO20, CLV20, CLV21, CGŠV21, FGYZ21, JPV21, Liu21, BCC⁺21]. Hypercontractivity is classically connected to the *Log-Sobolev inequality*, which gives strong control over the mixing time of its associated random walk. Applied to the hypercube, for instance, this connection improves the standard spectral mixing bound from $O(n^2)$ to the optimal $\Theta(n \log(n))$ [DSC96]. Recent analysis of entropic notions of high dimensional expansion and a *modified Log-Sobolev inequality* have led to a slew of analogous improvements on important sampling problems [CLV21, BCC⁺21, AJK⁺21]. These results, however, usually only apply to dense objects and need stronger assumptions. Given these connections, it is natural to ask whether our theory of hypercontractivity can improve mixing times for general local-spectral expanders in some analogous fashion.

4.2 Related Work

Hypercontractivity on Extended Domains: Nearly 20 years after its introduction, Kahn, Kalai, and Linial [KKL88] revolutionized the study of boolean functions with hypercontractivity. Not long after, a significant interest grew in the development and application of hypercontractivity beyond the hypercube, with a particular focus on product distributions and especially the p -biased hypercube [BKK⁺92, Tal94, FK96, Fri98]. These works offered a general theory of hypercontractivity for such domains, but their strength depended on the underlying distributions in the product space. While this was sufficient for many important applications including a number of sharp threshold theorems [FK96, FB99], the results become meaningless for unbalanced products such as the p -biased cube for $p \leq o_n(1)$. This issue was resolved only recently by Keevash, Lifshitz, Long, and Minzer [KLLM19], who showed a hypercontractivity theorem for pseudorandom functions on product spaces that is independent of the underlying distributions. This result offered the missing piece for a number of classical applications, including a tight variant of the KKL Theorem (for monotone functions) [KLLM19], Majority is Stablest [LM19], and had a number of other applications in extremal combinatorics [KLLM21].

Another line of work has examined hypercontractivity on what are often called “exotic” domains: specific objects beyond products such as the slice [KMMS18], multislice [FOW18], Grassmannian [KMS18] (or similarly the degree-two short code [BKS18]), and symmetric group [FKLM20]. Like KLLM’s improved result for product distributions, most of these examples are only hypercontractive for pseudorandom functions (with the multislice being the only exception).¹² The main application of this line of work has been to agreement testing and hardness of approximation. In particular, hypercontractivity for the Grassmannian was used to prove the soundness of an agreement tester in the “1% regime” needed for the proof of the 2-2 Games Conjecture [KMS17, DKK⁺18b, DKK⁺18a, BKS18, KMMS18, KMS18]. It is worth noting that agreement testing theorems are also known for local-spectral expanders [DK17, DD19, KM20] (indeed the objects were originally introduced in this context). These results, however, lie in the “99% regime,” so it is interesting to ask whether our theory of hypercontractivity can be used to build a bounded degree agreement tester in the more difficult 1% regime.

Finally, we should note that our overarching proof structure for hypercontractivity builds on KMMS’ work on the slice (i.e. the complete complex). Their techniques, however, rely heavily on the fact that the slice is close in ℓ_1 -distance to a product. This is far from true on local-spectral expanders, especially those of

¹²We note that higher degrees of the short code are also hypercontractive, but only on low Fourier levels for general functions [BGH⁺12].

bounded degree which may essentially be as far as possible from products. As previously discussed, this lack of structure is a challenging barrier broken for the first time in this work.

Fourier Analysis on HDX: Fourier analysis on HDX was originally studied by Diksein, Dinur, Filmus, and Harsha [DDFH18], who introduced the HD-Level-Set Decomposition, analyzed its spectral properties, and used it to prove an FKN Theorem for HDX. A similar decomposition was also proposed around the same time by Kaufman and Oppenheim [KO20], though their work was more focused on understanding the spectral structure of higher order random walks than on developing a theory of Fourier analysis. In the years since, the HD-Level-Set Decomposition has seen some further development [AJT19, KS20, BHKL20], and the nascent theory has helped build efficient approximation algorithms for certain k -CSPs [AJT19] and unique games [BHKL20], but the restriction to second moment methods seems to have limited its use otherwise. Towards breaking this same barrier, Gur, Lifshitz, and Liu [GLL21] have also (independently) developed a similar theory of hypercontractivity on local-spectral expanders. While their work certainly shares some connections to ours, its main proof techniques differ substantially and we believe the two works are of independent interest.

4.3 Roadmap

Having concluded introductory discussion of our work, we lay out a brief roadmap for the rest of the paper. In Section 5 we give preliminaries and formally define local-spectral expansion and higher order random walks. In Section 6 we discuss our new local-to-global method for higher moments that allows us to move beyond product distributions. In Section 7 we discuss our new explicit Fourier Decomposition, its basic properties, and behavior under restriction. In Section 8 we prove hypercontractivity for pseudorandom functions (Theorem 3.6). In Section 9 we apply this result to characterize edge expansion in HD-walks (Theorem 3.11). Finally in Section 10 we introduce analogs of classic Fourier analytic notions such as influence and the noise operator and use them to prove both a KKL Theorem (Theorem 3.12) and noise-sensitivity of pseudorandom functions.

5 Preliminaries

Before moving into proofs and further discussion of our main results, we take a moment to cover the theory of local-spectral expanders and higher order random walks in more detail.

5.1 Simplicial Complexes

Our main objects of interest in this work are a family of expanding hypergraphs known as **local-spectral expanders**. In this context, it will be useful to think of d -uniform hypergraphs as objects called **pure simplicial complexes**.

Definition 5.1 (Weighted, Pure Simplicial Complex). *A d -dimensional, pure simplicial complex $X = X(0) \cup \dots \cup X(d)$ on n vertices is the downward closure of a hypergraph $X(d) = \binom{[n]}{d}$ where*

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

We call the elements of $X(i)$ i -faces. A weighted pure simplicial complex (X, Π) is a simplicial complex X endowed with a distribution Π over $X(d)$. This induces a distribution over each $X(i)$ by downward closure:

$$\pi_i(x) = \frac{1}{i+1} \sum_{y \in X(i+1): y \supset x} \pi_{i+1}(y), \quad (5)$$

where $\pi_d = \Pi$.

Weighted pure simplicial complexes are equivalent to weighted hypergraphs, and we will adopt the former viewpoint throughout the rest of this work. We note that our definition of dimension is off by one from some of the literature which adopts the convention that an i -face has $i + 1$ vertices. While this is natural from a topological viewpoint, it makes less sense in our combinatorial context.

Weighted simplicial complexes also come equipped with a natural set of inner products. Recall that $C_i = C_i(X)$ denotes the space of functions $f : X(i) \rightarrow \mathbb{R}$. The distribution $\Pi = (\pi_d, \dots, \pi_0)$ induces a natural inner product on each level:

$$\forall f, g \in C_i : \quad \langle f, g \rangle_{X(i)} = \mathbb{E}_{\tau \sim \pi_i} [f(\tau)g(\tau)].$$

When clear from context, we drop $X(i)$ from the notation. Just like on the hypercube, these associated products are a core component of function analysis and the development of Fourier analysis on HDX.

5.2 Local Spectral Expansion

In this work we focus on a recent spectral notion of high dimensional expansion called *two-sided local-spectral expansion* introduced by Dinur and Kaufman [DK17]. The definition hinges crucially on a form of local structure in simplicial complexes called **links**.

Definition 5.2 (Link). *Let (X, Π) be a d -dimensional weighted, pure simplicial complex. The **link** of an i -face $s \in X(i)$ is a $(d - i)$ -dimensional pure simplicial complex given by the restriction of X to faces containing s , that is:*

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

We call X_s an ***i*-link**. Throughout the rest of the paper, X_s will always refer to its weighted version (X_s, Π_s) where Π_s is induced by the original distribution Π by normalizing over top level faces of X_s .

When analyzing a particular level k of the complex, we will often abuse notation and write X_s to mean the set of k -faces in X which contain s when clear from context.

Much of the high dimensional expansion literature centers around what is called the **local-to-global paradigm**, where properties on links are lifted to global properties on a complex. Local-spectral expansion can be seen as a definitional formalization of this notion: a complex is said to be expanding if all its local parts are expanding.

Definition 5.3 (Local-spectral expansion [DK17]). *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 graph¹³ of X_s is a two-sided γ -spectral expander.¹⁴*

5.3 Higher Order Random Walks

Just like expander graphs are inextricably tied to their underlying random walks, local-spectral expanders are similarly connected to an analogous set of random processes known as higher order random walks (HD-walks). In Section 2, we discussed one example of these objects called the canonical walks that move between k -faces via a shared $(k + i)$ -face, and saw that these could be defined by the averaging operators. We'll now extend these definitions to the more general setting of weighted simplicial complexes and, as well as define HD-walks in full generality. We'll start with the weighted averaging operators.

¹³The underlying graph of a complex X is $G = (V = X(1), E = X(2))$.

¹⁴A weighted graph is a two-sided γ -spectral expander if $\max\{|\lambda_2|, |\lambda_n|\} \leq \gamma$.

Definition 5.4 (Averaging Operators). *Let (X, Π) be a d -dimensional weighted, pure simplicial complex. For every $0 \leq k < d$, the **Up Operator** U_k lifts functions from C_k to C_{k+1} by averaging:*

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

*Similarly, the **Down Operator** lowers functions from C_{k+1} to C_k by averaging:*

$$\forall \tau \in X(k) : D_{k+1} f(\tau) = \frac{1}{\pi_{k+1}(X_\tau)} \sum_{\sigma \in X_\tau} \pi_{k+1}(\sigma) f(\sigma),$$

where $\pi_{k+1}(X_\tau) = \sum_{\sigma \in X_\tau} \pi_{k+1}(\sigma)$, and the sum is over $k+1$ faces of X containing τ .

It is worth noting that the averaging operators are *adjoint* with respect to the associated inner products mentioned in the previous section, that is for any $f \in X(i)$ and $g \in X(i-1)$:

$$\langle f, U_i g \rangle_{X(i)} = \langle D_i f, g \rangle_{X(i-1)} = \mathbb{E}_{(\sigma, \tau) \sim (\pi_i, \pi_{i-1})} [f(\sigma)g(\tau)].$$

This means that basic combinations of the operators such as the canonical walks discussed in Section 2 are *self-adjoint* and therefore have a spectral decomposition.

Let's now formalize the notion of higher order random walks. We'll start with a basic version called **pure walks** that are simply a composition of the averaging operators.

Definition 5.5 (Pure Walk [AJT19]). *Given a weighted, pure simplicial complex (X, Π) , a k -dimensional pure walk $Y : C_k \rightarrow C_k$ on (X, Π) of height $h(Y)$ is a composition:*

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

where each Z_i is a copy of D or U .

For the moment we won't force these walks to be self adjoint, but as we noted basic examples such as $N_k^i = D_k^{k+i} U_k^{k+i}$ do satisfy this constraint.

We define general higher order random walks to be any linear combinations of pure walks which is stochastic and self-adjoint.

Definition 5.6 (HD-walk [AJT19]). *Let (X, Π) be a pure, weighted simplicial complex, and \mathcal{Y} a family of pure walks $Y : C_k \rightarrow C_k$ on (X, Π) . We call a linear combination*

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

a k -dimensional HD-walk on (X, Π) as long as it is stochastic and self-adjoint. We call $w(M) := \sum |\alpha_Y|$ the weight of M , and $h(M) = \max\{h(Y)\}$ its height.

5.4 Rectangular Swap Walks

Definition 5.6 only captures walks which stay on some fixed level of the complex. While these are certainly our main object of study, it turns out that in analysis it is often useful to consider *rectangular walks* which move between levels of the complex. We will be particularly interested in a rectangular walk introduced independently by Alev, Jeronimo, and Tusiani [AJT19], and Dikstein and Dinur [DD19] called the *swap walk*. Informally, the swap walk from $X(i)$ to $X(j)$ moves from an i -face τ to a j -face σ through a shared $(i+j)$ -face, but *swaps out all original elements in τ* . In other words, the intersection between τ and σ must be empty, and the shared $(i+j)$ -face is exactly $\tau \cup \sigma$. To formalize this, it is useful to first introduce the more basic rectangular walk moving between $X(i)$ and $X(j)$ with no such restrictions.

Definition 5.7 (Rectangular Canonical Walks). *Let (X, Π) be a d -dimensional, pure, weighted simplicial complex. For any $i + j \leq d$, the rectangular canonical walk $N_{i,j}$ is the natural operator moving between $X(i)$ and $X(j)$ through $X(i + j)$:*

$$N_{i,j} = D_i^{i+j} U_j^{i+j}.$$

Swap walks are then defined by forcing the down steps in a canonical walk to remove only vertices in the initial face.

Definition 5.8 (Rectangular Swap Walks). *Let (X, Π) be a d -dimensional, pure, weighted simplicial complex. For any $i + j \leq d$, the rectangular swap walk $S_{i,j}$ is the (normalized) restriction of $N_{i,j}$ to pairs $(\tau, \sigma) \in X(i) \times X(j)$ such that $|\tau \cap \sigma| = 0$.*

Swap walks appear naturally in a number of areas, including agreement testing [DD19], coding theory [JST21], and approximation algorithms [AJT19, BHKL20] and were well studied even before their formal introduction on HDX. On the complete complex, for instance, rectangular swap walks are exactly the bipartite Kneser graphs. Swap walks are particularly useful in these contexts because unlike their canonical counterpart canonical walks, they are actually great expanders.

Theorem 5.9 (Theorem 7.1 [DD19]). *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander. Then for any $i + j \leq d$, the spectral expansion of $S_{i,j}$ is at most:*

$$\lambda(S_{i,j}) \leq ij\gamma,$$

where $\lambda(S_{i,j})$ is the second largest singular value of $S_{i,j}$.

We note that this result was concurrently proved by AJT [AJT19], albeit with a quantitatively worse bound.

6 Localization Beyond the Second Moment

Localization is one of the (if not the) most important technique in the analysis of high dimensional expanders. Classic results, often grouped together under the name *Garland's method*, show how global functions on simplicial complexes can be broken down into an average over local parts. There are two forms of Garland's method that will be relevant to our work. The first handles *restrictions* of a function f in C_k to any $\tau \in X(i)$, s.t. $f|_\tau \in C_{k-i}(X_\tau)$ satisfies:

$$\forall \sigma \in X_\tau(k-i) : f|_\tau(\sigma) = f(\tau \cup \sigma).$$

Lemma 6.1 (Garland's method (restrictions) [KO20]). *Let (X, Π) be a weighted, pure simplicial complex, and $f \in C_k$. Then for any $i \leq k$, $\|f\|^2$ is equal to its average second moment restricted to i -links:*

$$\langle f, f \rangle = \mathbb{E}_{\tau \in X(i)} [\langle f|_\tau, f|_\tau \rangle].$$

The second form of interest handles *localizations* of f to $\tau \in X(i)$, where $f_\tau \in C_k(X_\tau)$ lifts f from $X(k)$ to $X_\tau(k)$:

$$\forall \sigma \in X_\tau(k) : f_\tau(\sigma) = f(\sigma).$$

Lemma 6.2 (Garland's method (localizations) [Opp18]). *Let (X, Π) be a d -dimensional, weighted, pure simplicial complex, and $f \in C_k$. Then for any $k + i \leq d$, $\|f\|_2^2$ is equal to its average second moment localized to i -links:*

$$\langle f, f \rangle = \mathbb{E}_{\tau \in X(i)} [\langle f_\tau, f_\tau \rangle].$$

Garland's method will play an important role in the analysis of Theorem 3.6, but the results are generally only useful once a problem has been reduced to analyzing second moments. Since we are mainly interested in hypercontractivity and analyzing higher moments, Garland's method alone won't be sufficient.

To this end, we introduce a new technique for analyzing higher moments on two-sided local-spectral expanders. At its core, the strategy relies on a deceptively simple observation: the difference between the global expectation of f and its localized expectation over links is exactly given by an application of the swap walk minus its stationary operator.

Lemma 6.3. *Let (X, Π) be a d -dimensional pure, weighted simplicial complex and $f \in C_i$. Then for any $v \in X(j)$ such that $i + j \leq d$, we have:*

$$\mathbb{E}_{X_v} [f_v] - \mathbb{E}[f] = (S_{j,i} - U_0^j D_0^i) f(v).$$

Proof. This is essentially immediate from expanding the left-hand side. We have:

$$\begin{aligned} \mathbb{E}_{X_v} [f_v] - \mathbb{E}[f] &= \sum_{w \in X(i)} \pi_{v,i}(w) f(w) - \sum_{w \in X(i)} \pi_i(w) f(w) \\ &= \sum_{w \in X(i)} (\pi_{v,i}(w) - \pi_i(w)) f(w) \\ &= (S_{j,i} - U_0^j D_0^i) f(v), \end{aligned}$$

where $\pi_{v,i}(w) = 0$ for any $w \notin X_v(i)$. □

Lemma 6.3 is particularly powerful on two-sided local-spectral expanders, since the spectral norm of $\|S_{j,i} - U_0^j D_0^i\|$ is small on every link by Theorem 5.9.

Corollary 6.4. *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander, $f \in C_i$, and $\tau \in X(\ell)$ for any $\ell < i$. Then for any $v \in X_\tau(j)$ such that $i + j - \ell \leq d$, the global and localized expectations of f differ by:*

$$\mathbb{E}_{X_\tau \cup v} [f|_\tau] - \mathbb{E}_{X_\tau} [f|_\tau] = \Gamma f|_\tau(v)$$

where $\Gamma : C_{i-\ell}(X_\tau) \rightarrow C_j(X_\tau)$ is an operator with spectral norm at most $\|\Gamma\| \leq (i - \ell)j\gamma$.

Proof. Applying Lemma 6.3 to the restricted function $f|_\tau$ on X_τ (which is also a two-sided γ -local-spectral expander), we have that the left-hand side is exactly given by $(S_{j,i-\ell} - U_0^j D_0^{i-\ell})_\tau f|_\tau(v)$. Since $U_0^j D_0^{i-\ell}$ is the stationary operator of $S_{j,i-\ell}$, the spectral norm $\|S_{j,i-\ell} - U_0^j D_0^{i-\ell}\|$ is exactly the second largest singular value of $S_{j,i-\ell}$. As discussed in Theorem 5.9, Dikstein and Dinur [DD19, Theorem 7.1] proved that this quantity is at most $(i - \ell)j\gamma$ on any two-sided γ -local-spectral expander. □

7 The Bottom-Up Decomposition

In this section, we introduce the Bottom-Up Decomposition, an explicit combinatorial decomposition on simplicial complexes which approximates Dikstein, Dinur, Filmus, and Harsha's HD-Level-Set Decomposition [DDFH18]. This is particularly useful since the latter decomposition essentially corresponds to the eigenspaces of HD-walks [DDFH18, AJT19, BHKL20].

Definition 7.1 (Level Functions (Recursive Form)). *Let (X, Π) be a d -dimensional pure simplicial complex and $f \in C_k$ any function. The i th level function of the Bottom-Up Decomposition is given by*

$$g_{\uparrow i} = D_i^k f - \sum_{j=0}^{i-1} \binom{i}{j} U_j^i g_{\uparrow j}$$

The Bottom-Up Decomposition is given by lifting the level functions via the up operator.

Theorem 7.2 (Bottom-Up Decomposition (Explicit Form)). *Let (X, Π) be a d -dimensional pure simplicial complex and $f \in C_k$ any function. Let $f_{\uparrow i} = \binom{k}{i} U_i^k g_{\uparrow i}$ be the lift of the i th level function to C_k . Then the following statements hold:*

1. *The lifted level functions give a decomposition of f :*

$$f = \sum_{i=0}^k f_{\uparrow i}$$

2. *The lifted level functions have the following explicit form:*

$$f_{\uparrow i} = \binom{k}{i} \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} U_k^i D_j^k f,$$

or equivalently for all $\tau \in X(i)$:

$$g_{\uparrow i}(\tau) = \sum_{\sigma \subseteq \tau} (-1)^{|\tau \setminus \sigma|} \mathbb{E}_{X_\sigma} [f], \quad f_{\uparrow i}(\tau) = \sum_{\sigma \in X(i): \sigma \subset \tau} g_{\uparrow i}(\sigma)$$

Proof. (1) can be proved directly by the explicit form given in (2):

$$\begin{aligned} \sum_{i=0}^k f_{\uparrow i} &= \sum_{i=0}^k \binom{k}{i} U_i^k \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} U_j^i D_j^k f \\ &= \sum_{j=0}^k \left(\sum_{i=j}^k (-1)^{i-j} \binom{k}{i} \binom{i}{j} \right) U_j^k D_j^k f \\ &= U_k^k D_k^k f \\ &= f. \end{aligned}$$

where we've used the fact that:

$$\sum_{i=j}^k (-1)^{i-j} \binom{k}{i} \binom{i}{j} = \delta_{jk}.$$

It is left to prove (2). We proceed by induction. Note that the equality clearly holds for $i = 0$, where both sides are simply the global expectation $\mathbb{E}[f]$. Now assume by induction that the equivalence holds up to $i - 1$.

We then have:

$$\begin{aligned}
g_{\uparrow i} &= D_i^k f - \sum_{j=0}^{i-1} \binom{i}{j} U_j^i g_{\uparrow j} \\
&= D_i^k f - \sum_{j=0}^{i-1} \binom{i}{j} U_j^i \sum_{\ell=0}^j (-1)^{j-\ell} \binom{j}{\ell} U_\ell^j D_\ell^k f \\
&= D_i^k f - \sum_{\ell=0}^{i-1} \left(\sum_{j=\ell}^{i-1} (-1)^{j-\ell} \binom{i}{j} \binom{j}{\ell} \right) U_\ell^i D_\ell^k f \\
&= D_i^k f - \sum_{\ell=0}^{i-1} (-1)^{i-1-\ell} \binom{i}{\ell} U_\ell^i D_\ell^k f \\
&= \sum_{\ell=0}^i (-1)^{i-\ell} \binom{i}{\ell} U_\ell^i D_\ell^k f.
\end{aligned}$$

The explicit form of $f_{\uparrow i}$ then follows simply by applying $\binom{k}{i} U_i^k$ to $g_{\uparrow i}$, and the equivalent form is immediate from the definition of the down and up operators. \square

7.1 Bottom-Up vs. HD-Level-Set

Dikstein, Dinur, Filmus, and Harsha's HD-Level-Set Decomposition [DDFH18] is an elegant linear-algebraic decomposition for functions on local-spectral expanders. Like the Bottom-Up Decomposition, it breaks $f \in C_k$ down into $k + 1$ Fourier levels, but differs in that it does so in a *top-down* fashion.

Theorem 7.3 (HD-Level-Set Decomposition, Theorem 8.2 [DDFH18]). *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander, $\gamma < \frac{1}{d}$, $0 \leq k \leq d$, and let:*

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

Then:

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

In other words, every $f \in C_k$ has a unique decomposition $f = \sum f_{\downarrow i}$ such that $f_{\downarrow i} = U_i^k g_{\downarrow i}$ for $g_{\downarrow i} \in \text{Ker}(D_i)$.

While the HD-Level-Set Decomposition is certainly useful in its own right, it has no known explicit form. This can make the analysis of standard Fourier-analytic techniques like restriction difficult, and hampers the analysis of higher moments. We will show that the Bottom-Up Decomposition provides an explicit approximation of the HD-Level-Set Decomposition that circumvents these issues while maintaining the latter's useful properties.

Before jumping into the details, however, it is worth reviewing an elegant technical tool of [DDFH18] that will be crucial for our analysis. They prove that local-spectral expansion is equivalent to a global notion of spectral expansion on complexes that relates the upper and lower walks.

Theorem 7.4 (DDFH Claim 8.8). *Let (X, Π) be a d -dimensional two-sided γ -local-spectral expander. Then for any $1 \leq i \leq j \leq d$:*

$$D_i U_j^i = \frac{j}{i} U_{j-1}^{i-1} D_j + \frac{i-j}{i} U_j^{i-1} + E_{i,j}, \quad (6)$$

where $\|E_{i,j}\| \leq (i-j)\gamma$.

It is worth noting that this result (and the HD-Level-Set Decomposition) hold more generally for any “expanding poset”—the difference lies in the exact coefficients in the above relation.

The crucial observation for proving that the Bottom-Up Decomposition is a Fourier basis (that explicitly approximates the HD-Level-Set Decomposition) is that while $g_{\uparrow i}$ may not lie directly in $\text{Ker}(D_i)$ like $g_{\downarrow i}$, it is fairly close to doing so. Proving this actually relies crucially on the exact coefficients in Equation (6) which correspond to working over a simplicial complex. As a result, it is not clear the Bottom-Up Decomposition is a Fourier basis at all for general expanding posets.

Lemma 7.5. *Let (X, Π) be a two-sided γ -local-spectral expander, $f \in C_k$, and $g_{\uparrow i}$ be given as in the Bottom-Up Decomposition. Then:*

$$\|D_i g_{\uparrow i}\|_2 \leq 2^{O(i)} \gamma \left\| D_i^k f \right\|_2.$$

Proof. The result follows from directly expanding $D_i g_{\uparrow i}$:

$$\begin{aligned} D_i g_{\uparrow i} &= \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} D_i U_j^i D_j^k f \\ &= \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} \left(\frac{j}{i} U_{j-1}^{i-1} D_j + \frac{i-j}{i} U_j^{i-1} + E_{i,j} \right) D_j^k f. \end{aligned}$$

The key is then to notice that the main terms cancel. That is, setting $c_{i,j} = (-1)^{i-j} \binom{i}{j}$ we have:

$$\sum_{j=0}^i c_{i,j} \left(\frac{j}{i} U_{j-1}^{i-1} D_j + \frac{i-j}{i} U_j^{i-1} \right) D_j^k f = \sum_{j=0}^{i-1} \left(\frac{j+1}{i} c_{i,j+1} + \frac{i-j}{i} c_{i,j} \right) U_j^{i-1} D_j^k f = 0.$$

Finally by the triangle inequality we have:

$$\begin{aligned} \|D_i g_{\uparrow i}\|_2 &\leq \sum_{j=0}^i \binom{i}{j} \|E_{i,j} D_j^k f\|_2 \leq \sum_{j=0}^i \binom{i}{j} (i-j) \gamma \|D_j^k f\|_2 \\ &\leq O(i 2^i \gamma) \|D_i^k f\|_2 \end{aligned}$$

where we have used the fact that D contracts ℓ_2 -norm in the final step. \square

We will also need approximate orthogonality of both decompositions.

Lemma 7.6. *Let (X, Π) be a two-sided γ -local-spectral expander. Then the following three approximate orthogonality relations hold for all $i \neq j$:*

$$|\langle f_{\downarrow i}, f_{\downarrow j} \rangle| \leq 2^{O(k)} \gamma \|f\|_2^2 \quad (7)$$

$$|\langle f_{\downarrow i}, f_{\uparrow j} \rangle| \leq 2^{O(k)} \gamma \|f\|_2^2 \quad (8)$$

$$|\langle f_{\uparrow i}, f_{\uparrow j} \rangle| \leq \min\{k^{O(i+j)}, 2^{O(k)}\} \gamma \|f\|_2^2 \quad (9)$$

Proof. All relations follow from Equation (6) and Lemma 7.5. The first relation is proved in [DDFH18]. The latter relations follow similarly, but we give the third for completeness. In particular, assuming $i > j$ we

have:

$$\begin{aligned}
|\langle f_{\uparrow i}, f_{\uparrow j} \rangle| &= \left| \binom{k}{i} \binom{k}{j} \langle U_i^k g_{\uparrow i}, U_j^k g_{\uparrow j} \rangle \right| \\
&= \binom{k}{i} \binom{k}{j} \left| \langle D_{i-1}^k U_i^k g_{\uparrow i}, U_j^{i-1} g_{\uparrow j} \rangle \right| \\
&\leq \binom{k}{i} \binom{k}{j} \left\| D_{i-1}^k U_i^k g_{\uparrow i} \right\|_2 \left\| U_j^{i-1} g_{\uparrow j} \right\|_2 \\
&\leq \binom{k}{i} \binom{k}{j} \left\| D_{i-1}^k U_i^k g_{\uparrow i} \right\|_2 \|g_{\uparrow j}\|_2
\end{aligned}$$

where we have applied Cauchy-Schwarz and used the fact that averaging operators contract ℓ_2 -norm. We now separately bound both norms. The second is the simpler of the two, and we claim it is at most $2^{O(i)} \|f\|$. In fact a more general claim holds.

Claim 7.7. *For any ℓ_p -norm, we have:*

$$\|g_{\uparrow i}\|_p \leq 2^i \left\| D_i^k f \right\|_p$$

Proof. This follows from direct expansion of the ℓ_p -norm:

$$\begin{aligned}
\|g_{\uparrow i}\|_p &\leq \sum_{j=0}^i \binom{i}{j} \|U_j^i D_j^k f\|_p \\
&\leq \sum_{j=0}^i \binom{i}{j} \|D_i^k f\|_p \\
&= 2^i \|D_i^k f\|_p
\end{aligned}$$

where we have applied the triangle inequality and the fact that U and D contract p -norms. \square

Plugging this back into the above, we have:

$$\langle f_{\uparrow i}, f_{\uparrow j} \rangle \leq 2^i \binom{k}{i} \binom{k}{j} \|f\|_2 \left\| D_{i-1}^k U_i^k g_{\uparrow i} \right\|_2.$$

To complete the proof, it is therefore enough to argue that $\left\| D_{i-1}^k U_i^k g_{\uparrow i} \right\|_2$ is at most $c_k \gamma \|f\|$ for some small enough c_k . This follows from $k - i$ repeated applications of Equation (6) (one for each instance of the up operator). Informally, each application of Equation (6) incurs two error terms, one stemming from the matrix $E_{i,j}$, and the other from $\frac{i}{j} U_{j-1}^{i-1} D_i$ applied to $g_{\uparrow i}$, which we know is small by Lemma 7.5. The final remaining term is then proportional to the original term, but with one DU pair removed. For instance, for the first application we have:

$$\begin{aligned}
\left\| D_{i-1}^k U_i^k g_{\uparrow i} \right\| &= \left\| D_{i-1}^{k-1} (D_k U_i^k) g_{\uparrow i}, U_j^{i-1} g_{\uparrow j} \right\| \\
&= \left\| D_{i-1}^{k-1} \left(\frac{i}{k} U_{i-1}^{k-1} D_i + \frac{k-i}{k} U_i^{k-1} + E_{k,i} \right) g_{\uparrow i} \right\| \\
&= \frac{k-i}{k} \left\| D_{i-1}^{k-1} U_i^{k-1} g_{\uparrow i} \right\| + \frac{i}{k} \left\| D_{i-1}^{k-1} U_{i-1}^{k-1} D_i g_{\uparrow i} \right\| + \left\| D_{i-1}^{k-1} E_{k,i} g_{\uparrow i} \right\| \\
&\leq \frac{k-i}{k} \left\| D_{i-1}^{k-1} U_i^{k-1} g_{\uparrow i} \right\| + 2^{O(i)} \gamma \|f\| + k\gamma \|g_{\uparrow i}\| \\
&\leq \frac{k-i}{k} \left\| D_{i-1}^{k-1} U_i^{k-1} g_{\uparrow i} \right\| + k2^{O(i)} \gamma \|f\|
\end{aligned}$$

where we have used the facts that by Equation (6) and Lemma 7.5, $\|E_{k,i}\| \leq k\gamma$, and $\|D_i g_{\uparrow i}\| \leq 2^{O(i)}\gamma \|f\|$. A basic inductive argument then implies that $\|D_{i-1}^k U_i^k g_{\uparrow i}\| \leq 2^{O(i)}k^2\gamma \|f\|$, which completes the proof. For a more formal induction following exactly the same strategy, see [DDFH18, BHKL20]. \square

With these lemmas in hand, proving that the Bottom-Up Decomposition ℓ_2 -approximates the HD-Level-Set Decomposition is elementary.

Theorem 7.8. *Let (X, Π) be a two-sided γ -local-spectral expander and $f \in C_k$. Then the Bottom-Up and HD-Level-Set Decomposition are close in ℓ_2 -norm:*

$$\|f_{\uparrow i} - f_{\downarrow i}\|_2^2 \leq 2^{O(k)}\gamma \|f\|_2^2.$$

Similarly:

$$\left| \|f_{\uparrow i}\|_2^2 - \|f_{\downarrow i}\|_2^2 \right| \leq 2^{O(k)}\gamma \|f\|_2^2$$

Proof. By Lemma 7.6, we have:

$$\begin{aligned} \langle f_{\downarrow i} - f_{\uparrow i}, f_{\downarrow i} \rangle &= \langle f - f_{\uparrow i}, f_{\downarrow i} \rangle \pm 2^{O(k)}\gamma \|f\|_2^2 \\ &= \pm 2^{O(k)}\gamma \|f\|_2^2 \end{aligned}$$

Similarly:

$$\langle f_{\downarrow i} - f_{\uparrow i}, f_{\uparrow i} \rangle = \pm 2^{O(k)}\gamma \|f\|_2^2$$

Therefore:

$$\langle f_{\downarrow i} - f_{\uparrow i}, f_{\downarrow i} - f_{\uparrow i} \rangle \leq 2^{O(k)}\gamma \|f\|_2^2$$

as desired. To prove the second inequality, note that:

$$\begin{aligned} |\langle f_{\uparrow i}, f_{\uparrow i} \rangle - \langle f_{\downarrow i}, f_{\downarrow i} \rangle| &= |\langle f_{\uparrow i} - f_{\downarrow i}, f_{\uparrow i} - f_{\downarrow i} \rangle + 2\langle f_{\downarrow i}, f_{\uparrow i} - f_{\downarrow i} \rangle| \\ &\leq 2^{O(k)}\gamma \|f\|_2^2 \end{aligned}$$

by our previous observations. \square

Note that since the HD-Level-Set Decomposition satisfies the Fourier-analytic properties in Theorem 3.2 [DDFH18, AJT19, BHKL20], Theorem 7.8 implies that the Bottom-Up Decomposition does as well.

7.2 Properties of the Bottom-Up Decomposition

Our proof of hypercontractivity (Theorem 3.6) relies on a number of important structural properties of and relations between g_i and f_i . The first (and most basic) of these is the analog of a classic result for the HD-Level-Set Decomposition relating to ℓ_2 -norms of g_i and f_i .

Lemma 7.9. *Let (X, Π) be a two-sided γ -local-spectral expander and $f \in C_k$. Then:*

$$\langle g_{\uparrow i}, g_{\uparrow i} \rangle = \frac{1}{\binom{k}{i}} \langle f_{\uparrow i}, f_{\uparrow i} \rangle \pm c_{k,i}\gamma \|D_i^k f\|_2^2$$

where $c_{k,i} \leq k^{O(i)}$.

Proof. The proof is essentially the same as for the HD-Level-Set Decomposition and as our analysis above, though we repeat the idea for completeness. The key is again to apply Equation (6). In particular, recall that:

$$\begin{aligned}\langle f_{\uparrow i}, f_{\uparrow i} \rangle &= \binom{k}{i}^2 \langle U_i^k g_{\uparrow i}, U_i^k g_{\uparrow i} \rangle \\ &= \binom{k}{i}^2 \langle g_{\uparrow i}, D_i^k U_i^k g_{\uparrow i} \rangle\end{aligned}$$

by adjointness of D and U [DDFH18]. The proof is then essentially the same as Lemma 7.6. Repeated application of Equation (6) gives an error term of $O(k^2)\gamma \|D_i^k f\|_2$. The only difference is that the main term no longer has an extra occurrence of D at the end. Thus instead of becoming another error term, the main term becomes:

$$\binom{k}{i}^2 \left(\prod_{j=0}^{k-i-1} \frac{k-j-i}{k-j} \right) \langle g_{\uparrow i}, g_{\uparrow i} \rangle = \binom{k}{i} \langle g_{\uparrow i}, g_{\uparrow i} \rangle$$

which gives the result. \square

We now cover a few important bounds on g_i for *pseudorandom functions*.

Definition 7.10 (Pseudorandom). *Let (X, Π) be a simplicial complex. We say that $f \in C_k$ is (ε, i) -pseudorandom if it is sparse across all i -links in two senses:*

1. For all $\tau \in X(i)$:

$$\left| \mathbb{E}_{X_\tau}[f] \right| \leq \varepsilon \|f\|_\infty$$

2. For all $\tau \in X(i)$:

$$\langle f|_\tau, f|_\tau \rangle \leq \varepsilon \|f\|_\infty^2$$

We note that if f is non-negative, the former condition implies the latter:

$$\begin{aligned}\langle f|_\tau, f|_\tau \rangle &\leq \|f|_\tau\|_1 \|f\|_\infty \\ &= \mathbb{E}[f|_\tau] \|f\|_\infty \\ &\leq \varepsilon \|f\|_\infty^2.\end{aligned}$$

It is also worth noting that any (ε, i) -pseudorandom function is automatically (ε, j) -pseudorandom for $j \leq i$.

We now cover the first property of the Bottom-Up Decomposition that does not follow from standard HDX analysis, the behavior of level functions under restriction. Analysis of restrictions is a classic Fourier analytic tool, and the fact that our decomposition behaves nicely under restriction is a major advantage over previous decompositions which have no clear local structure in this sense. For this particular work, we'll mostly be interested in the following bound on the ℓ_2 -norm of restrictions.

Proposition 7.11. *Let (X, Π) be a two-sided γ -local-spectral expander and $f \in C_k$ an (ε, i) -pseudorandom function. Then for any $j \leq i \leq k$ and $\tau \in X(j)$:*

$$\langle g_{\uparrow i}|_\tau, g_{\uparrow i}|_\tau \rangle \leq \left(\frac{\varepsilon}{\binom{k-j}{i-j}} + c_{k,i}\gamma \right) \|f\|_\infty^2,$$

where $c_{k,i} \leq k^{O(i)}$.

Proving this, however, requires a more general understanding of the Bottom-Up Decomposition under restriction. The key observation is that the restriction of our level functions is closely related to the level functions of the restriction. More formally, for any $\tau \in X(j)$ let $g_{\uparrow \ell}^{(\tau)}$ denote the Bottom-Up Decomposition of $f|_{\tau}$. Then following relation between $g_i|_{\tau}$ and the $g_{\uparrow \ell}^{(\tau)}$ holds.

Lemma 7.12. *Let (X, Π) be a pure, weighted simplicial complex, and $f \in C_k$. Then for any $j \leq i \leq k$ and $\tau \in X(j)$:*

$$g_{\uparrow i}|_{\tau} = \sum_{\sigma \subseteq \tau} (-1)^{|\sigma|} g_{\uparrow i-j}^{(\tau \setminus \sigma)}.$$

Proof. This follows almost immediately from directly expanding the definition of $g_{\uparrow i}|_{\tau}$. In particular, recall that for all $I \in X(i-j)$, we have by Theorem 7.2:

$$g_{\uparrow i}|_{\tau}(I) = \sum_{T \subseteq I \cup \tau} (-1)^{|(\tau \cup I) \setminus T|} \mathbb{E}_{X_T} [f].$$

The trick is to notice that we can divide up this sum over $T \subseteq I \cup \tau$ by T 's intersection with τ . It will be convenient to phrase this in the following way. Let \mathcal{T} denote the set of all sub-faces $T \subseteq I \cup \tau$, and for each $\sigma \subset \tau$, let \mathcal{T}_{σ} be the set of sub-faces $T \subset I \cup \tau$ such that $T \cap \tau = \tau \setminus \sigma$. Notice that for any $\sigma \neq \sigma'$, \mathcal{T}_{σ} and $\mathcal{T}_{\sigma'}$ are disjoint, and that the union of these families is exactly \mathcal{T} . Together, this means that we can break up the above sum by first summing over σ , and then every $T \in \mathcal{T}_{\sigma}$:

$$g_{\uparrow i}|_{\tau}(I) = \sum_{\sigma \subseteq \tau} \left(\sum_{T \in \mathcal{T}_{\sigma}} (-1)^{|(I \cup \tau) \setminus T|} \mathbb{E}_{X_T} [f] \right)$$

By definition, every $T \in \mathcal{T}_{\sigma}$ can be written as $T' \cup (\tau \setminus \sigma)$. Plugging this into the above gives the result:

$$\begin{aligned} \sum_{\sigma \subseteq \tau} \left(\sum_{T \in \mathcal{T}_{\sigma}} (-1)^{|(I \cup \tau) \setminus T|} \mathbb{E}_{X_T} [f] \right) &= \sum_{\sigma \subseteq \tau} \left(\sum_{T' \cup (\tau \setminus \sigma) \in \mathcal{T}_{\sigma}} (-1)^{|(I \cup \tau) \setminus (T' \cup (\tau \setminus \sigma))|} \mathbb{E}_{X_{T' \cup (\tau \setminus \sigma)}} [f] \right) \\ &= \sum_{\sigma \subseteq \tau} (-1)^{|\sigma|} \left(\sum_{T' \cup (\tau \setminus \sigma) \in \mathcal{T}_{\sigma}} (-1)^{|I \setminus T'|} \mathbb{E}_{X_{T' \cup (\tau \setminus \sigma)}} [f] \right) \\ &= \sum_{\sigma \subseteq \tau} (-1)^{|\sigma|} g_{\uparrow i-j}^{(\tau \setminus \sigma)}(I), \end{aligned}$$

where the final step comes from the fact that the inner summation over \mathcal{T}_{σ} is equivalent to summing over all T' in the link of $\tau \setminus \sigma$. \square

We note that the same result was known to hold for the Bottom-Up Decomposition over the complete complex [KMMS18], who proved the result by induction using the recursive form of the decomposition. The same strategy will work for general simplicial complexes, but we find using the explicit form as above to be a bit simpler.

With this in hand, proving Proposition 7.11 is fairly elementary and follows similarly to its analogous statement for the complete complex (see [KMMS18, Corollary 3.4]).

Proof of Proposition 7.11. An application of Lemma 7.12 and Cauchy-Schwarz implies that:

$$\begin{aligned} g_{\uparrow i}|_{\tau}(T)^2 &= \left(\sum_{\sigma \subseteq \tau} (-1)^{|\sigma|} g_{\uparrow i-j}^{(\tau \setminus \sigma)}(T) \right)^2 \\ &\leq 2^{O(i)} \sum_{\sigma \subseteq \tau} g_{\uparrow i-j}^{(\tau \setminus \sigma)}(T)^2 \end{aligned}$$

Then applying Lemma 7.9 gives:

$$\begin{aligned} \langle g_{\uparrow i}|_{\tau}, g_{\uparrow i}|_{\tau} \rangle &\leq 2^{O(i)} \sum_{\sigma \subseteq \tau} \langle g_{\uparrow i-j}^{(\tau \setminus \sigma)}, g_{\uparrow i-j}^{(\tau \setminus \sigma)} \rangle \\ &\leq 2^{O(i)} \sum_{\sigma \subseteq \tau} \frac{\langle f|_{(\tau \setminus \sigma)}, f|_{(\tau \setminus \sigma)} \rangle}{\binom{k-j+|\sigma|}{i-j}} + c_1 \gamma \|f\|_{\infty}^2 \\ &\leq 2^{O(i)} \sum_{\sigma \subseteq \tau} \frac{\varepsilon \|f\|_{\infty}^2}{\binom{k-j+|\sigma|}{i-j}} + c_1 \gamma \|f\|_{\infty}^2 \\ &\leq 2^{O(i)} \frac{\varepsilon \|f\|_{\infty}^2}{\binom{k-j}{i-j}} + c_2 \gamma \|f\|_{\infty}^2 \end{aligned}$$

where $c_1, c_2 \leq k^{O(i)}$. □

Finally, it will also be useful to bound the infinity norm of $g_{\uparrow i}$ as well. Our final property shows that $\|g_{\uparrow i}\|_{\infty}$ is particularly small when f is pseudorandom.

Lemma 7.13. *Let (X, Π) be a two-sided γ -local-spectral expander and $f \in C_k$ be any (ε, i) -pseudorandom function satisfying $\mathbb{E}[f] \leq \varepsilon \|f\|_{\infty}$. Then the infinity norm of $g_{\uparrow i}$ is small:*

$$\|g_{\uparrow i}\|_{\infty} \leq 2^i \varepsilon \|f\|_{\infty}$$

Proof. This is immediate from combining the explicit form of $g_{\uparrow i}$ with (ε, i) -pseudorandomness.

$$\begin{aligned} |g_{\uparrow i}(w)| &= \left| \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} U_j^i D_j^k f(w) \right| \\ &\leq \sum_{j=0}^i \binom{i}{j} \varepsilon \|f\|_{\infty} \\ &= 2^i \varepsilon \|f\|_{\infty} \end{aligned}$$

where we have used the observation that since f is (ε, i) -pseudorandom, for all $w \in X(i)$:

$$|U_j^i D_j^k f(w)| \leq \varepsilon \|f\|_{\infty}$$

□

8 Hypercontractivity on HDX

In this section, we prove a hypercontractivity theorem for the Bottom-Up Decomposition on two-sided local-spectral expanders. Since we will work only with the Bottom-Up Decomposition in this section, we drop the \uparrow for simplicity and simply write $f = \sum f_i$ for $f_i = \binom{k}{i} U_i^k g_i$ and $g_i = g_{\uparrow i}$ as defined in the Bottom-Up Decomposition.

Theorem 8.1. *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq k^{-\Omega(i)}$, and $f \in C_k$ an (ε, i) -pseudorandom function. If $f = f_0 + \dots + f_k$ is the Bottom-Up Decomposition of f , then:*

$$\mathbb{E}[f_i^4] \leq 2^{O(i)} \varepsilon \mathbb{E}[f_i^2] \|f\|_\infty^2 + c_{k,i} \varepsilon \gamma^{1/2} \left\| D_i^k f \right\|_2^2 \|f\|_\infty^2$$

where $c_{k,i} \leq k^{O(i)}$.¹⁵

The proof of Theorem 8.1 can get a bit technical at points, so for simplicity of notation, we note it is sufficient to prove the result assuming $\|f\|_\infty = 1$. Given a general function f , applying this to $\frac{f}{\|f\|_\infty}$ gives the general form in Theorem 8.1. Keeping this in mind, we'll start by laying out our general strategy for analyzing the fourth moment. Let $[\tau]_i = \{a \subseteq \tau : a \in X(i)\}$, and note that $f_i(\tau) = \sum_{a \in [\tau]_i} g_i(a)$. Using this notation, we can expand out the 4th moment of f_i :

$$\mathbb{E}[f_i^4] = \mathbb{E}_{\tau \in X(k)} \sum_{a,b,c,d \in [\tau]_i} g_i(a)g_i(b)g_i(c)g_i(d) = \sum_{a,b,c,d \in X(i)} \pi_k(X_{a \cup b \cup c \cup d}) g_i(a)g_i(b)g_i(c)g_i(d),$$

where the indices a, b, c, d are ordered. We can further simplify this by grouping the terms by size of $a \cup b \cup c \cup d$:

$$\mathbb{E}[f_i^4] = \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{e \in X(\ell)} \pi_\ell(e) \sum_{a,b,c,d \in X(i): a \cup b \cup c \cup d = e} g_i(a)g_i(b)g_i(c)g_i(d).$$

Analyzing the RHS directly is difficult, so taking after [KMMS18], we will partition the term even further by summing over fixed **intersection patterns** of a, b, c , and d (an intersection pattern fixes the intersection size of every subset of $\{a, b, c, d\}$). Denote the set of such patterns where $|a \cup b \cup c \cup d| = \ell$ by Σ_ℓ , and for any $e \in X(\ell)$, and $\sigma \in \Sigma_\ell$, let $\sigma(e)$ denote all tuples (a, b, c, d) such that $a \cup b \cup c \cup d = e$, $(a, b, c, d) \in \sigma$. We may now write:

$$\mathbb{E}[f_i^4] = \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{\sigma \in \Sigma_\ell} \sum_{e \in X(\ell)} \pi_\ell(e) \sum_{a,b,c,d \in \sigma(e)} g_i(a)g_i(b)g_i(c)g_i(d).$$

We make one final simplification of the above before moving to analysis. Let x_1, \dots, x_ℓ be random variables which take on vertex values in the complex. For each intersection pattern $\sigma \in \Sigma_\ell$, let $I_1^\sigma, \dots, I_4^\sigma$ be size- i subsets of $\{x_1, \dots, x_\ell\}$ whose union is $\{x_1, \dots, x_\ell\}$ and which satisfy the intersection pattern σ . Then we can simplify the above as the following expectation over the x_i :

$$\mathbb{E}[f_i^4] = \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{\sigma \in \Sigma_\ell} \beta(\sigma) \mathbb{E}_{x_1 \in X(1)} \left[\mathbb{E}_{x_2 \in X_{x_1}(1)} \dots \left[\mathbb{E}_{x_\ell \in X_{x_1, \dots, x_{\ell-1}}(1)} [g_i(I_1^\sigma)g_i(I_2^\sigma)g_i(I_3^\sigma)g_i(I_4^\sigma)] \right] \right].$$

¹⁵Note that this can be improved to $c_{k,i} \leq \max\{2^{O(i)}, \binom{k}{i}^{O(1)}\}$, but since we generally consider the regime of $i \ll k$ we use $k^{O(i)}$ throughout for simplicity.

where $\beta(\sigma) \leq 2^{O(i)}$ is a parameter dependent on the intersection pattern that accounts for the new normalization of terms in the nested expectation.¹⁶ For simplicity of notation, we will instead write the right-hand side as:

$$\mathbb{E}[f_i^4] = \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{\sigma \in \Sigma_\ell} \beta(\sigma) \mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1^\sigma) g_i(I_2^\sigma) g_i(I_3^\sigma) g_i(I_4^\sigma)]$$

where it is understood that $\mathbb{E}_{x_1, \dots, x_\ell}$ is a shorthand for the nested expectation $\mathbb{E}_{x_1 \in X(1)} \mathbb{E}_{x_2 \in X_{x_1}(1)} \dots \mathbb{E}_{x_\ell \in X_{x_1, \dots, x_{\ell-1}}(1)}$. We will use this convention throughout the rest of the proof, as the nested notation is cumbersome to write otherwise.

Our goal is now to upper bound this sum to get a hypercontractive inequality. We do this by bounding each sign pattern independently.

Claim 8.2. *For every sign pattern σ , the corresponding expectation is bounded by:*

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1^\sigma) g_i(I_2^\sigma) g_i(I_3^\sigma) g_i(I_4^\sigma)] \leq \left(\frac{i}{k}\right)^\ell 2^{O(i)} \varepsilon \mathbb{E}[f_i^2] + c_{k,i} \varepsilon \gamma^{1/2} \left\| D_i^k f \right\|_2^2$$

where $c_{k,i} \leq k^{O(i)}$.

Before jumping into the proof of Claim 8.2, let's show how it can be used to prove Theorem 8.1.

Proof of Theorem 8.1. Recall it is sufficient to prove the result assuming $\|f\|_\infty = 1$. As discussed earlier in the section, expanding the 4th moment gives the following relation:

$$\mathbb{E}[f_i^4] = \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{\sigma} \beta(\sigma) \mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1^\sigma) g_i(I_2^\sigma) g_i(I_3^\sigma) g_i(I_4^\sigma)].$$

Applying Claim 8.2 to the righthand side gives:

$$\begin{aligned} \mathbb{E}[f_i^4] &\leq \sum_{\ell=i}^{4i} \binom{k}{\ell} \sum_{\sigma \in \Sigma_\ell} \beta(\sigma) \left(\left(\frac{i}{k}\right)^\ell 2^{O(i)} \varepsilon \mathbb{E}[f_i^2] + c_{k,i} \varepsilon \gamma^{1/2} \left\| D_i^k f \right\|_2^2 \right) \\ &\leq \sum_{\ell=i}^{4i} \left(\frac{ek}{\ell}\right)^\ell \left(\frac{i}{k}\right)^\ell \left(\sum_{\sigma \in \Sigma_\ell} \beta(\sigma) \right) \left(2^{O(i)} \varepsilon \mathbb{E}[f_i^2] + c_1 \varepsilon \gamma^{1/2} \left\| D_i^k f \right\|_2^2 \right) \\ &\leq 2^{O(i)} \varepsilon \mathbb{E}[f_i^2] + c_2 \varepsilon \gamma^{1/2} \left\| D_i^k f \right\|_2^2 \end{aligned}$$

where $c_1, c_2 \leq k^{O(i)}$ and the last step follows from noting that there are at most $\text{poly}(i)$ intersection patterns. \square

¹⁶The bound $2^{O(i)}$ follows from a simple counting argument. Fix an intersection pattern $\sigma = \{a_1, \dots, a_{15}\}$. Expanding out the expectation $\mathbb{E}_{x_1 \in X(1)} \left[\mathbb{E}_{x_2 \in X_{x_1}(1)} \dots \left[\mathbb{E}_{x_\ell \in X_{x_1, \dots, x_{\ell-1}}(1)} [g_i(I_1^\sigma) g_i(I_2^\sigma) g_i(I_3^\sigma) g_i(I_4^\sigma)] \right] \right]$. The coefficient of any term $g_i(a) g_i(b) g_i(c) g_i(d)$ is given by $\frac{\pi_\ell(e)}{\ell!}$ times the number of permutations of x_1, \dots, x_ℓ that fix $g_i(a) g_i(b) g_i(c) g_i(d)$. This latter value is the same for any term, and can be lower bounded by $\prod_{i=1}^{15} a_i!$, so the final coefficient is at least $\pi_\ell(e) \frac{\prod a_i!}{\ell!}$. To cancel this value, $\beta(\sigma)$ is therefore at most $\frac{\ell!}{\prod a_i!} \leq 2^{O(i)}$ as desired.

8.1 Proving Claim 8.2

The main technical work comes in proving Claim 8.2, which relies heavily on Garland's method and our new localization strategy for decorrelating variables (Corollary 6.4).

We split the proof into two parts. First, we will show that any pattern which has a unique element (i.e. some x_i which appears only in one of the four sets) may be disregarded.

Proposition 8.3. *If σ is a pattern in which any variable is unique (appears in only one I_j), then:*

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1^\sigma)g_i(I_2^\sigma)g_i(I_3^\sigma)g_i(I_4^\sigma)] \leq 2^{O(i)}\gamma\varepsilon^2 \left\| D_i^k f \right\|_2^2.$$

Proof. To simplify notations in the proof, let $I = \{x_1, \dots, x_\ell\}$ and $I_j = I_j^\sigma$. Assume without loss of generality that I_4 has a unique variable x_ℓ , and set $J = I \setminus \{x_\ell\}$ and $J_4 = I_4 \setminus \{x_\ell\}$. We can re-write our expectation as:

$$(*) = \mathbb{E}_J \left[g_i(I_1)g_i(I_2)g_i(I_3) \mathbb{E}_{x_\ell \in X_{J(1)}} [g_i|_{J_4}(x_\ell)] \right].$$

By Corollary 6.4, the inner expectation can be replaced with $D_i g_i(J_4)$ up to γ error in the following sense. Consider any fixing of the variables J (namely, fixing $x_1, \dots, x_{\ell-1}$), we have:

$$\begin{aligned} \mathbb{E}_{x_\ell \in X_{J(1)}} [g_i|_{J_4}(x_\ell)] &= \mathbb{E}_{x_\ell \in X_{J_4(1)}} [g_i|_{J_4}(x_\ell)] + \Gamma g_i|_{J_4}(J \setminus J_4) \\ &= D_i g_i(J_4) + \Gamma g_i|_{J_4}(I \setminus I_4) \end{aligned}$$

where $\|\Gamma\| \leq O(i\gamma)$ by Corollary 6.4. Plugging this back into our original expectation gives:

$$(*) = \mathbb{E}_J [g_i(I_1)g_i(I_2)g_i(I_3)D_i g_i(J_4)] + \mathbb{E}_J [g_i(I_1)g_i(I_2)g_i(I_3)\Gamma g_i|_{J_4}(I \setminus I_4)]$$

The idea is now to split each term into two parts: the first three terms $g_i(I_1)g_i(I_2)g_i(I_3)$ and the last term. Let's first split these portions by Cauchy-Schwarz to get:

$$\begin{aligned} (*) &\leq \mathbb{E}_J [g_i(I_1)^2 g_i(I_2)^2 g_i(I_3)^2]^{1/2} \\ &\quad \cdot \left(\mathbb{E}_J [D_i g_i(J_4)^2]^{1/2} + \mathbb{E}_{J_4} \left[\mathbb{E}_{I \setminus I_4 \in X_{J_4}} [\Gamma g_i|_{J_4}(I \setminus I_4)^2] \right]^{1/2} \right) \end{aligned}$$

where we have re-arranged variable for convenience in the last term. We now bound each term separately.

The first term can be bounded by the observation that $\|g_i\|_\infty \leq 2^{O(i)}\varepsilon \|f\|_\infty$ (Lemma 7.13), and hence:

$$\mathbb{E}_J [g_i(I_1)^2 g_i(I_2)^2 g_i(I_3)^2]^{1/2} \leq 2^{O(i)}\varepsilon^2 \|g_i\|_2$$

where we simply bounded two of the three g_i^2 terms by their infinity norm and applied Garland's lemma for localizations (Lemma 6.2) to remove the extra variables.

We next analyze the second term. The first summand is exactly $\|D_i g_i\|$, which by Lemma 7.5 is at most $O(\gamma \|D_i^k f\|)$. The second summand is a bit trickier, but can be analyzed through a combination of standard spectral bounds and Garland's lemma for restrictions (Lemma 6.1). In particular, re-writing the inner expectation as an inner-product we get:

$$\begin{aligned} \mathbb{E}_{J_4} \left[\mathbb{E}_{I \setminus I_4 \in X_{J_4}} [\Gamma g_i|_{J_4}(I \setminus I_4)^2] \right]^{1/2} &= \mathbb{E}_{J_4} [\langle \Gamma g_i|_{J_4}, \Gamma g_i|_{J_4} \rangle]^{1/2} \\ &\leq c\gamma \mathbb{E}_{J_4} [\langle g_i|_{J_4}, g_i|_{J_4} \rangle]^{1/2} \\ &= c\gamma \|g_i\|_2, \end{aligned}$$

where $c \leq O(i)$ and we have applied the fact that $\|\Gamma\| \leq O(i\gamma)$ and Garland's lemma for restrictions (Lemma 6.1). Recalling from Claim 7.7 that $\|g_i\|_2 \leq 2^i \|D_i^k f\|_2$ completes the result. \square

We may now restrict our analysis to patterns in which every variable appears at least twice. Note that this implies $\ell \leq 2i$, which is important because we expect our expectation to scale at best with k^{-2i} , so any terms with $\ell > 2i$ would cause difficulty. As in [KMMS18], we break this analysis into two steps. Let I_1, \dots, I_4 satisfy intersection pattern σ as above (we drop the σ superscript for convenience), and let H_i for $i \in \{2, 3, 4\}$ denote the set of variables that appear i times.

We'll start by handling H_2 through a combination of Cauchy-Schwarz, Garland's method, and our localization technique for higher moments. Unlike the case of the complete complex studied in [KMMS18], these latter components are necessary due to the fact that local-spectral expanders are generally far from product spaces (a crucial property of the complete complex exploited in [KMMS18]). The proof is fairly technical, so we'll start by laying out some convenient notation. For any $0 \leq m \leq \ell$, let $T^m = \{x_1, \dots, x_m\}$. Let $j = |H_3 \cup H_4|$ where $0 \leq j \leq \ell$. Noting that re-ordering the variables x_1, \dots, x_ℓ has no effect on the distribution, we may assume without loss of generality that $H_3 \cup H_4 = \{x_1, \dots, x_j\}$ (where if $j = 0$ then $H_3 \cup H_4$ is empty). Finally, we introduce two useful notations: for $m \leq \ell$ let $I_r^m = I_r \cap \{x_1, \dots, x_m\}$ and $s_r^m = i - |I_r^m|$.

Proposition 8.4.

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] \leq \mathbb{E}_{x_1, \dots, x_j} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{T^j}(s_r^j)} [g_i^2|_{I_r^j}(\tau_r)]} \right] + 2^{O(i)}\gamma^{1/2}\varepsilon^2 \|g_i\|_2^2$$

Proof. The proof follows from an inductive argument where we pull one variable in $x \in H_2$ inside the sum in each step by de-correlating the two copies of g_i which do not take x as an input, and then applying Cauchy-Schwarz. In particular, we will show by induction that for all $\ell \geq m \geq j$:

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] \leq \mathbb{E}_{x_1, \dots, x_m} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{T^m}(s_r^m)} [g_i^2|_{I_r^m}(\tau_r)]} \right] + 2^{O(i)}\gamma^{1/2}\varepsilon^2 \|g_i\|_2^2$$

The base case ($m = \ell$) is trivial. Since we also done if $m = j$, we may assume that $x_m \in H_2$ and therefore lies in exactly two of I_1, I_2, I_3, I_4 by definition. Assume without loss of generality that $x_m \in I_3, I_4$. We'd like to pull x_m inside the expectation. The issue is that despite the fact that x_m does not participate in I_1 or I_2 , these terms actually depend on x_m regardless since τ_1 and τ_2 are drawn from a link that includes x_m . To fix this, we can use Corollary 6.4 to de-correlate these terms from x_m :

$$\begin{aligned} \sqrt{\prod_{r=1}^2 \mathbb{E}_{\tau_r \sim X_{T^m}(s_r^m)} [g_i^2|_{I_r^m}(\tau_r)]} &= \sqrt{\prod_{r=1}^2 \left(\mathbb{E}_{\tau_r \sim X_{T^{m-1}}(s_r^{m-1})} [g_i^2|_{I_r^{m-1}}(\tau_r)] + \Gamma g_i^2|_{I_r^{m-1}}(x_m) \right)} \\ &\leq \sqrt{\mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2|_{I_1^{m-1}}(\tau_1)]} \sqrt{\mathbb{E}_{\tau_2 \sim X_{T^{m-1}}(s_2^{m-1})} [g_i^2|_{I_2^{m-1}}(\tau_1)]} \\ &\quad + \sqrt{\Gamma g_i^2|_{I_2^{m-1}}(x_m)} \sqrt{\mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2|_{I_1^{m-1}}(\tau_1)]} \\ &\quad + \sqrt{\Gamma g_i^2|_{I_1^{m-1}}(x_m)} \sqrt{\mathbb{E}_{\tau_2 \sim X_{T^{m-1}}(s_2^{m-1})} [g_i^2|_{I_2^{m-1}}(\tau_2)]} \\ &\quad + \sqrt{\Gamma g_i^2|_{I_1^{m-1}}(x_m)} \sqrt{\Gamma g_i^2|_{I_2^{m-1}}(x_m)} \end{aligned}$$

where $\|\Gamma\| \leq O(i\gamma)$ and we have used the fact that by assumption $I_r^m = I_r^{m-1}$ for $r = 1, 2$. For the moment, denote the last 3 terms by $err(g)$. Then by the inductive hypothesis we have:

$$\begin{aligned}
& \mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] \\
& \leq \mathbb{E}_{x_1, \dots, x_m} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{T^m}(s_r)} [g_i^2 |_{I_r^m}(\tau_r)]} \right] + 2^{O(i)} \gamma^{1/2} \varepsilon^2 \|g_i\|_2^2 \\
& \leq \mathbb{E}_{x_1, \dots, x_{m-1}} \left[\sqrt{\prod_{r=1}^2 \mathbb{E}_{\tau_r \sim X_{T^{m-1}}(s_r)} [g_i^2 |_{I_r^{m-1}}(\tau_r)]} \mathbb{E}_{x_m} \left[\sqrt{\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3)} [g_i^2 |_{I_3^m}(\tau_3)]} \sqrt{\mathbb{E}_{\tau_4 \sim X_{T^m}(s_4)} [g_i^2 |_{I_4^m}(\tau_4)]} \right] \right] \\
& + \mathbb{E}_{x_1, \dots, x_m} \left[err(g) \sqrt{\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3^m)} [g_i^2 |_{I_3^m}(\tau_3)]} \sqrt{\mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 |_{I_4^m}(\tau_4)]} \right] + 2^{O(i)} \gamma^{1/2} \varepsilon^2 \|g_i\|_2^2
\end{aligned}$$

By Cauchy-Schwarz, the first term can be bounded by:

$$\mathbb{E}_{x_1, \dots, x_{m-1}} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{T^{m-1}}(s_r^{m-1})} [g_i^2 |_{I_r^{m-1}}(\tau_r)]} \right],$$

so it is enough to show that the latter error term is small. We'll analyze each term in $err(g)$ independently using Cauchy-Schwarz, Garland's method, and our bound on $\|g\|_\infty$. Starting with the first term, an application of Cauchy-Schwarz gives:

$$\begin{aligned}
& \mathbb{E}_{x_1, \dots, x_m} \left[\sqrt{\Gamma g_i^2 |_{I_2^{m-1}}(x_m)} \sqrt{\mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2 |_{I_1^{m-1}}(\tau_1)]} \sqrt{\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3^m)} [g_i^2 |_{I_3^m}(\tau_3)]} \sqrt{\mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 |_{I_4^m}(\tau_4)]} \right] \\
& \leq \mathbb{E}_{x_1, \dots, x_m} \left[\Gamma g_i^2 |_{I_2^{m-1}}(x_m) \mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2 |_{I_1^{m-1}}(\tau_1)] \right]^{1/2} \mathbb{E}_{x_1, \dots, x_m} \left[\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3^m)} [g_i^2 |_{I_3^m}(\tau_3)] \mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 |_{I_4^m}(\tau_4)] \right]^{1/2}.
\end{aligned}$$

The righthand expectation is easy to analyze using the fact that $\|g_i\|_\infty \leq 2^{O(i)} \varepsilon$:

$$\begin{aligned}
\mathbb{E}_{x_1, \dots, x_m} \left[\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3^m)} [g_i^2 |_{I_3^m}(\tau_3)] \mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 |_{I_4^m}(\tau_4)] \right]^{1/2} & \leq 2^{O(i)} \varepsilon \mathbb{E}_{x_1, \dots, x_m} \left[\mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 |_{I_4^m}(\tau_4)] \right]^{1/2} \\
& = 2^{O(i)} \varepsilon \mathbb{E}_{\tau \sim X(I_4^m)} [\langle g_i |_\tau, g_i |_\tau \rangle]^{1/2} \\
& = 2^{O(i)} \varepsilon \|g_i\|_2
\end{aligned}$$

where the last two equalities follow from Garland's method. Turning our attention to the lefthand expectation, we can apply Cauchy-Schwarz to get:

$$\begin{aligned}
& \mathbb{E}_{x_1, \dots, x_m} \left[\Gamma g_i^2 |_{I_2^{m-1}}(x_m) \mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2 |_{I_1^{m-1}}(\tau_1)] \right]^{1/2} \\
& \leq \mathbb{E}_{x_1, \dots, x_{m-1}} \left[\langle \Gamma g_i^2 |_{I_2^{m-1}}, \Gamma g_i^2 |_{I_2^{m-1}} \rangle \right]^{1/4} \mathbb{E}_{x_1, \dots, x_{m-1}} \left[\mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2 |_{I_1^{m-1}}(\tau_1)]^2 \right]^{1/4} \\
& \leq 2^{O(i)} \gamma^{1/2} \mathbb{E}_{x_1, \dots, x_{m-1}} \left[\langle g_i^2 |_{I_2^{m-1}}, g_i^2 |_{I_2^{m-1}} \rangle \right]^{1/4} \mathbb{E}_{x_1, \dots, x_{m-1}} \left[\langle g_i^2 |_{I_1^{m-1}}, g_i^2 |_{I_1^{m-1}} \rangle \right]^{1/4}
\end{aligned}$$

where in the last step we have applied the fact that $\|\Gamma\| \leq O(i\gamma)$. Analysis of the remaining expectations follows exactly as before. In particular, re-arranging variables by symmetry and applying Garland's method, we can continue the above inequality as follows:

$$\begin{aligned}
&= \gamma^{1/2} 2^{O(i)} \mathbb{E}_{\tau \sim X(I_2^{m-1})} [\langle g_i^2 | \tau, g_i^2 | \tau \rangle]^{1/4} \mathbb{E}_{\tau \sim X(I_1^{m-1})} [\langle g_i^2 | \tau, g_i^2 | \tau \rangle]^{1/4} \\
&= \gamma^{1/2} 2^{O(i)} \langle g_i^2, g_i^2 \rangle^{1/2} \\
&\leq \gamma^{1/2} 2^{O(i)} \varepsilon \|g_i\|_2
\end{aligned}$$

where in the final step we have again applied our bound on $\|g_i\|_\infty$. Putting the analysis of these two terms together, we get the desired bound on the first summand of $err(g)$:

$$\begin{aligned}
&\mathbb{E}_{x_1, \dots, x_m} \left[\sqrt{\Gamma g_i^2 | I_2^{m-1}(x_m)} \sqrt{\mathbb{E}_{\tau_1 \sim X_{T^{m-1}}(s_1^{m-1})} [g_i^2 | I_1^{m-1}(\tau_1)]} \sqrt{\mathbb{E}_{\tau_3 \sim X_{T^m}(s_3^m)} [g_i^2 | I_3^m(\tau_3)]} \sqrt{\mathbb{E}_{\tau_4 \sim X_{T^m}(s_4^m)} [g_i^2 | I_4^m(\tau_4)]} \right] \\
&\leq \gamma^{1/2} 2^{O(i)} \varepsilon^2 \|g_i\|_2^2
\end{aligned}$$

Thankfully, the analysis of second summand in $err(g)$ is exactly the same, and the third term differs only in that the lefthand expectation in the previous analysis becomes:

$$\mathbb{E}_{x_1, \dots, x_m} \left[\Gamma g_i^2 | I_1^{m-1}(x_m) \Gamma g_i^2 | I_2^{m-1}(x_m) \right]^{1/2} \leq \gamma 2^{O(i)} \varepsilon \|g_i\|_2$$

by the same arguments. Combining these together, we get that our error term is bounded by $\gamma^{1/2} 2^{O(i)} \varepsilon^2 \|g_i\|_2^2$, which completes the proof. \square

It is left to analyze H_3 and H_4 . Recalling that we've assumed $\{x_1, \dots, x_j\} = H_3 \cup H_4$, Proposition 8.4 can be restated as:

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] \leq \mathbb{E}_{H_3 \cup H_4} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 | I_r^j(\tau_r)]} \right] + 2^{O(i)} \gamma^{1/2} \varepsilon^2 \|g_i\|_2^2.$$

The key is now to apply Proposition 7.11, which says that the maximum of the inner restricted expectations are small, where the factor is better the fewer variables we restrict. In order to minimize the number of restrictions, we use Cauchy-Schwarz to separate out I_1 and I_2 from I_3 and I_4 :

$$\begin{aligned}
\mathbb{E}_{H_3 \cup H_4} \left[\sqrt{\prod_{r=1}^4 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 | I_r^j(\tau_r)]} \right] &\leq \mathbb{E}_{H_3 \cup H_4} \left[\prod_{r=1}^2 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 | I_r^j(\tau_r)] \right]^{1/2} \\
&\cdot \mathbb{E}_{H_3 \cup H_4} \left[\prod_{r=3}^4 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 | I_r^j(\tau_r)] \right]^{1/2}.
\end{aligned}$$

Analysis of these two terms is the same, so we focus on the former. The idea is to bound one of the two inner expectations (say I_1) by its maximum, and note that the other term then simply returns $\|g_i\|$. Unfortunately, there is a slight issue with this strategy naively: H_3 may contain variables that are not in I_1 , so we cannot directly apply Proposition 7.11. Thankfully, localization again comes to our rescue: we can apply Proposition 7.11 as long as we first de-correlate I_1 from the extraneous variables in H_3 as in Proposition 8.4.

More formally, letting $B_{12} = (H_3 \cap I_1 \cap I_2) \cup H_4$ for simplicity of notation, by exactly the same inductive argument used in Proposition 8.4 we have:

$$\mathbb{E}_{H_3 \cup H_4} \left[\prod_{r=1}^2 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 |_{I_r^j}(\tau_r)] \right]^{1/2} \leq \mathbb{E}_{B_{12} \sim X} \left[\mathbb{E}_{\tau_1 \sim X_{B_{12}}} [g_i^2 |_{I_1^{B_{12}}}(\tau_1)] \mathbb{E}_{\tau_2 \sim X_{B_{12}}} [g_i^2 |_{I_2^{B_{12}}}(\tau_2)] \right]^{1/2} + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|g_i\|_2$$

where for the moment we have omitted the sizes of B_{12} , τ_1 , and τ_2 for simplicity (these will be computed soon). Pulling out the maximal I_1 term and applying Proposition 7.11 with $j = |B_{12}|$, we then get:

$$\begin{aligned} \mathbb{E}_{H_3 \cup H_4} \left[\prod_{r=1}^2 \mathbb{E}_{\tau_r \sim X_{H_3 \cup H_4}(s_r^j)} [g_i^2 |_{I_r^j}(\tau_r)] \right]^{1/2} &\leq \max_{B_{12}} \left(\mathbb{E}_{\tau_1 \sim X_{B_{12}}} [g_i^2 |_{I_1^{B_{12}}}(\tau_1)] \right)^{1/2} \|g_i\|_2 + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|g_i\|_2 \\ &\leq \frac{2^{O(i)} \varepsilon^{1/2}}{\binom{k-|B_{12}|}{i-|B_{12}|}^{1/2}} \|g_i\|_2 + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|g_i\|_2 \end{aligned}$$

The same argument holds for the latter product over I_3 and I_4 . Letting $B_{12} = (H_3 \cap I_3 \cap I_4) \cup H_4$, and putting everything together, we finally get the bound:

$$\begin{aligned} \mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] &\leq \frac{2^{O(i)} \varepsilon}{\binom{k-|B_{12}|}{i-|B_{12}|}^{1/2} \binom{k-|B_{34}|}{i-|B_{34}|}^{1/2}} \|g_i\|_2^2 + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|g_i\|_2^2 \\ &\leq \frac{2^{O(i)} \varepsilon}{\binom{k-|B_{12}|}{i-|B_{12}|}^{1/2} \binom{k-|B_{34}|}{i-|B_{34}|}^{1/2}} \frac{1}{\binom{k}{i}} \|f_i\|_2^2 + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|D_i^k f\|_2^2 \\ &\leq 2^{O(i)} \varepsilon \left(\frac{k}{i} \right)^{\frac{|B_{12}|+|B_{34}|-2i}{2}} \|f_i\|_2^2 + 2^{O(i)} \varepsilon^2 \gamma^{1/2} \|D_i^k f\|_2^2 \end{aligned}$$

where we have applied the basic binomial bound $\binom{n}{p} \geq \left(\frac{n}{p}\right)^p$. To complete the result, it suffices show that $\frac{|B_{12}|+|B_{34}|}{2} = 2i - \ell$. This follows similarly to the analogous argument in [KMMS18], but we'll give a simplification of their proof for completeness. Recall that B_{12} consists of variables in H_3 and H_4 that appear in both I_1 and I_2 , and B_{34} similarly consists of variables in H_3 and H_4 that appear in both I_3 and I_4 . Since every variable in H_3 occurs in exactly one of $(I_1 \cap I_2)$ and $(I_3 \cap I_4)$ by definition, we get that

$$\frac{|B_{12}| + |B_{34}|}{2} = |H_4| + \frac{|H_3|}{2}.$$

To compute the righthand side, note that by definition we have the following two relations:

1. Since each term has $4i$ total variables (with repetition):

$$4|H_4| + 3|H_3| + 2|H_2| = 4i$$

2. Since there are ℓ unique variables:

$$|H_4| + |H_3| + |H_2| = \ell$$

Combining these equations gives the desired equality:

$$|H_4| + \frac{|H_3|}{2} = \frac{(4|H_4| + 3|H_3| + 2|H_2|) - 2(|H_4| + |H_3| + |H_2|)}{2} = 2i - \ell.$$

Putting everything together, we finally get

$$\mathbb{E}_{x_1, \dots, x_\ell} [g_i(I_1)g_i(I_2)g_i(I_3)g_i(I_4)] \leq 2^{O(i)} \varepsilon \left(\frac{i}{k}\right)^\ell \|f_i\|^2 + c_{k,i} \varepsilon \gamma^{1/2} \|D_i^k f\|^2,$$

as desired.

9 Characterizing Expansion in HD-walks

One traditional application of hypercontractivity on the discrete hypercube lies in showing that the noisy hypercube graph (given by randomizing each bit of a binary string x with some probability $1 - \rho$) is a small-set expander. This result is also often thought of as stating “sparse functions on the hypercube are noise-sensitive,” an interpretation we’ll discuss in the next section. Unlike the noisy hypercube, it is well known that HD-walks are far from being small set expanders [BHKL20]. Before we quantify this further, let’s recall the definition of (edge) expansion in the general weighted setting.

Definition 9.1 (Weighted Edge Expansion). *Let (X, Π) be a weighted simplicial complex, M a k -dimensional HD-Walk over (X, Π) , and $S \subset X(k)$ a subset. The weighted edge expansion of S is the average probability of leaving S after one step of M :*

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

where $\pi_k|_S$ is the (normalized) restriction of π_k to S ,

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

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

A small-set expander is simply a graph where all small sets expand. To understand why HD-walks fail this condition, let’s consider the Johnson graph. The Johnson graph $J(n, k, \ell)$ is the graph on $\binom{[n]}{k}$ whose edges are given by sets with intersection size ℓ . Well-studied object in their own right, the Johnson graphs are a fundamental example higher order random walks on the complete complex [AJT19]. In our context, we usually think of n as being much larger than k , and ℓ as being (at least) ck for some constant $0 < c < 1$. In this case, one can show by direct computation that the expansion of any i -link X_τ is bounded away from 1:

$$\Phi(X_\tau) \approx 1 - c^{-i},$$

despite the fact that its density is vanishingly small: $\mathbb{E}[1_{X_\tau}] \approx (k/n)^i$.

Recently, BHKL proved a general variant of this result for all HD-walks (see [BHKL20, Theorem 9.2]). They show that spectrum of any k -dimensional walk M on a sufficiently strong local-spectral expander is divided up into $k + 1$ strips of width $O_{k,M}(\gamma)^{17}$ centered around some set of approximate eigenvalues $\{\lambda_i(M)\}_{i=0}^k$, and that the expansion of any link at level i is almost exactly $1 - \lambda_i(M)$. They also prove a weak converse to this result: *any non-expanding set must be concentrated in a link*. It is convenient to state the contrapositive. For any $\delta > 0$, let $R_\delta(M) = r$ denote the number of approximate eigenvalues of M that are greater than δ (a quantity BHKL call the **ST-Rank** of M). BHKL [BHKL20, Theorem 9.5] prove that the expansion of any set $S \subset X(k)$ is at least:

$$\Phi(S) \gtrsim 1 - \delta - c_1 \binom{k}{r} \varepsilon - c_2 \gamma \tag{10}$$

¹⁷BHKL actually only prove the width is $O_{k,M}(\sqrt{\gamma})$, the improvement to $O_{k,M}(\gamma)$ was given soon after by Zhang [Zha20].

where S is (ε, r) -pseudorandom.¹⁸ This is great when $\varepsilon \ll \binom{k}{r}$, but for many applications of interest (e.g. in hardness of approximation), we think of ε as fixed and of k as going to infinity. In this regime, the above characterization is useless, as the bound reduces to the trivial fact $\Phi(S) \geq 0$. Using hypercontractivity, we can completely resolve this issue by offering a variant of Equation (10) with no dependence on k . Before we give the statement, however, we note that both BHKL and our result require the approximate eigenvalues of the HD-walk $\{\lambda_i(M)\}_{i=0}^k$ to decrease monotonically. BHKL proved that this property holds for a broad class of walks they call **complete walks**, which includes all HD-walks of interest studied in the literature.

Definition 9.2 (Complete HD-Walk ([BHKL20] Definition 7.10)). *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) . M is called 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 the d -dimensional complete complex on n_0 vertices.*

All walks we have seen so far (canonical walks, swap walks, pure walks, affine combinations thereof, etc.) are complete, so restricting to this class does not lose much generality. With this in mind, we can finally state our dimension independent bound on the expansion of pseudorandom sets.

Theorem 9.3 (Pseudorandom Sets Expand). *Let (X, Π) be a two-sided γ -local-spectral expander, M a complete k -dimensional HD-walk, and $S \subseteq X(k)$ of density α . Then for any $\delta > 0$ and $r = R_\delta(M) - 1$, the expansion of S is at least:*

$$\Phi(S) \geq 1 - \delta - (1 - \delta)2^{O(r)}\varepsilon^{1/3} - c\gamma$$

where $c \leq 2^{O(k)}w(M)h(M)^2$ and S is (ε, r) -pseudorandom.

The proof of Theorem 9.3 goes through a level- i inequality for pseudorandom functions of independent interest.

Theorem 9.4 (Level- i Inequality). *Let (X, Π) be a γ -local-spectral expander with $\gamma < 2^{-\Omega(k)}$ and $f \in C_k$ a boolean, (ε, i) -pseudorandom function. Then:*

$$\langle f, f_{\uparrow i} \rangle \leq 2^{O(i)}\varepsilon^{1/3}\mathbb{E}[f].$$

Let's first prove Theorem 9.3 given Theorem 9.4.

Proof of Theorem 9.3. The argument is standard and follows from the identity $\Phi(S) = 1 - \frac{1}{\alpha}\langle f, Mf \rangle$ (where $\alpha = \mathbb{E}[f]$), and expanding $f = \sum_{i=0}^k f_{\downarrow i}$ (the HD-Level-Set Decomposition). Namely, since $f_{\downarrow i}$ is an approximate eigenvector, we can write:

$$\begin{aligned} \Phi(S) &= 1 - \frac{1}{\alpha} \sum_{i=0}^k \langle f, Mf_{\downarrow i} \rangle \\ &\geq 1 - \frac{1}{\alpha} \sum_{i=0}^k \lambda_i(M) \langle f, f_{\downarrow i} \rangle - c_1\gamma \end{aligned}$$

¹⁸Note that we have simplified BHKL's result here somewhat for simplicity of presentation, but it is an accurate representation of their result in most cases of interest.

where $c_1 \leq w(M)h(M)^2 2^{O(k)}$. We can now apply Theorem 7.8 to switch between decompositions to get:

$$\Phi(S) \geq 1 - \frac{1}{\alpha} \sum_{i=0}^k \lambda_i(M) \langle f, f_{\uparrow i} \rangle - c_2 \gamma$$

where $c_2 = c_1 + 2^{O(k)}$. Since M is a complete walk, its eigenvalues decay monotonically [BHL20], we can therefore simplify the above to:

$$\begin{aligned} \Phi(S) &\geq 1 - \frac{1}{\alpha} \sum_{i=0}^r \lambda_i(M) \langle f, f_{\uparrow i} \rangle - \frac{\lambda_{r+1}(M)}{\alpha} \sum_{i=r+1}^k \langle f, f_{\uparrow i} \rangle - c_2 \gamma \\ &= 1 - \frac{1}{\alpha} \sum_{i=0}^r \lambda_i(M) \langle f, f_{\uparrow i} \rangle - \frac{\lambda_{r+1}(M)}{\alpha} \left(\alpha - \sum_{i=0}^r \langle f, f_{\uparrow i} \rangle \right) - c_2 \gamma. \end{aligned}$$

Recall that by definition $\lambda_{r+1}(M) \leq \delta$, and hence

$$\begin{aligned} \Phi(S) &\geq 1 - \delta - \frac{1 - \delta}{\alpha} \sum_{i=0}^r \langle f, f_{\uparrow i} \rangle - c_2 \gamma \\ &\geq 1 - \delta - (1 - \delta) 2^{O(r)} \varepsilon^{1/3} - c_2 \gamma \end{aligned}$$

where in the last step we have applied Theorem 9.4. \square

It is left to prove Theorem 9.4, which also follows from fairly standard arguments given Theorem 8.1.

Proof of Theorem 9.4. To simplify notations, we write f_i instead of $f_{\uparrow i}$. Notice that by Hölder's inequality for $p = 4/3, q = 1/4$ we have:

$$\langle f, f_i \rangle \leq \|f\|_{4/3} \|f_i\|_{1/4} = \alpha^{3/4} \mathbb{E}[f_i^4]^{1/4}$$

Combining this with Theorem 8.1 gives the following relation:

$$\frac{\langle f, f_i \rangle^4}{\alpha^3} \leq \mathbb{E}[f_i^4] \leq 2^{O(i)} \varepsilon \mathbb{E}[f_i^2] + c_1 \gamma^{1/2} \alpha \leq 2^{O(i)} \varepsilon \langle f, f_i \rangle + c_2 \gamma^{1/2} \alpha \quad (11)$$

where $c_1, c_2 \leq 2^{O(k)}$ by approximate orthogonality (Lemma 7.6). We can simplify the above via two observations. First, note that we can assume without loss of generality that $\gamma^{1/4} \leq c_1^{-1} \varepsilon$. This follows from observing that:

$$\langle f, f_i \rangle = \binom{k}{i} \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} \langle D_j^k f, D_j^k f \rangle.$$

Appealing to arguments from [BHL20, Lemma 8.8], we have that $\langle D_j^k f, D_j^k f \rangle \leq \varepsilon \alpha$, which gives the naive bound:

$$\langle f, f_i \rangle \leq \binom{k}{i} 2^{O(i)} \varepsilon \alpha.$$

If $\varepsilon \leq \binom{k}{i}^{-3/2}$ then $\binom{k}{i} \varepsilon \leq \varepsilon^{1/3}$ and our desired bound follows. Otherwise, we may assume from now on that $\varepsilon \geq \binom{k}{i}^{-3/2} \geq 2^{-(3/2)k}$. Since $\gamma \leq 2^{-\Omega(k)}$, we are therefore free to assume $\gamma^{1/4} \leq c_1^{-1} \varepsilon$ as well.

Second, we can also assume $\langle f, f_i \rangle \geq \gamma^{1/4}\alpha$, since otherwise we are done by our previous assumptions on γ and ε . Combining these with Equation (11) then gives:

$$\begin{aligned} \frac{\langle f, f_i \rangle^4}{\alpha^3} &\leq 2^{O(i)}\varepsilon\langle f, f_i \rangle + c_1\gamma^{1/2}\alpha \\ &\leq 2^{O(i)}\varepsilon\langle f, f_i \rangle + \varepsilon\gamma^{1/4}\alpha \\ &\leq 2^{O(i)}\varepsilon\langle f, f_i \rangle \end{aligned}$$

as desired. □

10 Fourier Analysis on HDX

In this section we further develop the theory of Fourier analysis on simplicial complexes, and show how hypercontractivity for pseudorandom functions (Theorem 8.1) recovers tight analogs of the KKL Theorem and noise-sensitivity of sparse functions. This requires introducing a number of new analog definitions of classic Fourier analytic quantities on simplicial complexes. To get an idea for what these should look like, it will be useful to start by considering a natural embedding of the hypercube itself into a simplicial complex.

Definition 10.1 (Hypercube Complex). *The hypercube complex $X = X_{\{0,1\}^n}$ is the complete n -partite complex on $X(1) = [n] \times \{0, 1\}$, where the first coordinate denotes the color of the vertex. That is, the top level faces are $X(n) = \{(1, x_1), \dots, (n, x_n) : x \in \{0, 1\}^n\}$.*

We make a few notes on this definition. First, it is clear from definition that $X(n)$ can equivalently be thought of as the hypercube $\{0, 1\}^n$, where each color in $[n]$ corresponds to a coordinate in $\{0, 1\}^n$. Further, classic graphs on $\{0, 1\}^n$ such as the hypercube or noisy hypercube can be expressed as simple higher order random walks. The hypercube graph, for instance, is simply the non-lazy lower walk $UD^+ = 2U_{n-1}D_n - I$. This embedding will serve as our guiding principle for developing analog Fourier-analytic definitions on simplicial complexes—whenever possible, our definitions will reduce to the standard notion when applied to the hypercube complex. We note that the same embedding can be used for any product distribution and all of our definitions will generalize appropriately. We focus on the simple case of the hypercube for ease of exposition.

10.1 Total Influence and the KKL Theorem

We'll start with a fundamental concept in classical Fourier analysis, *influence*. Let's first recall the definition of (total) influence on the discrete hypercube. Influence can be formalized in a number of equivalent ways. It is often thought of, for instance, as a measure of average sensitivity. In our context, it will be most convenient to view influence as a statement about the expansion of a function with respect to the hypercube graph. More formally, let Q_n denote the normalized adjacency matrix of the hypercube graph, and $Q_n^{\text{lazy}} = \frac{I+Q_n}{2}$ its lazy variant. We will write total influence in terms of the (un-normalized) Laplacian operator $L = n(I - Q_n^{\text{lazy}})$.

Definition 10.2 (Total Influence (hypercube)). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. The total influence of f , denoted $I[f]$, is:*

$$I[f] = \langle f, Lf \rangle.$$

Expressed in this sense, there is a natural generalization to simplicial complexes. It is not hard to see that on the hypercube complex, Q_n^{lazy} is exactly the lower walk UD . As a result, we'll define influence using the Laplacian of the lower walk.

Definition 10.3 (Total Influence). *Let (X, Π) be a pure, weighed simplicial complex and $f \in C_k$. The influence of f , denoted $I[f]$ is:*

$$I_{(X, \Pi)}[f] = \langle f, L_{UD}f \rangle$$

where $L_{UD} = k(I - U_{k-1}D_k)$. When clear from context, we will simply write $I[f]$.

When (X, Π) is sufficiently expanding, Definition 10.3 acts much like standard influence on the cube. For instance, recalling standard bounds on the spectral expansion of L_{UD} [KO20], it is not hard to see the total influence of any function on a γ -local-spectral expander lies between $(1 + O_k(\gamma))\text{Var}(f) \leq I_{(X, \Pi)}[f] \leq k\text{Var}(f)$, which returns the standard bounds as γ goes to 0. Similarly, it is obvious that the total influence of any function on the hypercube complex is equivalent to its total influence on the hypercube, as the lower walk $U_{n-1}D_n$ is exactly Q_n^{lazy} .

Observation 10.4. *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be any function and $f_X : X_{\{0,1\}^n}(n) \rightarrow \mathbb{R}$ its equivalent on the hypercube complex, that is:*

$$f(x_1, \dots, x_n) = f_X((1, x_1), \dots, (n, x_n))$$

Then:

$$I_{X_{\{0,1\}^n}}[f] = I[f].$$

One of the most well-studied problems in the analysis of boolean functions is understanding the structure of functions with low influence. The seminal result in this area is called the ‘‘KKL Theorem’’ [KKL88]. Informally, the KKL Theorem states that if a function has low total influence, there must exist an *influential coordinate* (in the sense that on average over $\{0, 1\}^n$, the coordinate has a large affect on the value of f). Morally, this can also be thought of as strong notion of the following statement: ‘‘functions with low influence are not pseudorandom.’’ While the KKL Theorem itself does not extend beyond the hypercube, this latter interpretation does. In particular, Bourgain [FB99] proved a similar statement over any product space: functions with low influence must have some influential *set* of coordinates, and are therefore not pseudorandom. We prove a variant of Bourgain’s result for local-spectral expanders.

Theorem 10.5 (Bourgain’s Theorem for HDX). *Let (X, Π) be a two-sided γ -local-spectral expander with $\gamma \leq 2^{-\Omega(k)}$ and $f \in C_k$ a boolean function. Then for any $0 \leq K \leq k$, if $I[f] \leq K\text{Var}(f)$, there exists an $(i \leq K)$ -link τ with large density:*

$$\mathbb{E}_{X_\tau}[f] \geq 2^{-O(K)}.$$

Proof. This follows without too much difficulty from the expansion of pseudorandom sets (Theorem 9.3). In particular, notice that our assumption on the influence implies the following bound on the expansion of f with respect to the lower walk $U_{k-1}D_k$:

$$\Phi(f) = \frac{\langle f, L_{UD}f \rangle}{k\mathbb{E}[f]} = \frac{I[f](1 - \mathbb{E}[f])}{k\text{Var}(f)} \leq \frac{K}{k}.$$

Recall that Theorem 9.3 states that for any $\delta > 0$ and $r = R_\delta(UD) - 1$, the expansion of an (ε, r) -pseudorandom boolean function g with respect to the lower walk is at least:

$$\Phi(g) \geq 1 - \delta - (1 - \delta)2^{O(r)}\varepsilon^{1/3} - c\gamma.$$

Using this fact, we’ll show that f cannot be (ε, K) -pseudorandom for $\varepsilon \leq 2^{-\Omega(K)}$, which gives the result.

To this end, assume f is (ε, K) -pseudorandom for some $\varepsilon = 2^{-\Omega(K)}$ to be determined soon (else we are done), and let δ be $1 - \frac{K(1+\varepsilon^{1/6})}{k}$ such that $1 - \delta < (K + 1)/k$. Since the eigenvalues of UD are concentrated

around $1, 1 - 1/k, 1 - 2/k, \dots, 1 - K/k$ for small enough γ [BHKL20], the ST-Rank $R_\delta(UD) = K + 1$, and $r = K$. Theorem 9.3 then implies:

$$\Phi(f) \geq \frac{K}{k} \cdot (1 + \varepsilon^{1/6})(1 - 2^{O(r)}\varepsilon^{1/3}),$$

where we have again used our assumption on the size of γ . Re-arranging the above using the upper bound on expansion then gives a lower bound on ε of:

$$\varepsilon^{1/3} \geq \frac{\varepsilon^{1/6}}{1 + \varepsilon^{1/6}} \cdot \frac{1}{2^{O(r)}} \geq \frac{1}{2^{O(r)}},$$

which implies the result. □

Before moving on, we'll prove that this result is tight.

Proposition 10.6. *Let $c \geq 1$ be any constant. Then for all integers $K, k > 1$ satisfying $k \geq \Omega_c(K)$ and any n sufficiently larger than k , there exists a Boolean function $f \in C_k$ on the k -dimensional complete complex on n vertices satisfying:*

1. *The influence of f is small:*

$$I[f] \leq K \text{Var}(f)$$

2. *For every $i \leq cK$, all i -links are sparse:*

$$\forall i \leq cK, \tau \in X(i) : \mathbb{E}_{X_\tau}[f] \leq 2^{-\Omega(K)}.$$

Proof. Our construction is based on a careful analysis of the anti-tribes function (a.k.a the AND of ORs function) similar to [KLLM19, Example 5.8]. Concretely, let T_1, \dots, T_m (called “tribes”) be $m = 2cK$ disjoint sets of $c_1 \frac{n}{k}$ vertices for some $c_1 \geq \log(\Omega(c))$. We define our candidate function $f \in C_k$ to be 1 on a k -face S exactly when S contains some vertex from each T_i :

$$f(S) = \begin{cases} 1 & \text{if } \forall 1 \leq i \leq m: |S \cap T_i| > 0 \\ 0 & \text{else.} \end{cases}$$

For simplicity, it will actually be more convenient to analyze f as a function over $[n]^k$ rather than $X(k) = \binom{[n]}{k}$. Since the probability of repeated vertices in the former is $o_{n,k}(1)$, this has no effect on our final result when n is sufficiently larger than k .

Let's start by proving the density of every cK -link is at most $2^{-\Omega(k)}$.¹⁹ Note that this implies the same for every i -link for $i \leq cK$. It is not hard to see that the largest density link comes from fixing an element in each of cK tribes. For simplicity, fix such a cK -link T with a vertex in each T_i for $cK + 1 \leq i \leq m$ (all such links are symmetric, so it suffices to analyze this case). For a uniformly drawn element $S \in [n]^k$, let E_i denote the event that S contains a vertex in T_i . We'd like to bound:

$$\mathbb{E}_{X_T}[f] = \Pr_{S \sim [n]^k} \left[\bigcap_{i=1}^{cK} E_i \mid S \supset T \right] = \Pr_{S \sim [n]^{k-cK}} \left[\bigcap_{i=1}^{cK} E_i \right],$$

¹⁹Formally we note this should really be shown for $\lfloor cK \rfloor$ -links, but this makes no significant difference in the analysis so we ignore it for simplicity.

where we have used the fact that $S \setminus T$ is independent of T since we are working over $[n]^k$. Since the E_i are negatively correlated, we can bound this probability by:

$$\begin{aligned} \Pr_{S \sim [n]^{k-cK}} \left[\bigcap_{i=1}^{cK} E_i \right] &\leq \prod_{i=1}^{cK} \Pr_{S \sim [n]^{k-cK}} [E_i] \\ &\leq (1 - (1 - c_1/k)^k)^{cK} \\ &\leq \left(1 - \frac{1}{O(c)}\right)^{cK} \\ &\leq 2^{-\Omega(K)} \end{aligned}$$

where we've used the fact that $e^{-x} \geq 1 - x \geq e^{-x/(1-x)}$ for $x < 1$ and our assumptions on the size of k .

We now move on to analyzing the influence of f , which will follow from similar computations. To start, notice that it is instead sufficient to bound the expansion of f with respect to the lower walk by:

$$\Phi(f) \leq \frac{K \operatorname{Var}(f)}{k \mathbb{E}[f]} = \frac{K(1 - \mathbb{E}[f])}{k},$$

as then:

$$I[f] = \langle f, L_{UD} f \rangle = k \Phi(f) \mathbb{E}[f] \leq K \operatorname{Var}(f),$$

where we recall L_{UD} is the un-normalized Laplacian of UD .

To this end, recall that the expansion of f can also be defined as the average probability of leaving $\operatorname{supp}(f)$ after applying the walk, that is:

$$\Phi(f) = \mathbb{E}_{S \sim \operatorname{supp}(f)} [\phi(S)],$$

where $\phi(S)$ denotes the probability of leaving S in a single step of the lower walk. To compute this value, recall that in the down step of the walk, a uniformly random vertex is removed from S . In order to leave the support of f in the up step, the removed element must have been selected from a tribe T_i such that $|S \cap T_i| = 1$. The idea is then to show that for most samples, only a small fraction of tribes have exactly one element. With this in mind, let B_i be the event $|S \cap T_i| = 1$ over the randomness of $S \sim \operatorname{supp}(f)$. Formalizing the above argument, we can bound $\phi(S)$ by the sum over B_i :

$$\phi(S) \leq \sum_{i=1}^m \frac{B_i(S)}{k},$$

and therefore the expansion $\Phi(f)$ by:

$$\Phi(f) \leq \frac{1}{k} \mathbb{E}_{S \sim \operatorname{supp}(f)} [B_i(S)].$$

By a similar argument to our density calculations, the probability that any fixed tribe T_i has exactly one element from $S \sim \operatorname{supp}(f)$ is at most:

$$\begin{aligned} \mathbb{E}[B_i] &= \left(1 - \frac{c_1}{k}\right)^{k-m} \\ &\leq e^{-c_1 \frac{k-m}{k}} \\ &\leq \frac{1}{\Omega(c)} \end{aligned}$$

since we have by assumption that k is much larger than m . Plugging this into our expression for expansion then gives:

$$\Phi(f) \leq \frac{m}{k} \cdot \frac{1}{\Omega(c)} \leq c_2 \frac{K}{k}$$

for some $c_2 < 1$. Noting that $\mathbb{E}[f] = 2^{-\Omega(K)}$ then implies the result for the appropriate setting of constants. \square

10.2 Stability and the Noise Operator

Another fundamental notion in boolean Fourier analysis is the *noise operator* T_ρ . It is convenient to express the definition in terms of the following process on an element $x \in \{0, 1\}^n$:

1. Remove each bit with probability $1 - \rho$.
2. Replace each removed bit uniformly²⁰ at random.

We write the distribution over y given by this process as $N_\rho(x)$. The noise operator T_ρ is simply the averaging operator over ρ -correlated strings.

Definition 10.7 (Noise Operator (Hypercube)). *Let $f : \{0, 1\} \rightarrow \mathbb{R}$ be any function. The noise operator T_ρ averages f over N_ρ :*

$$T_\rho f(x) = \mathbb{E}_{y \sim N_\rho(x)} [f(y)].$$

Extending the noise operator to simplicial complexes is a bit tricky naively since there is no notion of coordinates. To do this, consider the following reformulation of the distribution $N_\rho(x)$, instead of removing each coordinate independently with probability $1 - \rho$, we remove a uniformly random *set of i coordinates* with probability $\binom{n}{i} \rho^{n-i} (1 - \rho)^i$, and replace them uniformly at random. This equivalent process does have a natural analog on simplicial complexes: simply replace “uniformly random set of i coordinates” with “uniformly random i -face.” We can formalize this through the averaging operators.

Definition 10.8 (Noise Operator (Simplicial Complex)). *Let (X, Π) be a pure, weighted simplicial complex. The noise operator $T_\rho^k(X, \Pi)$ at level k of the complex is:*

$$T_\rho^k(X, \Pi) = \sum_{i=0}^k \binom{k}{i} (1 - \rho)^i \rho^{k-i} U_{k-i}^k D_{k-i}^k.$$

We write T_ρ when the level and complex are clear from context.

Let’s take a moment to check that, as with influence, when applied to the hypercube complex this definition recovers T_ρ .

Observation 10.9. *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be any function and $f_X : X_{\{0,1\}^n}(n) \rightarrow \mathbb{R}$ its equivalent on the hypercube complex, then:*

$$T_\rho^n(X_{\{0,1\}^n}) f_X = T_\rho f.$$

Proof. $T_\rho^n(X_{\{0,1\}^n})$ is also an averaging operator, so it is enough to confirm it averages over the ρ -noisy distribution N_ρ . We claim this is clear from definition. In particular, notice that $U_{n-i}^n D_{n-i}^n$ on $X_{\{0,1\}^n}$ is exactly the process of removing i -coordinates uniformly at random, and replacing them with uniformly random bits. As we mentioned above, this is an equivalent way to define $N_\rho(x)$, is applying this process with probability $\binom{n}{i} (1 - \rho)^i \rho^{n-i}$, which exactly matches the definition of $T_\rho(X_{\{0,1\}^n})$. \square

²⁰In more general settings like the p -biased cube, this is replaced with respect to the underlying distribution.

The noise operator has a wide variety of applications across boolean Fourier analysis. One classical application is to analyze the *noise-sensitivity* of a boolean function, that is the likelihood that the function flips on a noisy input. It is convenient to define the opposite concept first, *stability*.

Definition 10.10 (Stability (Hypercube)). *Let $f : \{0, 1\} \rightarrow \mathbb{R}$ be any function. The stability of f with respect to ρ , denoted $Stab_\rho(f)$, is:*

$$Stab_\rho(f) = \langle f, T_\rho f \rangle.$$

Since we already defined T_ρ on simplicial complexes, stability has an obvious analog.

Definition 10.11 (Stability (Simplicial Complex)). *Let (X, Π) be a weighted, pure simplicial complex and $f \in C_k$. The noise stability of f with respect to ρ , denoted $Stab_\rho(f)$, is:*

$$Stab_\rho^{(X, \Pi)}(f) = \langle f, T_\rho^k(X, \Pi)f \rangle.$$

We drop (X, Π) from the notation when clear from context.

Similarly, it is clear that our definition of stability for complexes returns the original definition when applied to the hypercube complex.

Observation 10.12. *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be any function and $f_X : X_{\{0,1\}^n}(n) \rightarrow \mathbb{R}$ its equivalent on the hypercube complex, then:*

$$Stab_\rho(f) = Stab_\rho^{X_{\{0,1\}^n}}(f_X)$$

A function is called *noise-sensitive* if it has poor stability. One classical result in boolean Fourier analysis is that sparse functions on the hypercube are noise-sensitive, which is equivalent to saying that the noisy hypercube graph is a small-set expander. Since the noise operator is just a specific instance of a (complete) higher order random walk, Theorem 9.3 implies an analogous statement for functions on HDX: pseudorandom functions are noise-sensitive.

Corollary 10.13 (Pseudorandom functions are noise sensitive). *Let (X, Π) be a two-sided γ -local-spectral expander, $f \in C_k$ an (r, δ) -pseudorandom boolean function for $r = \log(2/\varepsilon)/\log(1/\rho) + 2$ and $\delta \leq 2^{-\Omega(r)}\varepsilon^3$. Then f is noise sensitive:*

$$Stab_\rho(f) \leq (\varepsilon + c\gamma)\mathbb{E}[f]$$

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

Proof. One can directly compute from [BHKL20, Corollary 7.6] that the approximate eigenvalues of T_ρ are exactly $\lambda_i = \rho^i$. As a result, for small enough γ , the $(\varepsilon/2)$ -ST-Rank of T_ρ is at most:

$$R_{\varepsilon/2}(T_\rho) \leq \log(\varepsilon/2)/\log(1/\rho) + 2.$$

Since T_ρ is a higher order random walk, Theorem 9.3 states that the non-expansion of any (δ, r) -pseudorandom function f of density α is at most:

$$\begin{aligned} \frac{1}{\alpha} \langle f, T_\rho f \rangle &\leq \alpha + (1 - \alpha)\varepsilon/2 + 2^{O(r)}\delta + c\gamma \\ &\leq \alpha + (1 - \alpha)\varepsilon/2 + \varepsilon/4 + c\gamma \\ &\leq \varepsilon + c\gamma, \end{aligned}$$

where we've used the fact that $\alpha \leq \delta \leq \varepsilon/4$. □

The noise operator is actually also commonly used to define hypercontractivity. In this form, the standard hypercontractive inequality generally states:

$$\|T_\rho f\|_4 \leq \|f\|_2$$

for some $\rho = \Theta(1)$. It is well known that on the hypercube this statement is in fact equivalent to Bonami's lemma. We can show a similar equivalence between our variant of Bonami's lemma (Theorem 8.1) and a noise operator based form of hypercontractivity for pseudorandom functions. To state the strongest form of the result, it will be useful to extend the classic notion of *degree* to simplicial complexes.

Definition 10.14 (Function Degree). *Let (X, Π) be a pure, weighted simplicial complex, and $f \in C_k$ any function. The degree of f , denoted $\deg(f)$, is the largest i such that $f_{\uparrow i}$ is non-zero.*

We now show how to translate Theorem 8.1 into noise operator form for degree i , (ε, i) -pseudorandom functions.

Proposition 10.15. *Let (X, Π) be a two-sided γ -local-spectral expander satisfying $\gamma \leq 2^{-\Omega(k)}$ and $f \in C_k$ a degree i ,²¹ (ε, i) -pseudorandom function. Then for some constant $\rho = \Theta(1)$, we have:*

$$\|T_\rho f\|_4^4 \leq \varepsilon \|f\|_2^2 \|f\|_\infty^2.$$

Proof. The overall proof follows a fairly standard reduction from hypercontractivity to Bonami's lemma (see e.g. [O'D14, Exercise 9.6]), but requires some extra work due to the fact that the Bottom-Up decomposition is only approximately an eigenbasis for T_ρ (in an ℓ_2 -sense). Namely, by [BHKL20, Proposition 7.5] and Lemma 7.5, we can write:

$$T_\rho f_{\uparrow j} = \rho^k f_{\uparrow j} + \text{err}_j$$

where $\|\text{err}_j\|_2 \leq k^{O(j)} \gamma \|f\|_2$, and $\|\text{err}_j\|_\infty \leq k^{O(j)} \varepsilon \|f\|_\infty$. The last of these facts is slightly less standard, and follows from noting that err_j is really a linear combination of at most $k^{O(j)}$ averaging operators applied to g_j (see [BHKL20, Proposition 7.5]), and that $\|g_j\|_\infty \leq 2^{O(j)} \varepsilon \|f\|_\infty$. With this in mind, we can expand out $\|T_\rho f\|$ by the Bottom-Up Decomposition and apply Theorem 8.1 to get:

$$\begin{aligned} \|T_\rho f\|_4 &\leq \sum_{j=0}^i \|T_\rho f_{\uparrow j}\|_4 \\ &\leq \sum_{j=0}^i \rho^j \|f_{\uparrow j}\|_4 + \|\text{err}_j\|_4 \\ &\leq \frac{1}{2} \varepsilon^{1/4} \|f\|_2^{1/2} \|f\|_\infty^{1/2} + \sum_{j=0}^i \|\text{err}_j\|_4 \\ &\leq \frac{1}{2} \varepsilon^{1/4} \|f\|_2^{1/2} \|f\|_\infty^{1/2} + \sum_{j=0}^i \|\text{err}_j\|_2^{1/2} \|\text{err}_j\|_\infty^{1/2} \\ &\leq \varepsilon^{1/4} \|f\|_2^{1/2} \|f\|_\infty^{1/2} \end{aligned}$$

where we have assumed that ρ is a sufficiently small constant. Taking the fourth power of both sides completes the proof. \square

²¹When $i \ll k$, we can replace the condition on γ with $\gamma \leq k^{-\Omega(i)}$

References

- [AJK⁺21] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong, *Entropic independence in high-dimensional expanders: Modified log-sobolev inequalities for fractionally log-concave polynomials and the ising model*, arXiv preprint arXiv:2106.04105 (2021).
- [AJT19] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani, *Approximating constraint satisfaction problems on high-dimensional expanders*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 180–201.
- [AL20] Vedat Levi Alev and Lap Chi Lau, *Improved analysis of higher order random walks and applications*, arXiv preprint arXiv:2001.02827 (2020).
- [ALO20] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, arXiv preprint arXiv:2001.00303 (2020).
- [ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant, *Log-concave polynomials ii: high-dimensional walks and an fpras for counting bases of a matroid*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 1–12.
- [BCC⁺21] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda, *On mixing of markov chains: Coupling, spectral independence, and entropy factorization*, arXiv preprint arXiv:2103.07459 (2021).
- [Bec75] William Beckner, *Inequalities in fourier analysis*, Annals of Mathematics **102** (1975), no. 1, 159–182.
- [BGH⁺12] Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer, *Making the long code shorter*, 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, IEEE, 2012, pp. 370–379.
- [BHKL20] Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett, *High dimensional expanders: eigenstripping, pseudorandomness, and unique games*, arXiv preprint arXiv:2011.04658 (2020).
- [BKK⁺92] Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial, *The influence of variables in product spaces*, Israel Journal of Mathematics **77** (1992), no. 1, 55–64.
- [BKS18] Boaz Barak, Pravesh K Kothari, and David Steurer, *Small-set expansion in shortcode graph and the 2-to-2 conjecture*, arXiv preprint arXiv:1804.08662 (2018).
- [Bon70] Aline Bonami, *Étude des coefficients de fourier des fonctions de $L^p(G)$* , Annales de l’institut Fourier, vol. 20, 1970, pp. 335–402.
- [CGŠV21] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda, *Rapid mixing for colorings via spectral independence*, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 1548–1557.
- [CLV20] Zongchen Chen, Kuikui Liu, and Eric Vigoda, *Rapid mixing of glauber dynamics up to uniqueness via contraction*, arXiv preprint arXiv:2004.09083 (2020).
- [CLV21] ———, *Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1537–1550.

- [DD19] Yotam Dikstein and Irit Dinur, *Agreement testing theorems on layered set systems*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 1495–1524.
- [DDFH18] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha, *Boolean function analysis on high-dimensional expanders*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [DDHRZ20] Yotam Dikstein, Irit Dinur, Prahladh Harsha, and Noga Ron-Zewi, *Locally testable codes via high-dimensional expanders*, arXiv preprint arXiv:2005.01045 (2020).
- [DEL⁺21] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes, *Locally testable codes with constant rate, distance, and locality*, arXiv preprint arXiv:2111.04808 (2021).
- [DHK⁺19] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta Shma, *List decoding with double samplers*, Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2019, pp. 2134–2153.
- [DK17] Irit Dinur and Tali Kaufman, *High dimensional expanders imply agreement expanders*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 974–985.
- [DKK⁺18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra, *On non-optimally expanding sets in grassmann graphs*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 940–951.
- [DKK⁺18b] ———, *Towards a proof of the 2-to-1 games conjecture?*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 376–389.
- [DSC96] Persi Diaconis and Laurent Saloff-Coste, *Logarithmic sobolev inequalities for finite markov chains*, The Annals of Applied Probability **6** (1996), no. 3, 695–750.
- [EKZ20] Shai Evra, Tali Kaufman, and Gilles Zémor, *Decodable quantum LDPC codes beyond the square root distance barrier using high dimensional expanders*, 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020, 2020, pp. 218–227.
- [FB99] Ehud Friedgut and Jean Bourgain, *Sharp thresholds of graph properties, and the k-sat problem*, Journal of the American mathematical Society **12** (1999), no. 4, 1017–1054.
- [FGYZ21] Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang, *Rapid mixing from spectral independence beyond the boolean domain*, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 1558–1577.
- [FK96] Ehud Friedgut and Gil Kalai, *Every monotone graph property has a sharp threshold*, Proceedings of the American mathematical Society **124** (1996), no. 10, 2993–3002.
- [FKLM20] Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer, *Hypercontractivity on the symmetric group*, arXiv preprint arXiv:2009.05503 (2020).
- [FKN02] Ehud Friedgut, Gil Kalai, and Assaf Naor, *Boolean functions whose fourier transform is concentrated on the first two levels*, Advances in Applied Mathematics **29** (2002), no. 3, 427–437.

- [FOW18] Yuval Filmus, Ryan O’Donnell, and Xinyu Wu, *A log-sobolev inequality for the multislice, with applications*, arXiv preprint arXiv:1809.03546 (2018).
- [Fri98] Ehud Friedgut, *Boolean functions with low average sensitivity depend on few coordinates*, *Combinatorica* **18** (1998), no. 1, 27–35.
- [GLL21] Tom Gur, Noam Lifshitz, and Siqi Liu, Personal communication, June 2021.
- [Gro75] Leonard Gross, *Logarithmic sobolev inequalities*, *American Journal of Mathematics* **97** (1975), no. 4, 1061–1083.
- [JPV21] Vishesh Jain, Huy Tuan Pham, and Thuy Duong Vuong, *Spectral independence, coupling with the stationary distribution, and the spectral gap of the glauber dynamics*, arXiv preprint arXiv:2105.01201 (2021).
- [JQST20] Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani, *Unique decoding of explicit epsilon-balanced codes near the gilbert-varshamov bound*, arXiv preprint arXiv:2011.05500 (2020).
- [JST21] Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani, *Near-linear time decoding of ta-shma’s codes via splittable regularity*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1527–1536.
- [KKL88] J. Kahn, G. Kalai, and N. Linial, *The influence of variables on boolean functions*, 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (Los Alamitos, CA, USA), IEEE Computer Society, oct 1988, pp. 68–80.
- [KLLM19] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer, *Hypercontractivity for global functions and sharp thresholds*, arXiv preprint arXiv:1906.05568 (2019).
- [KLLM21] ———, *Global hypercontractivity and its applications*, arXiv preprint arXiv:2103.04604 (2021).
- [KM16] Tali Kaufman and David Mass, *High dimensional combinatorial random walks and colorful expansion*, arXiv preprint arXiv:1604.02947 (2016).
- [KM20] ———, *Local-to-global agreement expansion via the variance method*, 11th Innovations in Theoretical Computer Science Conference (ITCS 2020), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [KMMS18] Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra, *Small set expansion in the johnson graph.*, *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 25, 2018, p. 78.
- [KMS17] Subhash Khot, Dor Minzer, and Muli Safra, *On independent sets, 2-to-2 games, and grassmann graphs*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 576–589.
- [KMS18] ———, *Pseudorandom sets in grassmann graph have near-perfect expansion*, 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2018, pp. 592–601.
- [KO18] Tali Kaufman and Izhar Oppenheim, *Construction of new local spectral high dimensional expanders*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 773–786.

- [KO20] ———, *High order random walks: Beyond spectral gap*, *Combinatorica* (2020), 1–37.
- [KO21] Tali Kaufman and Izhar Oppenheim, *High dimensional expansion implies amplified local testability*, *CoRR* (2021).
- [KS20] Tali Kaufman and Ella Sharakanski, *Chernoff bound for high-dimensional expanders*, To Appear *APPROX/RANDOM 2020*, 2020.
- [KT21a] Tali Kaufman and Ran J Tessler, *Local to global high dimensional expansion and garland’s method for general posets*, arXiv preprint arXiv:2101.12621 (2021).
- [KT21b] Tali Kaufman and Ran J. Tessler, *New cosystolic expanders from tensors imply explicit quantum LDPC codes with $\Omega(\sqrt{n} \log^k n)$ distance*, *STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, Virtual Event, Italy, June 21-25, 2021, 2021, pp. 1317–1329.
- [Liu21] Kuikui Liu, *From coupling to spectral independence and blackbox comparison with the down-up walk*, arXiv preprint arXiv:2103.11609 (2021).
- [LM19] Noam Lifshitz and Dor Minzer, *Noise sensitivity on the p -biased hypercube*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 1205–1226.
- [LSV05] Alexander Lubotzky, Beth Samuels, and Uzi Vishne, *Explicit constructions of ramanujan complexes of type ad* , *European Journal of Combinatorics* **26** (2005), no. 6, 965–993.
- [MOO05] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz, *Noise stability of functions with low influences: invariance and optimality*, 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05), IEEE, 2005, pp. 21–30.
- [O’D14] Ryan O’Donnell, *Analysis of boolean functions*, Cambridge University Press, 2014.
- [Opp18] Izhar Oppenheim, *Local spectral expansion approach to high dimensional expanders part i: Descent of spectral gaps*, *Discrete & Computational Geometry* **59** (2018), no. 2, 293–330.
- [Tal94] Michel Talagrand, *On russo’s approximate zero-one law*, *The Annals of Probability* (1994), 1576–1587.
- [Zha20] Ruizhe Zhang, Personal communication, December 2020.