Hypercontractivity on High Dimensional Expanders

Mitali Bafna
Harvard University
Cambridge, MA, USA

Max Hopkins
UC San Diego
La Jolla, CA, USA

Tali Kaufman
Bar-Ilan University
Ramat-Gan, Israel

Shachar Lovett
UC San Diego
La Jolla, CA, USA

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.

This is an extended abstract. The full paper may be found at https://arxiv.org/abs/2111.09444.

CCS CONCEPTS

• **Theory of computation → Expander graphs and randomness extractors**; **Pseudorandomness and derandomization**.

KEYWORDS

High Dimensional Expanders, Hypercontractivity, Small-Set Expansion

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06.
https://doi.org/10.1145/3519935.3520040

ACM Reference Format:
Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. 2022. Hypercontractivity on High Dimensional Expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3519935.3520040

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 [40] and Majority is Stablest [58], the study of hypercontractivity has since seen a resurgence on extended domains such as the p-biased cube [50], slice [52], and Grassmannian [54]. 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 [54]. 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 and qLDPC codes [21, 60], and efficient approximate sampling for many important systems (e.g. for matroid bases [6], independent sets [5], Ising models [4], 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 [54]. 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 $f : \{0,1\}^n \to \mathbb{R}$ 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* $f : \{0,1\}^n \to \{0,1\}$. 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 [12] (and later independently by Beckner [10] and Gross [34]), 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 ith level of a boolean function f, denoted f_i, should satisfy:

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

This deceptively simple observation, known in the above form as "Bonami's Lemma" [12], led to many landmark results including the KKL Theorem [40], noise-sensitivity of sparse functions [40], Friedgut's Junta Theorem [30], and Majority is Stablest [58]. 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 [13, 30, 32, 62], but it wasn't until recently that tight analogues of Equation (1) were developed for general product spaces [50] (generalizing work of Friedgut and Bourgain [31] and Hatami [36]) as well as for other structured domains such as the symmetric group [28] and Grassmannian [54]. 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 analogues of classical results, a variety of sharp threshold theorems [31, 50, 51, 55], and perhaps most famously the proof of the 2-2 Games Conjecture [9, 24, 25, 52–54]. 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 [20–22, 26, 38, 39, 46, 49, 60] and approximate sampling [3, 5, 6, 11, 14–16, 27, 37, 56], but also in agreement testing [18, 23, 42], CSP-approximation [2, 7], and (implicitly) hardness of approximation [52, 54]. In this work, we study a central notion of high dimensional expansion called *two-sided local-spectral expansion*, originally developed by Dinur and Kaufman [23] 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 due to Dikstein, Dinur, Filmus, and Harsha (DDFH) [19], and Kaufman and Oppenheim (KO) [45]. 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 [19, 33], efficient CSP-approximation [2, 7]), 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 [54]. On the other hand, Bafna, Hopkins, Kaufman, and Lovett (BHKL) [7] 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 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

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

[2] The 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.

[3] We 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.

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 analogues 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 [7], *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 [54], and opens yet another avenue towards the use of HDX in hardness of approximation. We also introduce natural analogues[5] 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 [44, 57], 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 [19]. While these classes of local-spectral expanders are essentially equivalent in an ℓ_2 sense [7, 19, 48], 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.

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

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

[6]Formally, it is more accurate to say 'locally sparse' or 'bounded degree' here. While previous settings such as highly imbalanced products may be sparse in the sense that most of their weight is concentrated on relatively few faces, they are not sparse in the much stronger sense of a bounded-degree HDX. The former, for instance, will always have some very dense restrictions, whereas every vertex in the latter sees only a tiny fraction of the full complex.

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 [23] 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**:

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

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 [41] 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 $\{f : X(k) \to \mathbb{R}\}$ 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:

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

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 $D^k_i = D_i \circ \ldots \circ D_k$ and $U^k_i = U_k \circ \ldots \circ U_i$. Informally, HD-walks are simply affine combinations of composed averaging operators. For instance, the basic composition $N^i_k = U^{k+i}_k D^{k+i}_k$, 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. Formal versions and proofs of all results are available in the full version of the paper at https://arxiv.org/abs/2111.09444.

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

[7]A graph is a γ-spectral expander if the second largest eigenvalue of its weighted adjacency matrix (also called the random walk matrix) is at most γ in absolute value.

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.

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 [19, 45] 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 [2, 7, 19], 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.

Definition 3.1 (Bottom-Up Decomposition). *Let X be a d-dimensional pure simplicial complex and $f \in C_k$ any function. For all $0 \leq i \leq k$ and $\tau \in X(i)$, define the ith level function(s) to be:*

$$g_{\uparrow i}(\tau) = \sum_{\sigma \subseteq \tau} (-1)^{|\tau \setminus \sigma|} \mathop{\mathbb{E}}_{X_\sigma}[f], \quad f_{\uparrow i} = \binom{k}{i} U_i^k g_{\uparrow i}.$$

One can check that $f = \sum_{i=0}^{k} f_{\uparrow i}$.

Here, $g_{\uparrow i}(\tau)$ should be thought of as the contribution to f coming from τ (where contributions from $\sigma \subsetneq \tau$ have been removed by inclusion/exclusion). The Fourier level $f_{\uparrow i}$ is then defined by summing over these contributions. It is worth noting that the Bottom-Up Decomposition also has a simple recursive form:

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

In fact, it should be noted that while the consideration of this basis is new over general simplicial complexes, the above recursive form was first studied for the special case of the complete complex by [52]. 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). *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 similar to an analogous result for the HD-Level-Set Decomposition in [19, Theorem 1.3]. 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 close in ℓ_2-norm.

Theorem 3.3 (Bottom-Up Approximates HD-Level-Set). *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:*

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

Similarly,

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

The main advantage of the Bottom-Up Decomposition then lies in its simple explicit and recursive forms. In the full paper, we show how these properties are useful for analyzing finer-grained structure 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 [40], FKN [33], Majority is Stablest [58], sharp threshold theorems [31], 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 [12], 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 ith Fourier level of a boolean function. The corresponding statement in our context is therefore that the ith level of the Bottom-Up Decomposition should satisfy an analogous inequality:

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

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 [50], where **pseudorandom** functions satisfy a form of Equation (2).

Definition 3.4 (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 $\tau \in X(i)$:

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

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

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

While the use of $\|f\|_\infty$ here may initially seem unnatural, it is in fact the appropriate scaling factor on a bounded-degree complex (at least up to constants). Namely since restrictions are of constant size, doubling the largest value in f leads to a $(1+\delta)$ multiplicative increase in density on links including that face for some constant $\delta > 0$.

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 3.4 are also sometimes called *global* since they are not concentrated in any local structure [50, 55]. We call them pseudorandom in keeping with prior literature on the Johnson and Grassmann graphs [52, 54], and because they cannot be distinguished from an (ε-sparse) random function by examining density inside links. Finally, note that Definition 3.4 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. Keevash, Lifshitz, Long, and Minzer's (KLLM) result [50], for instance, led to the resolution of Majority is Stablest in the p-biased setting [55], and the resolution of several conjectures in extremal combinatorics as well [51]. While previous results to this effect were restricted by their reliance on product structure or strong symmetry, we show such assumptions are not necessary and prove an analogous form of hypercontractivity for HDX.

Theorem 3.5 (Hypercontractivity on HDX). *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 $c_k \leq \min\{2^{O(k)}, k^{O(i)}\}$.

Crucially Theorem 3.5 is independent of k for small enough γ. This means our bounds remain meaningful even when k grows large (roughly speaking, one should think of the bound as being non-trivial in the regime where $k \ll \log(|X(1)|)$).[8] This was a crucial property in the analogous result on the Grassmann in the proof of the 2-2 Games Conjecture [54].

Our overall framework for proving Theorem 3.5 roughly follows Khot, Minzer, Moshkovitz, and Safra's [52] strategy for the complete complex. However, even with analogous results for the Bottom-Up Decomposition in hand, most of their techniques fail in our setting due to local-spectral expanders' distinct lack of product structure. In fact, Theorem 3.5 gives the first general class of hypercontractive objects beyond product spaces, and combined with known bounded degree constructions of local-spectral expanders [44, 57], 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 3.5 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 3.5, 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:

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

for some constant ρ. Despite the fact that coordinates do not exist on a simplicial complex, there is still a natural analog of T_ρ 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 3.6 (Noise Operator). *Let X be a d-dimensional pure simplicial complex. The noise operator $T_\rho^k(X) : C_k \to C_k$ 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,[9] 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 3.7. *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 $\delta \leq \varepsilon \|f\|_2^2 / \|f\|_\infty^2$. Then for some constant $\rho = \Theta(1)$:*

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

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 3.5 to give an analog on HDX for pseudorandom functions.

Theorem 3.8 (Level-i inequality). *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:*

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

[8] In reality, there is a more subtle trade-off here between the expansion parameters, degree, and dimension of the complex. The stated relation is for the complete complex where one optimizes expansion at the cost of degree.

[9] The 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.

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 3.9. *Let M be a walk on the kth 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:*

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

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* [7, 52]. Using Theorem 3.8, we can prove a converse to this result: *any* non-expanding set must be concentrated in a link.

Theorem 3.10 (Characterizing non-expansion). *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-walk[10] 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:*

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

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:

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

where e_i is the ith standard basis vector. One of the most celebrated results in the analysis of boolean functions is the KKL Theorem [40], 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**. While the KKL Theorem does not hold over arbitrary product spaces,[11] a useful analog known as "Bourgain's Sharp Threshold Theorem" [31, 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.

[10]Formally, 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 full paper that holds for all HD-walks.

[11]More accurately, it does hold but decays with the minimum probability of any marginal, becoming trivial e.g. for the p-biased cube for small enough p.

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. Using Theorem 3.8, we prove that any function with low total influence must be concentrated in a link.

Theorem 3.11 (Bourgain's Theorem for HDX). *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 $I[f] \leq K Var(f)$, there exists $i \leq K$ and an i-face τ such that the link of τ is dense:*

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

Note that Theorem 3.11 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.[12]

Proposition 3.12 (Bourgain's Theorem Lower Bound). *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:

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

(2) For every $i \leq cK$, all i-links are sparse:

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

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 3.13. *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:*

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

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

[12]A similar tight version of Bourgain's Theorem for sparse functions on the p-biased cube was proved by [50].

We emphasize that the first expectation in this definition is given by *restricting* rather than *localizing* 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 (as in say Definition 3.4). This latter notion of localization is also very important in analysis of HDX. Theorem 3.13 is proved through the machinery of *swap walks*, introduced independently by Alev, Jeronimo, and Tulsiani [2], and Dikstein and Dinur [18]. These walks, which crucially exhibit a very good spectral gap, have related applications in analyzing agreement tests [18], approximation algorithms [2], codes [1, 39], and most recently constructions of group-independent co-systolic expanders [43].

Theorem 3.13 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 $g \in C_2$:

$$\mathop{\mathbb{E}}_{a\sim X(1)} \mathop{\mathbb{E}}_{b\sim X_a(1)} \mathop{\mathbb{E}}_{c\sim X_{ab}(1)} \left[g(a,b)g(a,c)\right] \tag{4}$$

$$= \mathop{\mathbb{E}}_{a\sim X(1)} \mathop{\mathbb{E}}_{b\sim X_a(1)} \left[g(a,b) \mathop{\mathbb{E}}_{c\sim X_{ab}(1)} [g(a,c)]\right]. \tag{5}$$

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:

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

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 3.13 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 3.13 promises that

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

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

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

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.

We emphasize that while Equation (5) in particular could also have been analyzed through a more direct application of the swap walk, such techniques fail when additional copies of g are added. Since there are j copies of g in analysis of the jth moment, this means the traditional HDX tool kit 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 book [59] 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:*

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

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 [50, 52, 54], 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 [19, 48], where hypercontractivity was crucial to resolving the 2-2 Games Conjecture [54]. The spectral and ℓ_2-structure of these expanding posets (eposets) is well understood [2, 7, 19], and essentially has no dependence on the underlying poset structure.[13] In stark contrast, our results break down over

[13]Different poset parameters result in different eigenvalues, but the structure is otherwise the same.

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. On the other hand, variants of hypercontractivity are known for some special eposets 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 [3, 5, 6, 11, 14–16, 27, 37, 56]. 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 $\Theta(n\log(n))$ [17]. 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 [4, 11, 16]. 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 [40] 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 [13, 30, 32, 62]. These works offered a general theory of hypercontractivity for such domains, but their strength depended on the underlying distributions in the product space. This issue was addressed to an extent in work of Friedgut and Bourgain [31], and later Hatami [36], who showed analogues to the KKL theorem in product spaces for certain pseudorandom functions, but it was not until the recent work of Keevash, Lifshitz, Long, and Minzer [50] (and independently O'Donnell and Zhao [63]) that a true hypercontractive inequality was developed in this setting. This offered the missing piece for a number of classical applications including a tight variant of the KKL Theorem (for monotone functions) [50], Majority is Stablest [55], as well as a number of interesting applications to extremal combinatorics [51].

Another line of work has examined hypercontractivity on what are often called "exotic" domains: specific objects beyond products such as the slice [52], multislice [29, 61], Grassmannian [54] (or similarly the degree-two short code [9]), and symmetric group [28]. Like KLLM's improved result for product distributions, in unbalanced settings these examples are only hypercontractive for pseudorandom functions.[14] 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 [9, 24, 25, 52–54]. It is worth noting that agreement testing theorems are also known for local-spectral expanders [18, 23, 42] (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 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 (and independently in [35]).

Fourier Analysis on HDX:. Fourier analysis on HDX was originally studied by Diksein, Dinur, Filmus, and Harsha [19], 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 [45], 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 [2, 7, 47], and the nascent theory has helped build efficient approximation algorithms for certain k-CSPs [2] and unique games [7], but the restriction to second moment methods seems to have limited its use otherwise. Towards breaking this same barrier, Gur, Lifshitz, and Liu [35] 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.

ACKNOWLEDGEMENTS

MH was supported by NSF Award DGE-1650112. KT was supported by ERC and BSF awards. SL was supported by NSF Award CCF-1953928.

REFERENCES

[1] Vedat Levi Alev, Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. 2020. List decoding of direct sum codes. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1412–1425.

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

[3] Vedat Levi Alev and Lap Chi Lau. 2020. Improved Analysis of Higher Order Random Walks and Applications. *arXiv preprint arXiv:2001.02827* (2020).

[4] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. 2021. Entropic Independence in High-Dimensional Expanders: Modified

[14] We note that higher degrees of the short code are also hypercontractive, but only on low Fourier levels for general functions [8].

Log-Sobolev Inequalities for Fractionally Log-Concave Polynomials and the Ising Model. *arXiv preprint arXiv:2106.04105* (2021).
[5] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. 2020. Spectral Independence in High-Dimensional Expanders and Applications to the Hardcore Model. *arXiv preprint arXiv:2001.00303* (2020).
[6] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. 2019. Log-concave polynomials II: high-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 1–12.
[7] Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. 2020. High Dimensional Expanders: eigenstripping, Pseudorandomness, and Unique Games. *arXiv preprint arXiv:2011.04658* (2020).
[8] Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. 2012. Making the long code shorter. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE, 370–379.
[9] Boaz Barak, Pravesh K Kothari, and David Steurer. 2018. Small-set expansion in shortcode graph and the 2-to-2 conjecture. *arXiv preprint arXiv:1804.08662* (2018).
[10] William Beckner. 1975. Inequalities in Fourier analysis. *Annals of Mathematics* 102, 1 (1975), 159–182.
[11] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda. 2021. On mixing of Markov chains: Coupling, spectral independence, and entropy factorization. *arXiv preprint arXiv:2103.07459* (2021).
[12] Aline Bonami. 1970. Étude des coefficients de Fourier des fonctions de $L^p(G)$. In *Annales de l'institut Fourier*, Vol. 20. 335–402.
[13] Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. 1992. The influence of variables in product spaces. *Israel Journal of Mathematics* 77, 1 (1992), 55–64.
[14] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. 2021. Rapid mixing for colorings via spectral independence. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1548–1557.
[15] Zongchen Chen, Kuikui Liu, and Eric Vigoda. 2020. Rapid Mixing of Glauber Dynamics up to Uniqueness via Contraction. *arXiv preprint arXiv:2004.09083* (2020).
[16] Zongchen Chen, Kuikui Liu, and Eric Vigoda. 2021. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 1537–1550.
[17] Persi Diaconis and Laurent Saloff-Coste. 1996. Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability* 6, 3 (1996), 695–750.
[18] Yotam Dikstein and Irit Dinur. 2019. Agreement testing theorems on layered set systems. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1495–1524.
[19] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha. 2018. Boolean function analysis on high-dimensional expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
[20] Yotam Dikstein, Irit Dinur, Prahladh Harsha, and Noga Ron-Zewi. 2020. Locally testable codes via high-dimensional expanders. *arXiv preprint arXiv:2005.01045* (2020).
[21] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. 2021. Locally Testable Codes with constant rate, distance, and locality. *arXiv preprint arXiv:2111.04808* (2021).
[22] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta Shma. 2019. List decoding with double samplers. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2134–2153.
[23] Irit Dinur and Tali Kaufman. 2017. High dimensional expanders imply agreement expanders. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 974–985.
[24] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. 2018. On non-optimally expanding sets in Grassmann graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 940–951.
[25] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. 2018. Towards a proof of the 2-to-1 games conjecture?. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 376–389.
[26] Shai Evra, Tali Kaufman, and Gilles Zémor. 2020. Decodable quantum LDPC codes beyond the square root distance barrier using high dimensional expanders. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. 218–227.
[27] Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang. 2021. Rapid mixing from spectral independence beyond the Boolean domain. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1558–1577.
[28] Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer. 2020. Hypercontractivity on the symmetric group. *arXiv preprint arXiv:2009.05503* (2020).
[29] Yuval Filmus, Ryan O'Donnell, and Xinyu Wu. 2018. A log-Sobolev inequality for the multislice, with applications. *arXiv preprint arXiv:1809.03546* (2018).
[30] Ehud Friedgut. 1998. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica* 18, 1 (1998), 27–35.
[31] Ehud Friedgut and Jean Bourgain. 1999. Sharp thresholds of graph properties, and the k-sat problem. *Journal of the American mathematical Society* 12, 4 (1999), 1017–1054.
[32] Ehud Friedgut and Gil Kalai. 1996. Every monotone graph property has a sharp threshold. *Proceedings of the American mathematical Society* 124, 10 (1996), 2993–3002.
[33] Ehud Friedgut, Gil Kalai, and Assaf Naor. 2002. Boolean functions whose Fourier transform is concentrated on the first two levels. *Advances in Applied Mathematics* 29, 3 (2002), 427–437.
[34] Leonard Gross. 1975. Logarithmic sobolev inequalities. *American Journal of Mathematics* 97, 4 (1975), 1061–1083.
[35] Tom Gur, Noam Lifshitz, and Siqi Liu. 2021. Hypercontractivity on high dimensional expanders. *arXiv preprint arXiv:2111.09375* (2021).
[36] Hamed Hatami. 2012. A structure theorem for Boolean functions with small total influences. *Annals of Mathematics* (2012), 509–533.
[37] Vishesh Jain, Huy Tuan Pham, and Thuy Duong Vuong. 2021. Spectral independence, coupling with the stationary distribution, and the spectral gap of the Glauber dynamics. *arXiv preprint arXiv:2105.01201* (2021).
[38] Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. 2020. Unique Decoding of Explicit ε-balanced Codes Near the Gilbert-Varshamov Bound. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 434–445.
[39] Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani. 2021. Near-linear time decoding of Ta-Shma's codes via splittable regularity. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 1527–1536.
[40] J. Kahn, G. Kalai, and N. Linial. 1988. The influence of variables on Boolean functions. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Los Alamitos, CA, USA, 68–80. https://doi.org/10.1109/SFCS.1988.21923
[41] Tali Kaufman and David Mass. 2016. High dimensional combinatorial random walks and colorful expansion. *arXiv preprint arXiv:1604.02947* (2016).
[42] Tali Kaufman and David Mass. 2020. Local-To-Global Agreement Expansion via the Variance Method. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
[43] Tali Kaufman and David Mass. 2021. Unique-Neighbor-Like Expansion and Group-Independent Cosystolic Expansion. In *32nd International Symposium on Algorithms and Computation, ISAAC 2021, December 6-8, 2021, Fukuoka, Japan (LIPIcs, Vol. 212)*, Hee-Kap Ahn and Kunihiko Sadakane (Eds.). 56:1–56:17.
[44] Tali Kaufman and Izhar Oppenheim. 2018. Construction of new local spectral high dimensional expanders. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 773–786.
[45] Tali Kaufman and Izhar Oppenheim. 2020. High order random walks: Beyond spectral gap. *Combinatorica* (2020), 1–37.
[46] Tali Kaufman and Izhar Oppenheim. 2021. High dimensional expansion implies amplified local testability. *CoRR* (2021). https://arxiv.org/abs/2107.10488
[47] Tali Kaufman and Ella Sharakanski. 2020. Chernoff Bound for High-Dimensional Expanders. (2020). To Appear APPROX/RANDOM 2020.
[48] Tali Kaufman and Ran J Tessler. 2021. Local to global high dimensional expansion and Garland's method for general posets. *arXiv preprint arXiv:2101.12621* (2021).
[49] Tali Kaufman and Ran J. Tessler. 2021. New cosystolic expanders from tensors imply explicit Quantum LDPC codes with $\Omega(\sqrt{n}\log^k n)$ distance. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*. 1317–1329.
[50] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. 2019. Hypercontractivity for global functions and sharp thresholds. *arXiv preprint arXiv:1906.05568* (2019).
[51] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. 2021. Global hypercontractivity and its applications. *arXiv preprint arXiv:2103.04604* (2021).
[52] Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra. 2018. Small Set Expansion in The Johnson Graph.. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 25. 78.
[53] Subhash Khot, Dor Minzer, and Muli Safra. 2017. On independent sets, 2-to-2 games, and Grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 576–589.
[54] Subhash Khot, Dor Minzer, and Muli Safra. 2018. Pseudorandom sets in grassmann graph have near-perfect expansion. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 592–601.
[55] Noam Lifshitz and Dor Minzer. 2019. Noise sensitivity on the p-biased hypercube. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1205–1226.
[56] Kuikui Liu. 2021. From Coupling to Spectral Independence and Blackbox Comparison with the Down-Up Walk. *arXiv preprint arXiv:2103.11609* (2021).
[57] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. 2005. Explicit constructions of Ramanujan complexes of type Ad. *European Journal of Combinatorics* 26, 6 (2005), 965–993.
[58] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. 2005. Noise stability of functions with low influences: invariance and optimality. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE,

21–30.
[59] Ryan O'Donnell. 2014. *Analysis of boolean functions*. Cambridge University Press.
[60] Pavel Panteleev and Gleb Kalachev. 2021. Asymptotically good quantum and locally testable classical LDPC codes. *arXiv preprint arXiv:2111.03654* (2021).
[61] Justin Salez. 2020. A sharp log-Sobolev inequality for the multislice. *arXiv preprint arXiv:2004.05833* (2020).
[62] Michel Talagrand. 1994. On Russo's approximate zero-one law. *The Annals of Probability* (1994), 1576–1587.
[63] Yu Zhao. 2021. *Generalizations and Applications of Hypercontractivity and Small-Set Expansion*. Ph. D. Dissertation. Columbia University.