

Improved Small Set Expansion in High Dimensional Expanders

Tali Kaufman*

David Mass†

September 6, 2026

Abstract

Small set high dimensional expansion over general sheaves has recently become central to several major breakthroughs. These include the construction of classical locally testable and quantum LDPC codes with constant rate and linear distance, $\Omega(n)$ -levels lower bounds for the Sum-of-Squares hierarchy, and the proof of the NLTS conjecture.

In this work, we improve upon the state-of-the-art results for small set expansion in simplicial complexes, for both constant and general sheaves. Compared to one line of work (e.g., [KM22, DD24]), we improve upon their result by obtaining *strong expansion* for small sets. Compared to another line of work (e.g., [KKL14, EK16, KM21, FK22]), we get an *exponential improvement* in the size of sets for which expansion is guaranteed.

Our result is based on a novel technique that bridges these two lines of work via a “local vs. global” case analysis: for “local” sets, which are concentrated within small links, we adapt ideas from the “fat machinery” of [KKL14, EK16]; for “global” sets, which are spread across the complex, we exploit global averaging operators in the spirit of [KM22, DD24].

*Bar-Ilan University, Israel. Email: kaufmant@mit.edu.

†The Academic College of Tel Aviv-Yaffo, Israel. Email: dudimass@gmail.com.

1 Introduction

1.1 Background

High dimensional expansion. High dimensional expansion is a generalization of expansion in graphs to high dimensional objects. In this work we focus on *simplicial complexes*, which can be thought of as hypergraphs that are downwards closed. Specifically, A d -dimensional simplicial complex X is a $(d+1)$ -hypergraph with the property that if $\sigma \in X$ and $\tau \subset \sigma$ then also $\tau \in X$. A hyperedge of size $k+1$ is called a k -face of the complex.

We briefly describe the two main notions of high dimensional expansion, which generalize graph expansion to higher dimensions. The first, *local spectral expansion*, is concerned with the spectral expansion of the underlying graph of each local piece of the complex. These local pieces are called *links*, where the link of a face $\sigma \in X$ is the subcomplex obtained by taking all faces $\tau \supseteq \sigma$ and removing σ from them.

The other notion, generalizing edge expansion in graphs, is called *coboundary* or *cosystolic expansion*. For the sake of introduction, we describe this notion only over $\mathbb{F}_2$.¹ Consider a k -cochain $f : X(k) \rightarrow \mathbb{F}_2$, which is a $\{0, 1\}$ -assignment to the k -faces of the complex. The *coboundary* of f is a $\{0, 1\}$ -assignment to the $(k+1)$ -faces defined by $\delta f(\sigma) = \sum_{v \in \sigma} f(\sigma \setminus \{v\})$ for every $\sigma \in X(k+1)$. The coboundary of f can be thought of as a set of unsatisfied constraints, where each $(k+1)$ -face corresponds to a parity check of its k -subfaces. A k -cochain that satisfies all the constraints is called a k -cocycle, and a k -cochain that is obtained as the coboundary of a $(k-1)$ -cochain is called a k -coboundary.

Coboundary and cosystolic expansion² are concerned with the relationship between the fraction of unsatisfied constraints and the distance of the cochain from a satisfying assignment. More formally, a complex is said to be a β -cosystolic expander in dimension k if for every k -cochain f ,

$$\|\delta f\| \geq \beta \min_{g: \delta(g)=0} \|f - g\|, \quad (1.1)$$

where $\|\cdot\|$ is the normalized hamming weight.

For example, when $X = (X(0), X(1))$ is a (connected) graph, this notion coincides exactly with edge expansion. It can be seen by mapping a set of vertices $S \subseteq X(0)$ to its indicator function $\mathbf{1}_S$. Then $\delta(\mathbf{1}_S) = \mathbf{1}_{E(S, \bar{S})}$ is exactly the indicator function of the set of edges leaving S . In this case, the only functions g that satisfy $\delta(g) = \mathbf{0}$ are the constant functions. Hence, β -cosystolic expansion in dimension 0 can be equivalently written as

$$\forall S \subset X(0), S \neq \emptyset : \frac{\|E(S, \bar{S})\|}{\min\{\|S\|, \|\bar{S}\|\}} = \frac{\|\delta(\mathbf{1}_S)\|}{\min\{\|\mathbf{1}_S\|, \|\mathbf{1}_{\bar{S}}\|\}} \geq \beta.$$

Small set expansion. With the previous example in mind, small set expansion in simplicial complexes can be viewed as a generalization of graph small set expansion to higher dimensions. A naive approach, similar to graphs, would require that every cochain with small weight has a proportionally large coboundary. However, this approach fails for k -cochains with $k > 0$. For instance, fix an arbitrary vertex $v \in X(0)$ and take the set of edges containing this vertex, i.e., define the 1-cochain $f : X(1) \rightarrow \mathbb{F}_2$ by $f(e) = 1$ for every $e \ni v$, and $f(e) = 0$ for all other edges. f is completely supported on a single vertex, so $\|f\|$ is small. However, $\|\delta f\| = 0$, since every

¹See §2.4 for the general definition.

²The difference between coboundary and cosystolic expansion is whether there exist cocycles that are not coboundaries, i.e., satisfying assignments which are not obtained by the coboundary of an assignment of one dimension lower. For the sake of introduction, we can think of them as similar notions.

triangle contains either zero or two edges that contain v . This phenomenon is actually a general property of chain complexes, known as $\delta \circ \delta = 0$. Note that in this example, $f = \delta(\mathbf{1}_{\{v\}})$.

Roughly speaking, this should be the only obstacle to expanding small cochains. To be more precise, a k -cochain is said to be *minimal* if its weight cannot be reduced by adding a k -coboundary to it. It is said to be *locally minimal* if it is minimal at every proper link, i.e., its weight cannot be reduced by adding a k -coboundary that is supported on a lower dimensional face (hence the term *locally minimal*). We say that a complex is a (c_1, c_2) -small set expander³ in dimension k , if for every locally minimal k -cochain f ,

$$\|f\| \leq c_1 \quad \Rightarrow \quad \|\delta f\| \geq c_2 \|f\|.$$

On the importance of small set expansion. In [KKL14, EK16], the authors introduced small set expansion in higher dimensions as a tool for establishing cosystolic expansion. [KKL14, EK16] have shown that expansion of small locally minimal cochains implies cosystolic expansion in one dimension lower. Interestingly, even *weak expansion* is sufficient for this purpose, i.e., having $\|\delta f\| > 0$ for every (non-zero) small locally minimal k -cochain f implies cosystolic expansion in dimension $k - 1$.

A sequence of works followed [KKL14, EK16] and proved small set expansion for broader sets of coefficients (e.g., [KM21, KM22, DD24] for groups, and [FK22] for sheaves, discussed below) as a method to deduce cosystolic expansion in the same way.

Recently, remarkable breakthrough works constructed locally testable codes [DEL⁺22, PK22, FK24] and quantum LDPC codes [PK22, LZ22, DHLV23] with constant rate and linear distance. Their constructions are based on certain expanding *square complexes* over *sheaves* (discussed in the next paragraph). Following their works, small set expansion has found wider applications in other areas. For instance, [HL22] showed that the construction of [LZ22] admits small set expansion, which in turn yields hard instances for $\Omega(n)$ levels of the Sum-of-Squares hierarchy, and [ABN23] proved the NLTS conjecture, again by small set expansion of similar square complexes.

Sheaves. Over the past few years, sheaf theory on high dimensional expanders has emerged, providing an essential algebraic framework for these major breakthroughs. Broadly speaking, sheaves generalize the notion of a constant coefficient group (such as $\mathbb{F}_2$ that we discussed above) by assigning a local vector space to each face of the complex. Again, for simplicity, we focus on the case where these vector spaces are over $\mathbb{F}_2$. Formally, a sheaf $\mathcal{F}$ on a simplicial complex X is defined by the following two ingredients:

- (1) An $\mathbb{F}_2$ -vector space $\mathcal{F}(\sigma)$ for every $\sigma \in X$.
- (2) An $\mathbb{F}_2$ -linear map $\text{res}_{\tau \leftarrow \sigma} : \mathcal{F}(\sigma) \rightarrow \mathcal{F}(\tau)$ for every $\sigma \subset \tau \in X$, subject to the requirement that $\text{res}_{\tau \leftarrow \sigma} \circ \text{res}_{\sigma \leftarrow \rho} = \text{res}_{\tau \leftarrow \rho}$ for every $\rho \subset \sigma \subset \tau \in X$.

Now, a k -cochain f assigns an element of $\mathcal{F}(\sigma)$ to every $\sigma \in X(k)$, i.e., $f \in \prod_{\sigma \in X(k)} \mathcal{F}(\sigma)$. The coboundary operator is defined as before, but it has to take into account the linear maps as well. Formally, for any $\tau \in X(k + 1)$,

$$\delta f(\tau) = \sum_{\substack{\sigma \in X(k) \\ \sigma \subset \tau}} \text{res}_{\tau \leftarrow \sigma} f(\sigma).$$

³In some other works, this property is explicitly termed “small set locally minimal expansion”. For brevity, we call it here just small set expansion.

With the coboundary operator in hand, the rest of the definitions (e.g., cosystolic, coboundary, and small set expansion) remain the same. In this language, defining $\mathcal{F}(\sigma) = \mathbb{F}_2$ for every $\sigma \in X$, and $\text{res}_{\tau \leftarrow \sigma} = \text{id}$ for every $\sigma \subset \tau \in X$, recovers the definitions we discussed at the beginning of the introduction.

1.2 Our contribution

In this work, we improve upon the state-of-the-art results for small set expansion in simplicial complexes, for *both constant and general sheaves*.

Theorem 1.1 (Main theorem, informal). *Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . If X is a local spectral expander and its links are β -coboundary expanders, then for any $0 < k < d$, X is a (c_1, c_2) -small set expander in dimension k , where⁴*

$$c_1 \simeq \frac{\beta^k}{(k+2)!} \quad \text{and} \quad c_2 = \frac{\beta^k}{k!(k+1)^4}.$$

At a high level, previous works can be divided into two main paradigms. One line of work [KKL14, EK16, KM21, FK22] uses what has been termed “fat machinery” to establish small set expansion. However, this framework is significantly affected by the weight of the cochain, consequently obtaining expansion only for cochains of *very small* weight. Technically, they establish (c_1, c_2) -small set expansion in dimension k , where $c_1 \lesssim (\beta^k/k!)^{2^k}$.

The other line of work [KM22, DD24] uses local-to-global averaging operators, which improve upon the weight of cochains that expand but at the cost of the expansion quality. More precisely, they show that every locally minimal k -cocycle f satisfies $\|f\| \gtrsim \beta^k/(k+1)!$, which in turn implies that any (non-zero) locally minimal cochain of weight $\|f\| \lesssim \beta^k/(k+1)!$ is not a cocycle, i.e., it satisfies $\|\delta f\| > 0$. In our notation, they improve upon the constant c_1 , but provide no guarantee for c_2 .

In this work, we obtain “the best of both worlds” and show strong expansion for locally minimal cochains of weight $\lesssim \beta^k/(k+2)!$. We achieve an exponential improvement in the weight of expanding cochains over the first line of work, and improve upon the expansion quality of the second line of work.

Our technique. In order to describe our technique, let us briefly discuss the strategies of previous works. The first line of work [KKL14, EK16, KM21, FK22] can be thought of as deducing the expansion of small locally minimal cochains from their local expansion in the links. Their analysis divides faces of smaller dimensions into “fat faces” and “non-fat faces”. Their main argument (with a few variations between the works) builds upon the observation that whenever a cochain contains many fat faces, their local coboundary as seen in the links is actually a global coboundary. The main caveat of this framework is that to benefit from these fat faces, the cochain must be *very* small.

The second line of work [KM22, DD24] can be thought of as deducing expansion from the global structure of the cochain. They utilize (with a few variations between the works) certain local-to-global averaging operators, which are applicable only to locally minimal *cocycles*. In this way, they obtain a larger lower bound on the weight of a locally minimal cocycle, which in turn implies only $\|\delta f\| > 0$ for any small locally minimal f , i.e., with no guarantee on the expansion quality. The main caveat of this approach is the dependency on the global structure of the cochain being a cocycle.

⁴The $\simeq$ notation hides factors that depend on the local spectral expansion of the complex, which is assumed to be arbitrarily small. See Theorem 3.2 for the exact constants.

Our result is based on a novel technique that bridges these two frameworks via a “local vs. global” case analysis.

Key idea: We argue that any cochain either has a “local” structure in the sense that it is concentrated within links of some smaller dimension, or it has a “global” structure in the sense that it looks much like a cocycle. In the local case, we adapt ideas from the “fat machinery” approach of the first link of work, and in the global case, we exploit the global averaging operators in the spirit of the second line of work.

Decoupling spectral and coboundary expansion. Interestingly, our technique decouples the spectral expansion from the coboundary expansion in the links. In the local case, we deduce expansion solely from coboundary expansion of the links, and in the global case, we deduce expansion relying only on spectral arguments. This stands in contrast to all previous works, which simultaneously utilized both spectral and coboundary expansion of the links throughout their proofs.

1.3 Discussion

Applying our technique to general posets. We improve upon the state-of-the-art results for small set expansion in simplicial complexes, for both constant and general sheaves. However, as mentioned above, recent breakthroughs rely not only on expansion over sheaves, but also on complexes which are not simplicial, e.g., square complexes. A natural direction is to investigate whether our “local vs. global” technique introduced in this work can be applied to other types of complexes, potentially improving upon the aforementioned results.

More applications for small set expansion in simplicial complexes. As far as we are aware, for most applications of small set expansion in simplicial complexes (e.g., as a tool for proving cosystolic expansion), having $\|\delta f\| > 0$ for every (non-zero) small locally minimal cochain is sufficient. However, the expansion quality plays a crucial role in recent applications that are derived from square complexes. For instance, it affects the number of levels that cannot be refuted by the Sum-of-Squares hierarchy [HL22], and it plays a central role in the proof of the NLTS conjecture [ABN23]. It would be interesting to explore whether this type of strong expansion in simplicial complexes yields further applications.

Specifically, in a recent work, [DLV24] constructs “almost good” quantum LTCs, using expanding cubical complexes over sheaves. They suggest that a possible direction toward obtaining “good” quantum LTCs is to venture away from cubical complexes and use sheaves on simplicial complexes instead (see the discussion in [DLV24]). It would be interesting to see whether our improved small set expansion over sheaves can be useful toward this goal.

1.4 Organization

In Section 2 we provide required preliminaries. In Section 3 we prove our strong small set expansion theorem. Propositions 3.4 and 3.5 demonstrate our case analysis of general cochains: If the cochain is “local”, i.e., it is concentrated within links of some smaller dimension, we get expansion from Proposition 3.4, and if the cochain is “global”, i.e., it looks like a cocycle, we get expansion from Proposition 3.5. We prove our main theorem assuming these two propositions and then we turn to prove each proposition.

As this work bridges between two lines of work, some lemmas with slight variations appear in previous works and some others appear with different notations. We comment through the work regarding the similarities and differences with previous works.

2 Preliminaries

2.1 Simplicial complexes

A pure d -dimensional simplicial complex X is a downward closed $(d+1)$ -hypergraph, namely, X is a collection of sets (hyperedges) of size $d+1$ together with all of their subsets. A hyperedge of size $k+1$ is called a k -face, and the set of all k -faces of X is denoted by $X(k)$. For instance, $X(0)$, $X(1)$, and $X(2)$ correspond to the sets of vertices, edges, and triangles, respectively. We also assume that every complex contains a unique (-1) -face which corresponds to the empty set, i.e., $X(-1) = \{\emptyset\}$. All complexes in this work are assumed to be pure and simplicial.

For any k -face $\sigma \in X(k)$, the *link* of σ is a $(d-k-1)$ -dimensional complex defined by $X_\sigma = \{\tau \setminus \sigma \mid \sigma \subseteq \tau \in X\}$, i.e., the subcomplex obtained by taking all faces $\tau \supseteq \sigma$ and removing σ from them.

We also consider *ordered* faces of X , where an ordered k -face is a $(k+1)$ -tuple $(v_0, \dots, v_k)$ such that $\{v_0, \dots, v_k\} \in X(k)$. The set of all ordered k -faces in X is denoted by $\vec{X}(k)$. We assume that the complex comes with an arbitrary ordering on its vertices, so there is a one-to-one correspondence between the unordered k -faces and the ordered k -faces of the complex. For convenience, when it is clear from the context, we might refer to $\sigma \in X$ as both an unordered face and its unique corresponding ordered face. When we wish to emphasize that a face is ordered, we denote it by $\vec{\sigma}$.

Let $\vec{\sigma} = (v_0, \dots, v_\ell)$ and $\vec{\tau} = (v_{\ell+1}, \dots, v_k)$ be two disjoint ordered faces. We denote by $\vec{\sigma\tau} = (v_0, \dots, v_k)$ the ordered k -face that is obtained by the concatenation of $\vec{\sigma}$ and $\vec{\tau}$. For any $0 \leq i \leq \ell$, we denote by $\vec{\sigma}_i = (v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_\ell)$ the ordered $(\ell-1)$ -face that is obtained by removing v_i from $\vec{\sigma}$.

Weights and probabilities. Let X be a pure d -dimensional simplicial complex. We assume that X comes with a probability distribution over its d -faces $P_d : X(d) \rightarrow [0, 1]$. Consider the following random process: choose a random d -face $\sigma^{(d)}$ according to P_d , and then, iteratively for $k = d-1, \dots, -1$, remove a uniformly random vertex from $\sigma^{(k+1)}$ to obtain $\sigma^{(k)}$. This random process naturally induces a probability distribution $P_k : X(k) \rightarrow [0, 1]$ over the k -faces for every $k < d$, where for any $\sigma \in X(k)$,

$$P_k(\sigma) = \Pr_{\sigma^{(d)}, \dots, \sigma^{(-1)}}[\sigma^{(k)} = \sigma].$$

For any set of k -faces $S \subseteq X(k)$, we define its weight by

$$\|S\| = \Pr_{\sigma^{(d)}, \dots, \sigma^{(-1)}}[\sigma^{(k)} \in S] = \Pr_{\sigma \in X(k)}[\sigma \in S],$$

where the probability distribution at the right-hand side (and throughout this work) is taken according to P_k .

For any ℓ -face $\sigma \in X(\ell)$, the link X_σ inherits the probability distribution from X by conditioning on σ , i.e., for any $\ell \leq k \leq d$ and any $\tau \in X_\sigma(k-\ell-1)$,

$$P_{k-\ell-1}^{X_\sigma}(\tau) = \Pr_{\sigma^{(d)}, \dots, \sigma^{(-1)}}[\sigma^{(k)} = \sigma \cup \tau \mid \sigma^{(\ell)} = \sigma].$$

For any set of k -faces $S \subseteq X(k)$ and any $\ell < k$, the following decomposition of S into links of ℓ -faces follows immediately from the law of total probability:

$$\|S\| = \sum_{\sigma \in X(\ell)} \Pr_{\sigma^{(d)}, \dots, \sigma^{(-1)}}[\sigma^{(k)} \in S \mid \sigma^{(\ell)} = \sigma] \cdot P_\ell(\sigma) = \mathbb{E}_{\sigma \in X(\ell)} \left[\Pr_{\tau \in X_\sigma(k-\ell-1)}[\sigma \cup \tau \in S] \right].$$

2.2 Spectral expansion

A weighted graph is called a λ -one-sided spectral expander if the second largest eigenvalue of its normalized adjacency matrix is upper bounded by λ .

A d -dimensional simplicial complex X is said to be a λ -one-sided local spectral expander if for every $k \leq d - 2$ and every $\sigma \in X(k)$, the underlying graph of X_σ ⁵ is a λ -one-sided spectral expander.

We will use the following standard expander mixing lemma, which first appeared in [AC88]:

Lemma 2.1 (Expander mixing lemma). *Let $G = (V, E)$ be a λ -one-sided spectral expander. For any set of vertices $S \subseteq V$,*

$$\|E(S)\| \leq \|S\|^2 + \lambda\|S\|,$$

where $E(S)$ is the set of edges with both endpoints in S .

2.3 Sheaves

Let X be a d -dimensional simplicial complex and $\mathbb{F}$ a field.⁶ A sheaf $\mathcal{F}$ on X consists of:

- (1) An $\mathbb{F}$ -vector space $\mathcal{F}(\sigma)$ for every $\sigma \in X$.
- (2) An $\mathbb{F}$ -linear map $\text{res}_{\tau \leftarrow \sigma} : \mathcal{F}(\sigma) \rightarrow \mathcal{F}(\tau)$ for every $\sigma \subset \tau \in X$, subject to the requirement that $\text{res}_{\tau \leftarrow \sigma} \circ \text{res}_{\sigma \leftarrow \rho} = \text{res}_{\tau \leftarrow \rho}$ for every $\rho \subset \sigma \subset \tau \in X$.

The maps $\text{res}_{\tau \leftarrow \sigma}$ are called the *restriction maps* of $\mathcal{F}$.

Sheaf cohomology. Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . For any $-1 \leq k \leq d$, a k -cochain with coefficients in $\mathcal{F}$ is a function f that for every $\vec{\sigma} \in \vec{X}(k)$ assigns an element of $\mathcal{F}(\sigma)$, subject to the requirement that for every $\vec{\sigma} = (v_0, v_1, \dots, v_k) \in \vec{X}(k)$ and every permutation $\pi \in \text{Sym}(k+1)$,

$$f(\pi\vec{\sigma}) := f((v_{\pi(0)}, v_{\pi(1)}, \dots, v_{\pi(k)})) = \text{sgn}(\pi)f(\vec{\sigma}).$$

The space of all k -cochains with coefficients in $\mathcal{F}$ is denoted by $C^k(X, \mathcal{F})$.

For any $-1 \leq k \leq d - 1$, any k -cochain $f \in C^k(X, \mathcal{F})$ induces a $(k+1)$ -cochain by the coboundary operator $\delta^k : C^k(X, \mathcal{F}) \rightarrow C^{k+1}(X, \mathcal{F})$, defined for any $\vec{\sigma} \in \vec{X}(k+1)$ by

$$\delta^k f(\vec{\sigma}) = \sum_{i=0}^{k+1} (-1)^i \text{res}_{\sigma \leftarrow \sigma_i} f(\vec{\sigma}_i).$$

A direct calculation shows that for any $f \in C^k(X, \mathcal{F})$, indeed $\delta^k f \in C^{k+1}(X, \mathcal{F})$, and also that $\delta^{k+1}\delta^k f = \mathbf{0}$. We usually omit the superscript k from δ^k and denote the coboundary operator simply by δ , where k is clear from context.

The kernel of δ^k is called the space of k -cocycles and denoted by

$$Z^k(X, \mathcal{F}) = \ker \delta^k = \{f \in C^k(X, \mathcal{F}) \mid \delta^k f = \mathbf{0}\}.$$

The image of δ^{k-1} is called the space of k -coboundaries and denoted by

$$B^k(X, \mathcal{F}) = \text{im } \delta^{k-1} = \{\delta^{k-1} f \mid f \in C^{k-1}(X, \mathcal{F})\},$$

with the convention that $\delta^{-2} = 0$. Since $\delta^k \circ \delta^{k-1} = 0$ for every $-1 \leq k \leq d - 1$, it follows that

$$B^k(X, \mathcal{F}) \subseteq Z^k(X, \mathcal{F}) \subseteq C^k(X, \mathcal{F}).$$

⁵The weighted graph whose vertices are $X_\sigma(0)$ and edges are $X_\sigma(1)$.

⁶One can similarly define sheaves over abelian groups or over R -modules for a ring R . For simplicity, we consider only sheaves over fields, which are called $\mathbb{F}$ -sheaves; see [FK24] for general definitions and examples.

2.4 Coboundary and cosystolic expansion

Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . For any $f \in C^k(X, \mathcal{F})$, the normalized hamming weight of f is defined by

$$\|f\| = \Pr_{\sigma \in X(k)} [f(\sigma) \neq 0],$$

where the probability is taken according to the distribution as described in §2.1.

Since the weight of a cochain is dependent only on its non-zero elements, it is often convenient to consider the set $\text{supp}(f)$, i.e., the set of non-zero elements in f , and define equivalently

$$\|f\| = \|\text{supp}(f)\| = \Pr_{\sigma \in X(k)} [\sigma \in \text{supp}(f)].$$

The distance between two k -cochains $f, g \in C^k(X, \mathcal{F})$ is defined by $\text{dist}(f, g) = \|f - g\|$.

Cosystolic expansion. For any $-1 \leq k \leq d-1$, X is called a (β, ε) -cosystolic expander in dimension k , if the following two conditions hold:

- (1) For every $f \in C^k(X, \mathcal{F}) \setminus Z^k(X, \mathcal{F})$,

$$\frac{\|\delta(f)\|}{\text{dist}(f, Z^k(X, \mathcal{F}))} \geq \beta,$$

where $\text{dist}(f, Z^k(X, \mathcal{F})) = \min\{\text{dist}(f, g) \mid g \in Z^k(X, \mathcal{F})\}$.

- (2) For every $f \in Z^k(X, \mathcal{F}) \setminus B^k(X, \mathcal{F})$,

$$\|f\| \geq \varepsilon.$$

The second condition ensures that there are no small cocycles which are not coboundaries.

Coboundary expansion. Coboundary expansion is similar to cosystolic expansion with the extra condition that $Z^k(X, \mathcal{F}) = B^k(X, \mathcal{F})$. Formally, for any $-1 \leq k \leq d-1$, X is called a β -coboundary expander in dimension k , if for every $f \in C^k(X, \mathcal{F}) \setminus B^k(X, \mathcal{F})$,

$$\frac{\|\delta(f)\|}{\text{dist}(f, B^k(X, \mathcal{F}))} \geq \beta,$$

where $\text{dist}(f, B^k(X, \mathcal{F})) = \min\{\text{dist}(f, g) \mid g \in B^k(X, \mathcal{F})\}$.

2.5 Locally minimal cochains

Restricting sheaves to links. Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . For any $\sigma \in X$, the sheaf $\mathcal{F}$ induces a sheaf $\mathcal{F}_\sigma$ restricted to the link X_σ , defined by:

- (1) $\mathcal{F}_\sigma(\tau) = \mathcal{F}(\sigma \cup \tau)$ for every $\tau \in X_\sigma$, and
(2) $\text{res}_{\rho \leftarrow \tau}^{\mathcal{F}_\sigma} = \text{res}_{\sigma \cup \rho \leftarrow \sigma \cup \tau}^{\mathcal{F}}$ for every $\tau \subset \rho \in X_\sigma$.

Localization. Let $f \in C^k(X, \mathcal{F})$ be a k -cochain. For any $0 \leq \ell \leq k$ and $\sigma \in X(\ell)$, the *localization* of f to the link X_σ , denoted by $f_\sigma \in C^{k-\ell-1}(X_\sigma, \mathcal{F}_\sigma)$, is a $(k-\ell-1)$ -cochain in X_σ , defined by $f_\sigma(\vec{\tau}) = f(\vec{\sigma}\vec{\tau})$ for every ordered $(k-\ell-1)$ -face $\vec{\tau} \in \vec{X}_\sigma(k-\ell-1)$ (where $\vec{\sigma}$ is the unique ordered face corresponding to σ).

Minimal and locally minimal cochains. A k -cochain $f \in C^k(X, \mathcal{F})$ is said to be *minimal* if its weight cannot be reduced by adding a coboundary to it. Namely, for every coboundary $g \in B^k(X, \mathcal{F})$ it holds that $\|f\| \leq \|f - g\|$. Recall that the distance of f from the coboundaries is defined by $\text{dist}(f, B^k(X, \mathcal{F})) = \min\{\|f - g\| \mid g \in B^k(X, \mathcal{F})\}$. Since $\mathbf{0} \in B^k(X, \mathcal{F})$, it follows that $\|f\| \geq \text{dist}(f, B^k(X, \mathcal{F}))$ for every $f \in C^k(X, \mathcal{F})$. Hence, f is said to be minimal if and only if $\|f\| = \text{dist}(f, B^k(X, \mathcal{F}))$.

A k -cochain $f \in C^k(X, \mathcal{F})$ is said to be *locally minimal* if the localization of f to the link of σ is minimal for every non-empty $\sigma \in X$, i.e., f_σ is minimal in X_σ for every $\sigma \in X \setminus \{\emptyset\}$.

3 Small set expansion

We recall the definition of small set expansion from the introduction.

Definition 3.1 (Small set expansion). Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . Let $k \in \{0, \dots, d-1\}$, and $c_1, c_2 \in [0, 1]$. X is said to be a (c_1, c_2) -small set expander in dimension k , if for every locally minimal k -cochain $f \in C^k(X, \mathcal{F})$,

$$\|f\| \leq c_1 \quad \Rightarrow \quad \|\delta f\| \geq c_2 \|f\|.$$

Our main theorem is a local criterion for small set expansion.

Theorem 3.2 (Main theorem). Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . Let $k \in \{1, \dots, d-1\}$, $\lambda \in [0, 1)$ and $\beta_0, \dots, \beta_k \in (0, 1]$. If X is a λ -one-sided local spectral expander and for every $\sigma \in X(\ell)$, $0 \leq \ell \leq k$, X_σ is a β_ℓ -coboundary expander in dimension $k - \ell - 1$, then X is a (c_1, c_2) -small set expander in dimension k , where

$$c_1 = \left(\frac{\beta_k}{k+2} - \lambda \right) \frac{\beta_0 \cdots \beta_{k-1}}{(k+1)!} - 2\lambda \quad \text{and} \quad c_2 = \frac{\beta_0 \cdots \beta_k}{k!(k+1)^4} \|f\|.$$

Remark. For the sake of comparison with previous works, when $\mathcal{F}$ is a constant sheaf (e.g., [KM21, KM22, DD24]), then $\beta_k = 1$. In the case of a general sheaf (e.g., [FK22]), β_k could be smaller than 1.

The main idea is to deal separately with two types of small cochains. If the cochain has a global structure in the sense that it looks like a cocycle, we would gain expansion by spectral arguments. Otherwise, if the cochain has a local structure in the sense that it is concentrated within links of some smaller dimension, we would gain expansion by the coboundary expansion of these links.

In order to distinguish between these two types of cochains, we provide another way of looking at neighborhoods of faces, in addition to the links.

Let X be a d -dimensional simplicial complex, $\mathcal{F}$ a sheaf on X , $f \in C^k(X, \mathcal{F})$, and $\sigma \in X(\ell)$, $\ell \leq k < d$. Recall that the link of σ is a $(d - \ell - 1)$ -dimensional simplicial complex defined by $X_\sigma = \{\tau \setminus \underline{\sigma} \mid \sigma \subseteq \tau \in X\}$, and the localization of f to X_σ is defined by $f_\sigma(\vec{\tau}) = f(\vec{\sigma}\vec{\tau})$, for every $\vec{\tau} \in \overrightarrow{X_\sigma}(k - \ell - 1)$.

We define Y_σ^k to be the set of k -faces that together with σ form a $(k+1)$ -face. Formally,

$$Y_\sigma^k := \{\tau \setminus \{v\} \mid v \in \sigma \subset \tau \in X(k+1)\}.$$

We denote by $f^\sigma = f|_{Y_\sigma^k}$ the restriction of f to the k -faces in Y_σ^k . We also define a probability distribution on Y_σ^k , $P_\sigma^{Y_\sigma^k} : Y_\sigma^k \rightarrow [0, 1]$, induced by the probability distribution of X : for every

k -face $\tau \setminus \{v\} \in Y_\sigma^k$, its probability corresponds to choosing the $(k - \ell)$ -face $\tau \setminus \sigma$ in the link X_σ , and then choosing $v \in \sigma$ uniformly at random. Formally,

$$P^{Y_\sigma^k}(\tau \setminus \{v\}) = \frac{P_{k-\ell}^{X_\sigma}(\tau \setminus \sigma)}{\ell + 1},$$

where we recall that $P_{k-\ell}^{X_\sigma}$ is the probability distribution over $(k - \ell)$ -faces in the link X_σ .

The definitions of Y_σ^k and f^σ are motivated by the following structural observation. Consider a $(k + 1)$ -face $\tau \supset \sigma$. Note that any k -face contained in τ either contains σ or does not contain σ . Consequently, the k -faces contained in τ can be decomposed into $(k - \ell - 1)$ -faces in the link X_σ and k -faces in Y_σ^k . We note that this observation and the definition of f^σ appeared in [KM22], and later also in [DD24] with a different notation.

Informally speaking, we expect that for a cochain f that looks like a cocycle, i.e., a “global” cochain, most faces σ will satisfy $\|f_\sigma\| \lesssim \|f^\sigma\|$. The following definition allows us to distinguish between cochains that are “global” and cochains that are concentrated within links of lower dimension, i.e., “local” cochains.

Definition 3.3 (Heavy faces⁷). Let X be a d -dimensional simplicial complex, $\mathcal{F}$ a sheaf on X , $f \in C^k(X, \mathcal{F})$, and $\sigma \in X(\ell)$, for $0 \leq \ell < k \leq d - 1$. Assume that X_σ is a β -coboundary expander in dimension $k - \ell - 1$. We say that f is *heavy* on σ if

$$\|f_\sigma\| > (\ell + 2)\beta^{-1}\|f^\sigma\|.$$

We denote by $\text{Heavy}_\ell(f) = \{\sigma \in X(\ell) \mid f \text{ is heavy on } \sigma\}$ the subset of ℓ -faces that are heavy.

The two aforementioned cases are dealt in the following two propositions.

Proposition 3.4 (Expansion in the local case). *Let X be a d -dimensional simplicial complex and $\mathcal{F}$ a sheaf on X . Let $0 \leq \ell < k \leq d - 1$, and assume that for every $\sigma \in X(\ell)$, X_σ is a β_ℓ -coboundary expander in dimension $k - \ell - 1$. Then for any locally minimal $f \in C^k(X, \mathcal{F})$,*

$$\|\delta f\| \geq \frac{\beta_\ell}{\ell + 2} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\|.$$

Proposition 3.5 (Expansion in the global case). *Let X be a d -dimensional λ -one-sided local spectral expander and $\mathcal{F}$ a sheaf on X . Let $k \in \{1, \dots, d - 1\}$ and assume that for every $\sigma \in X(\ell)$, $0 \leq \ell \leq k$, X_σ is a β_ℓ -coboundary expander in dimension $k - \ell - 1$. For any $f \in C^k(X, \mathcal{F})$, if*

$$\|f\| \leq \left(\frac{\beta_k}{k + 2} - \lambda \right) \frac{\beta_0 \cdots \beta_{k-1}}{(k + 1)!} - 2\lambda,$$

then

$$\|\delta f\| \geq \beta_k \|f\| - \sum_{\ell=0}^{k-1} \frac{(k + 1)(k + 2)!}{\beta_{\ell+1} \cdots \beta_{k-1}(\ell + 2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\|.$$

We note that the assumption of β_ℓ -coboundary expansion in the links of ℓ -faces, $0 \leq \ell < k$, is not required for the proof of Proposition 3.5, but only for heavy faces to be defined. In other words, one could choose any arbitrary constants β_ℓ for $0 \leq \ell < k$, and Proposition 3.5 would still hold.

We first prove Theorem 3.2 assuming Propositions 3.4 and 3.5, and then prove the propositions.

⁷We use the term “heavy faces” in order to distinguish it from the notion of “fat faces” of previous works. We say that a face σ is “heavy” if $\|f_\sigma\|$ is large *relative* to $\|f^\sigma\|$, whereas a face σ is said to be “fat” according to previous works, if $\|f_\sigma\|$ is larger than some *absolute constant*.

Proof of Theorem 3.2. If there exists $0 \leq \ell \leq k-1$ such that

$$\mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \geq \frac{\beta_{\ell+1} \cdots \beta_k (\ell+2)!}{k(k+1)(k+2)! + \beta_0 \cdots \beta_{k-1}} \|f\|,$$

then by Proposition 3.4

$$\|\delta f\| \geq \frac{\beta_\ell}{\ell+2} \cdot \frac{\beta_{\ell+1} \cdots \beta_k (\ell+2)!}{k(k+1)(k+2)! + \beta_0 \cdots \beta_{k-1}} \|f\| \geq \frac{\beta_0 \cdots \beta_k}{k!(k+1)^4} \|f\|,$$

where the second inequality holds since $\beta_\ell \leq 1$ for every $\ell \geq 0$.

Otherwise, it holds that

$$\sum_{\ell=0}^{k-1} \frac{(k+1)(k+2)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \leq \frac{\beta_k k(k+1)(k+2)!}{k(k+1)(k+2)! + \beta_0 \cdots \beta_{k-1}} \|f\|.$$

Then by Proposition 3.5

$$\begin{aligned} \|\delta f\| &\geq \beta_k \|f\| - \sum_{\ell=0}^{k-1} \frac{(k+1)(k+2)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \\ &\geq \beta_k \|f\| - \frac{\beta_k k(k+1)(k+2)!}{k(k+1)(k+2)! + \beta_0 \cdots \beta_{k-1}} \|f\| \\ &\geq \frac{\beta_0 \cdots \beta_k}{k!(k+1)^4} \|f\|. \end{aligned}$$

□

Proof of Proposition 3.4

In this case, when the cochain is concentrated within links of lower dimension $\ell < k$, we show that the local coboundary expansion inside the links of heavy ℓ -faces implies global expansion.

Let f be a locally minimal k -cochain. Consider a heavy ℓ -face $\sigma \in \text{Heavy}_\ell(f)$ and a $(k-\ell)$ -face $\tau \in X_\sigma(k-\ell)$ such that $\delta(f_\sigma)(\tau) \neq 0$. Denote the ordered faces corresponding to σ and τ by $\vec{\sigma} = (v_0, v_1, \dots, v_\ell)$ and $\vec{\tau} = (v_{\ell+1}, v_{\ell+2}, \dots, v_{k+1})$. By definition of the local coboundary,

$$0 \neq \delta(f_\sigma)(\vec{\tau}) = \sum_{i=\ell+1}^{k+1} (-1)^{i-\ell-1} \text{res}_{\vec{\tau} \leftarrow \tau_i}^{\mathcal{F}_\sigma} f_\sigma(\vec{\tau}_i) = \sum_{i=\ell+1}^{k+1} (-1)^{i-\ell-1} \text{res}_{\sigma \cup \tau \leftarrow \sigma \cup \tau_i}^{\mathcal{F}} f(\vec{\sigma} \vec{\tau}_i).$$

Now, if $f((\vec{\sigma} \vec{\tau})_i) = 0$ for every $0 \leq i \leq \ell$, then

$$\delta f(\vec{\sigma} \vec{\tau}) = \sum_{i=0}^{k+1} (-1)^i \text{res}_{\sigma \cup \tau \leftarrow (\sigma \cup \tau)_i} f((\vec{\sigma} \vec{\tau})_i) = (-1)^{\ell+1} \sum_{i=\ell+1}^{k+1} (-1)^{i-\ell-1} \text{res}_{\sigma \cup \tau \leftarrow \sigma \cup \tau_i} f(\vec{\sigma} \vec{\tau}_i) \neq 0.$$

Therefore, in this case, a local non-zero coboundary implies a global non-zero coboundary, i.e.,

$$\delta(f_\sigma)(\vec{\tau}) \neq 0 \wedge \forall_{0 \leq i \leq \ell} f((\vec{\sigma} \vec{\tau})_i) = 0 \Rightarrow \delta f(\vec{\sigma} \vec{\tau}) \neq 0. \quad (3.1)$$

Let $\sigma \in \text{Heavy}_\ell(f)$ be a heavy ℓ -face. We will show that many $(k-\ell)$ -faces in X_σ satisfy the hypothesis of (3.1). Recall the probability distribution over Y_σ^k : the probability of a k -face

$\rho \in Y_\sigma^k$ is the probability of obtaining $\rho = (\sigma \cup \tau) \setminus \{v\}$ by choosing a $(k - \ell)$ -face $\tau \in X_\sigma(k - \ell)$ according to the probability distribution of $X_\sigma(k - \ell)$ and a vertex $v \in \sigma$ uniformly at random. It holds that

$$\begin{aligned} \|f^\sigma\| &= \Pr_{\substack{\tau \in X_\sigma(k-\ell) \\ i \in \{0, \dots, \ell\}}} [f((\vec{\sigma\tau})_i) \neq 0] \\ &= \frac{1}{\ell + 1} \sum_{i=0}^{\ell} \Pr_{\tau \in X_\sigma(k-\ell)} [f((\vec{\sigma\tau})_i) \neq 0] \\ &\geq \frac{1}{\ell + 1} \Pr_{\tau \in X_\sigma(k-\ell)} [\exists_{0 \leq i \leq \ell} \text{ s.t. } f((\vec{\sigma\tau})_i) \neq 0]. \end{aligned} \tag{3.2}$$

Combining (3.1) and (3.2) yields

$$\begin{aligned} \Pr_{\tau \in X_\sigma(k-\ell)} [\delta f(\vec{\sigma\tau}) \neq 0] &\geq \Pr_{\tau \in X_\sigma(k-\ell)} [\delta(f_\sigma)(\vec{\tau}) \neq 0 \wedge \forall_{0 \leq i \leq \ell} f((\vec{\sigma\tau})_i) = 0] \\ &\geq \Pr_{\tau \in X_\sigma(k-\ell)} [\delta(f_\sigma)(\vec{\tau}) \neq 0] - \Pr_{\tau \in X_\sigma(k-\ell)} [\exists_{0 \leq i \leq \ell} \text{ s.t. } f((\vec{\sigma\tau})_i) \neq 0] \\ &\geq \|\delta(f_\sigma)\| - (\ell + 1) \|f^\sigma\| \\ &\geq \beta_\ell \|f_\sigma\| - \frac{(\ell + 1)\beta_\ell}{\ell + 2} \|f_\sigma\| \\ &= \frac{\beta_\ell}{\ell + 2} \|f_\sigma\|, \end{aligned}$$

where the first inequality follows by (3.1), the second inequality follows by the law of total probability, the third inequality follows by (3.2), and the last inequality follows since X_σ is a β_ℓ -coboundary expander in dimension $k - \ell - 1$, f is locally minimal, and $\sigma \in \text{Heavy}_\ell(f)$.

It follows that

$$\begin{aligned} \|\delta f\| &= \mathbb{E}_{\sigma \in X(\ell)} \left[\Pr_{\tau \in X_\sigma(k-\ell)} [\delta f(\vec{\sigma\tau}) \neq 0] \right] \\ &\geq \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} \left[\Pr_{\tau \in X_\sigma(k-\ell)} [\delta f(\vec{\sigma\tau}) \neq 0] \right] \cdot \|\text{Heavy}_\ell(f)\| \\ &\geq \frac{\beta_\ell}{\ell + 2} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\|, \end{aligned}$$

which completes the proof.

Proof of Proposition 3.5

In this case, when there are only a few heavy links in every dimension, we can show by spectral arguments that there are many $(k + 1)$ -faces that contain exactly one k -face σ for which $f(\sigma) \neq 0$.

The proof of Proposition 3.5 follows from the following two lemmas.

Lemma 3.6. *Let X be a d -dimensional λ -one-sided local spectral expander and $\mathcal{F}$ a sheaf on X . Let $k \in \{1, \dots, d - 1\}$ and assume that for every $\sigma \in X(k)$, X_σ is a β_k -coboundary expander in dimension -1 . Then for any $f \in C^k(X, \mathcal{F})$,*

$$\|\delta f\| \geq (k + 2) \left(\beta_k \|f\| - (k + 1) \mathbb{E}_{\sigma \in X(k)} [\|f_\sigma\| \|f^\sigma\|] \right).$$

Lemma 3.7. *Let X be a d -dimensional λ -one-sided local spectral expander and $\mathcal{F}$ a sheaf on X . For any $f \in C^k(X, \mathcal{F})$, $0 \leq k < d$ and $0 \leq \ell \leq k$,*

$$\mathbb{E}_{\sigma \in X(\ell)} [\|f_\sigma\| \|f^\sigma\|] \leq \mathbb{E}_{\sigma \in X(\ell-1)} [\|f_\sigma\|^2] + \lambda \|f\|.$$

We note that both lemmas are inspired from previous works. A slight variation of Lemma 3.6 (with regard to constant sheaves) has appeared in [KM22], and Lemma 3.7, with a different notation, has appeared in [DD24]. For completeness, we provide proofs for both lemmas after we show that Proposition 3.5 follows from them.

Proof of Proposition 3.5. First, by Lemma 3.7 applied to $\ell = k$,

$$\mathbb{E}_{\sigma \in X(k)} [\|f_\sigma\| \|f^\sigma\|] \leq \mathbb{E}_{\sigma \in X(k-1)} [\|f_\sigma\|^2] + \lambda \|f\|. \quad (3.3)$$

Then, for any $0 \leq \ell < k$,

$$\begin{aligned} \mathbb{E}_{\sigma \in X(\ell)} [\|f_\sigma\|^2] &= \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|^2] \cdot \|\overline{\text{Heavy}_\ell(f)}\| + \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|^2] \cdot \|\text{Heavy}_\ell(f)\| \\ &\leq \frac{\ell+2}{\beta_\ell} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\| \|f^\sigma\|] \cdot \|\overline{\text{Heavy}_\ell(f)}\| + \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \\ &\leq \frac{\ell+2}{\beta_\ell} \mathbb{E}_{\sigma \in X(\ell)} [\|f_\sigma\| \|f^\sigma\|] + \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \\ &\leq \frac{\ell+2}{\beta_\ell} \left(\mathbb{E}_{\sigma \in X(\ell-1)} [\|f_\sigma\|^2] + \lambda \|f\| \right) + \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\|, \end{aligned} \quad (3.4)$$

where the first inequality holds by the definition of non-heavy faces (for the left summand) and that $\|f_\sigma\|^2 \leq \|f_\sigma\|$ (for the right summand), the second inequality follows from the law of total expectation, and the last inequality holds by Lemma 3.7.

Applying (3.4) iteratively with $\ell = k-1, k-2, \dots, 0$ yields

$$\begin{aligned} \mathbb{E}_{\sigma \in X(k-1)} [\|f_\sigma\|^2] &\leq \frac{(k+1)!}{\beta_0 \cdots \beta_{k-1}} \left(\|f\|^2 + \sum_{\ell=0}^{k-1} \frac{\beta_0 \cdots \beta_{\ell-1}}{(\ell+1)!} \lambda \|f\| \right) + \\ &\quad \sum_{\ell=0}^{k-1} \frac{(k+1)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \\ &\leq \frac{(k+1)!}{\beta_0 \cdots \beta_{k-1}} (\|f\| + 2\lambda) \|f\| + \\ &\quad \sum_{\ell=0}^{k-1} \frac{(k+1)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\| \\ &\leq \left(\frac{\beta_k}{k+2} - \lambda \right) \|f\| + \\ &\quad \sum_{\ell=0}^{k-1} \frac{(k+1)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\|] \cdot \|\text{Heavy}_\ell(f)\|, \end{aligned} \quad (3.5)$$

where the second inequality holds since $\beta_\ell \leq 1$ for every $\ell \geq 0$, and the last inequality holds since $\|f\| \leq \left(\frac{\beta_k}{k+2} - \lambda \right) \frac{\beta_0 \cdots \beta_{k-1}}{(k+1)!} - 2\lambda$.

Substituting (3.5) in (3.3) implies

$$\mathbb{E}_{\sigma \in X(k)} [\|f_\sigma\| \|f^\sigma\|] \leq \frac{\beta_k}{k+2} \|f\| + \sum_{\ell=0}^{k-1} \frac{(k+1)!}{\beta_{\ell+1} \cdots \beta_{k-1} (\ell+2)!} \mathbb{E}_{\sigma \in \text{Heavy}_\ell(f)} [\|f_\sigma\| \cdot \|\text{Heavy}_\ell(f)\|]. \quad (3.6)$$

Finally, substituting (3.6) in Lemma 3.6 completes the proof. $\square$

Proofs of Lemmas 3.6 and 3.7

Proof of Lemma 3.6. Let $f \in C^k(X, \mathcal{F})$. For ease of notation, we identify f and δf with their support, i.e., we write $\sigma \in f$ and $\tau \in \delta f$ when $f(\sigma) \neq 0$ and $\delta f(\tau) \neq 0$, respectively.

Let us denote by $\delta_1(f)$ the set of $(k+1)$ -faces in δf that contain exactly one k -face in f . Formally,

$$\delta_1(f) = \{\tau \in \delta f \mid |\{\sigma \subset \tau \mid \sigma \in f\}| = 1\}.$$

Note that $\|\delta f\| \geq \|\delta_1(f)\|$, since by definition, $\tau \in \delta_1(f)$ implies $\tau \in \delta f$ (i.e., $\delta f(\tau) \neq 0$). So it is enough to lower bound $\|\delta_1(f)\|$.

First, for any $\tau \in X(k+1)$ such that $|\{\sigma \subset \tau \mid \sigma \in f\}| = 1$, if σ is the only k -face in τ that is in f (i.e., σ is the only k -face in τ for which $f(\sigma) \neq 0$), then σ is the only k -face that contributes to $\delta f(\tau)$. It follows that

$$\tau \in \delta f \iff \tau \setminus \sigma \in \delta(f_\sigma). \quad (3.7)$$

Now, let $\sigma \in f$. Recall the probability distribution over Y_σ^k (where σ is a k -face): the probability of a k -face $\rho \in Y_\sigma^k$ is the probability that $\rho = (\sigma \cup \{v\}) \setminus \{u\}$, where $v \in X_\sigma(0)$ is chosen according to the probability distribution of $X_\sigma(0)$ and $u \in \sigma$ is chosen uniformly at random. For readability, in what follows we slightly abuse notation and write σv and $\sigma v \hat{u}$ as shorthand for $(\sigma \cup \{v\})$ and $(\sigma \cup \{v\}) \setminus \{u\}$, respectively. It holds that

$$\begin{aligned} \|f^\sigma\| &= \Pr_{\substack{v \in X_\sigma(0) \\ u \in \sigma}} [\sigma v \hat{u} \in f] \\ &= \frac{1}{k+1} \sum_{u \in \sigma} \Pr_{v \in X_\sigma(0)} [\sigma v \hat{u} \in f] \\ &\geq \frac{1}{k+1} \Pr_{v \in X_\sigma(0)} [\exists u \in \sigma \text{ s.t. } \sigma v \hat{u} \in f]. \end{aligned} \quad (3.8)$$

Hence, by the definition of $\delta_1(f)$,

$$\begin{aligned} \Pr_{v \in X_\sigma(0)} [\sigma v \in \delta_1(f)] &= \Pr_{v \in X_\sigma(0)} [\sigma v \in \delta f \wedge |\{\rho \subset \sigma v \mid \rho \in f\}| = 1] \\ &= \Pr_{v \in X_\sigma(0)} [\sigma v \in \delta f \wedge \forall u \in \sigma, \sigma v \hat{u} \notin f] \\ &= \Pr_{v \in X_\sigma(0)} [v \in \delta(f_\sigma) \wedge \forall u \in \sigma, \sigma v \hat{u} \notin f] \\ &\geq \Pr_{v \in X_\sigma(0)} [v \in \delta(f_\sigma)] - \Pr_{v \in X_\sigma(0)} [\exists u \in \sigma \text{ s.t. } \sigma v \hat{u} \in f] \\ &\geq \beta_k - (k+1) \|f^\sigma\|, \end{aligned} \quad (3.9)$$

where the second equality holds since $\sigma \in f$, the third equality follows by (3.7), and the last inequality follows by (3.8), the β_k -coboundary expansion of X_σ in dimension -1 , and the assumption that $\sigma \in f$, so $\|f_\sigma\| = 1$.

Now, on one hand,

$$\begin{aligned}
\Pr_{\substack{\sigma \in X(k) \\ v \in X_\sigma(0)}} [\sigma \in f \wedge \sigma v \in \delta_1(f)] &= \mathbb{E}_{\sigma \in X(k)} \left[\Pr_{v \in X_\sigma(0)} [\sigma \in f \wedge \sigma v \in \delta_1(f)] \right] \\
&= \mathbb{E}_{\sigma \in X(k)} \left[\mathbf{1}_{\{\sigma \in f\}}(\sigma) \cdot \Pr_{v \in X_\sigma(0)} [\sigma v \in \delta_1(f)] \right] \\
&\geq \mathbb{E}_{\sigma \in X(k)} \left[\mathbf{1}_{\{\sigma \in f\}}(\sigma) \cdot \left(\beta_k - (k+1) \|f^\sigma\| \right) \right] \quad (3.10) \\
&= \mathbb{E}_{\sigma \in X(k)} \left[\|f_\sigma\| \left(\beta_k - (k+1) \|f^\sigma\| \right) \right] \\
&= \beta_k \|f\| - \mathbb{E}_{\sigma \in X(k)} (k+1) \|f_\sigma\| \|f^\sigma\|,
\end{aligned}$$

where $\mathbf{1}_{\{\sigma \in f\}}$ denotes the indicator function for the event $\sigma \in f$; the second equality holds since after fixing σ , the condition $\sigma \in f$ in the inner probability is always true for $\sigma \in f$ and always false for $\sigma \notin f$, the inequality follows by (3.9), and the third equality holds since f is a k -cochain, so $\|f_\sigma\| = \mathbf{1}_{\{\sigma \in f\}}(\sigma)$.

On the other hand,

$$\begin{aligned}
\Pr_{\substack{\sigma \in X(k) \\ v \in X_\sigma(0)}} [\sigma \in f \wedge \sigma v \in \delta_1(f)] &= \Pr_{\substack{\tau \in X(k+1) \\ \tau \supset \sigma \in X(k)}} [\tau \in \delta_1(f) \wedge \sigma \in f] \\
&= \Pr_{\tau \in X(k+1)} [\tau \in \delta_1(f)] \Pr_{\substack{\tau \in X(k+1) \\ \tau \supset \sigma \in X(k)}} [\sigma \in f \mid \tau \in \delta_1(f)] \quad (3.11) \\
&= \Pr_{\tau \in X(k+1)} [\tau \in \delta_1(f)] \frac{1}{k+2} \\
&= \frac{1}{k+2} \|\delta_1(f)\|,
\end{aligned}$$

where the first equality holds since both probability distributions are equivalent, the second equality is just conditional probability, and the third equality follows from the definition of $\delta_1(f)$.

Finally, combining (3.10) and (3.11) completes the proof. $\square$

Proof of Lemma 3.7. Let $f \in C^k(X, \mathcal{F})$. For ease of notation, as in the proof of previous lemma, we identify f with its support, i.e., we write $\sigma \in f$ when $f(\sigma) \neq 0$.

By the definition of f^σ ,

$$\mathbb{E}_{\sigma \in X(\ell)} [\|f_\sigma\| \|f^\sigma\|] = \mathbb{E}_{\sigma \in X(\ell)} \mathbb{E}_{u \in \sigma} [\|f_\sigma\| \|(f_{\sigma \setminus u})^u\|] = \mathbb{E}_{\sigma \in X(\ell-1)} \mathbb{E}_{u \in X_\sigma(0)} [\|f_{\sigma u}\| \|(f_\sigma)^u\|]. \quad (3.12)$$

Let $\sigma \in X(\ell-1)$. Define the following two random walks on the $(k-\ell)$ -faces in X_σ :

1. Move from a face $\tau \in X_\sigma(k-\ell)$ to a vertex $u \in \tau$, and then to a face $\tau' \in X_\sigma(k-\ell)$, such that $u \cup \tau' \in X_\sigma(k-\ell+1)$.

2. Move from a face $\tau \in X_\sigma(k - \ell)$ to a vertex $u \in X_\sigma(0)$, such that $u \cup \tau \in X_\sigma(k - \ell + 1)$, and then to a face $\tau' \in X_\sigma(k - \ell)$, such that $\tau' \ni u$.

Random walk #1 can be described by the following operators: $A_1 : \mathbb{R}^{X_\sigma(k-\ell)} \rightarrow \mathbb{R}^{X_\sigma(0)}$ which moves from a $(k - \ell)$ -face in X_σ to a vertex that is contained in it, and $B_1 : \mathbb{R}^{X_\sigma(0)} \rightarrow \mathbb{R}^{X_\sigma(k-\ell)}$ which then moves from a vertex to a $(k - \ell)$ -face that completes it to a $(k - \ell + 1)$ -face. Similarly, random walk #2 can be described by the following operators: $A_2 : \mathbb{R}^{X_\sigma(k-\ell)} \rightarrow \mathbb{R}^{X_\sigma(0)}$ which moves from a $(k - \ell)$ -face in X_σ to a vertex that completes it to a $(k - \ell + 1)$ -face, and $B_2 : \mathbb{R}^{X_\sigma(0)} \rightarrow \mathbb{R}^{X_\sigma(k-\ell)}$ which then moves from a vertex to a $(k - \ell)$ -face that contains it.

Therefore, random walk #1 is given by applying the operator $B_1 A_1 : \mathbb{R}^{X_\sigma(k-\ell)} \rightarrow \mathbb{R}^{X_\sigma(k-\ell)}$ and random walk #2 is given by applying the operator $B_2 A_2 : \mathbb{R}^{X_\sigma(k-\ell)} \rightarrow \mathbb{R}^{X_\sigma(k-\ell)}$. Note that if we consider the random walks which are defined by the same operators but with opposite order, i.e., $A_1 B_1 : \mathbb{R}^{X_\sigma(0)} \rightarrow \mathbb{R}^{X_\sigma(0)}$ and $A_2 B_2 : \mathbb{R}^{X_\sigma(0)} \rightarrow \mathbb{R}^{X_\sigma(0)}$ we get (in both cases) the well-known non-lazy random walk on the underlying graph of X_σ . For instance, $A_1 B_1$ moves from a vertex $u \in X_\sigma(0)$ to a $(k - \ell)$ -face $\tau \in X_\sigma(k - \ell)$ such that $u \cup \tau \in X_\sigma(k - \ell + 1)$ and then to a vertex $v \in \tau$, namely, it moves from u to a random neighbor $v \neq u$.

Since A_1, B_1, A_2 and B_2 are linear operators, the eigenvalues of $B_1 A_1$ and $B_2 A_2$ are equal to the eigenvalues of $A_1 B_1$ and $A_2 B_2$ (up to the multiplicity of the 0 eigenvalue). Therefore, $\lambda_2(B_1 A_1) = \lambda_2(A_1 B_1) = \lambda_2(A_2 B_2) = \lambda_2(B_2 A_2)$, where the equality in the middle is because both $A_1 B_1$ and $A_2 B_2$ describe the same random walk.

Now, consider the undirected graph $G(X_\sigma) = (V(X_\sigma), E(X_\sigma))$, where $V(X_\sigma)$ corresponds to the $(k - \ell)$ -faces of X_σ and $E(X_\sigma)$ corresponds to a step in either of the random walks #1 and #2, i.e., its random walk operator is given by $\frac{1}{2}(B_1 A_1 + B_2 A_2)$. Since X is a λ -one-sided local spectral expander and by the explanation above, it follows that $\lambda_2(\frac{1}{2}(B_1 A_1 + B_2 A_2)) = \lambda$.

It follows that

$$\begin{aligned} \mathbb{E}_{u \in X_\sigma(0)} [\|f_{\sigma u}\| \|(f_\sigma)^u\|] &= \mathbb{E}_{u \in X_\sigma(0)} \left[\Pr_{\substack{\tau \in X_\sigma(k-\ell) \\ u \in \tau}} [\tau \in f_\sigma] \cdot \Pr_{\substack{\tau' \in X_\sigma(k-\ell) \\ u \cup \tau' \in X_\sigma(k-\ell+1)}} [\tau' \in f_\sigma] \right] \\ &= \Pr_{\{\tau, \tau'\} \in E(X_\sigma)} [\tau \in f_\sigma \wedge \tau' \in f_\sigma] \\ &\leq \|f_\sigma\|^2 + \lambda \|f_\sigma\|, \end{aligned} \tag{3.13}$$

where the second equality follows by the probability distribution of $E(X_\sigma)$, and the inequality follows by Lemma 2.1.

Finally, substituting (3.13) in (3.12) completes the proof. $\square$

References

- [ABN23] Anurag Anshu, Nikolas P Breuckmann, and Chinmay Nirkhe. Nlts hamiltonians from good quantum codes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1090–1096, 2023.
- [AC88] Noga Alon and Fan RK Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.
- [DD24] Y. Dikstein and I. Dinur. Coboundary and Cosystolic Expansion Without Dependence on Dimension or Degree. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024)*, pages 62–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

- [DEL⁺22] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 357–374, 2022.
- [DHLV23] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good quantum ldpc codes with linear time decoders. In *Proceedings of the 55th annual ACM symposium on theory of computing*, pages 905–918, 2023.
- [DLV24] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of high-dimensional cubical complexes: with application to quantum locally testable codes. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 379–385. IEEE, 2024.
- [EK16] S. Evra and T. Kaufman. Bounded degree cosystolic expanders of every dimension. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 36–48, 2016.
- [FK22] Uriya A First and Tali Kaufman. On good 2-query locally testable codes from sheaves on high dimensional expanders. *arXiv preprint arXiv:2208.01778*, 2022.
- [FK24] Uriya A First and Tali Kaufman. Cosystolic expansion of sheaves on posets with applications to good 2-query locally testable codes and lifted codes. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1446–1457, 2024.
- [HL22] Max Hopkins and Ting-Chun Lin. Explicit lower bounds against $\omega(n)$ -rounds of sum-of-squares. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 662–673. IEEE, 2022.
- [KKL14] T. Kaufman, D. Kazhdan, and A. Lubotzky. Ramanujan Complexes and Bounded Degree Topological Expanders. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 484–493, 2014.
- [KM21] T. Kaufman and D. Mass. Unique-Neighbor-Like Expansion and Group-Independent Cosystolic Expansion. In *32nd International Symposium on Algorithms and Computation (ISAAC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [KM22] T. Kaufman and D. Mass. Double Balanced Sets in High Dimensional Expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [LZ22] Anthony Leverrier and Gilles Zémor. Quantum tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE, 2022.
- [PK22] P. Panteleev and G. Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, pages 375–388, 2022.