Cosystolic Expansion of Sheaves on Posets with Applications to Good 2-Query Locally Testable Codes and Lifted Codes

Uriya A. First Department of Mathematics University of Haifa Tali Kaufman Computer Science Department Bar-Ilan University

March 28, 2024

Abstract

We show that cosystolic expansion of sheaves on posets can be derived from local expansion conditions of the sheaf and the poset. When the poset at hand is a cell complex — typically a high dimensional expander — a sheaf may be thought of as generalizing coefficient groups used for defining homology and cohomology, by letting the coefficient group vary along the cell complex. Previous works, e.g. [KKL16], [EK16], established local criteria for cosystolic expansion only for *simplicial complexes* and with respect to *constant coefficients*. Our main technical contribution is providing a criterion that is more general in two ways: it applies to *posets* and *sheaves*, respectively.

The importance of working with sheaves on posets (rather than constant coefficients and simplicial complexes) stems from applications to locally testable codes (LTCs). It has been observed [KL14] that cosystolic expansion is related to property testing in the context of simplicial complexes and constant coefficients, but unfortunately, this special case does not give rise to interesting LTCs. We observe that this relation also exists in the much more general setting of sheaves on posets. As the language of sheaves is more expressive, it allows us to put this relation to use. Specifically, we apply our criterion for cosystolic expansion in two ways.

The first application: We show the existence of good 2-query LTCs. These codes are actually related to the good q-query LTCs of $[DEL^+22]$ and [PK22], being the formers' so-called *line codes*, but we get them from a new, more illuminating perspective. By realizing these codes as cycle codes of sheaves over posets, we can derive their good properties directly from our criterion for cosystolic expansion. The local expansion conditions that our criterion requires unfold to the conditions on the "small codes" in $[DEL^+22]$, [PK22], and hence give a conceptual explanation to why conditions such as *agreement testability* are required.

The second application: We show that local testability of a *lifted code* could be derived solely from local conditions, namely from agreement expansion properties of the *local* "small" codes which define it. In the work [DDHRZ20], it was shown that one can obtain local testability of lifted codes from a mixture of local and global conditions, namely from local testability of the local codes and *global* agreement expansion of an auxiliary 3-layer system called a *multilayered* agreement sampler. Our result achieves the same but using genuinely local conditions and a simpler 3-layer structure. It is derived neatly from our local criterion for cosystolic expansion, by interpreting the situation in the language of sheaves on posets.

This is a preliminary version. There may be mild typos and inconsistencies. The final version will appear soon.

Contents

1	Overview	4
	1.1 General	4
	1.2 Posets	4
	1.3 Sheaves	4
	1.4 Sheaf Cohomology.	5
	1.5 Locally Testable Codes From Sheaves on Posets	6
	1.6 A Criterion For Cosystolic Expansion of Sheaves	8
	1.7 First Application: Good 2-Query LTCs	12
	1.8 Second Application: A Local Criterion for Local Testablity of 2-Laver Lifted Codes.	15
	1.9 Conclusion	18
	1.10 Structure of This Paper	18
2	Preliminaries	18
	2.1 Expander Graphs	18
	2.2 Conventions about Codes	20
	2.3 Lifted Codes	20
	2.4 Agreement Testability	21
0		
3	Lifted Codes and Their Line Codes	22
4	Graded Partially Ordered Sets	27
	4.1 Graded Posets	28
	4.2 Weighted Posets	29
	$4.3 \text{Links} \dots \dots$	30
	4.4 Face Counting Constants and Lower-Regular Posets	30
	4.5 Degree and Upper-Regular Posets	33
	4.6 Orientation	33
5	Sheaves on Partially Ordered Sets	34
	5.1 Sheaves on Posets	35
	5.2 Sheaf Cohomology	36
	5.3 Restricting Sheaves to The Links	37
	5.4 Independence of The Orientation for Regular Cell Complexes	37
	5.5 Aside: Constraint Systems	39
6	Locally Testable Codes from Sheaves and Cosheaves	39
	6.1 Cocycle Codes	39
	6.2 Cosystolic Expansion	40
	6.3 Coboundary Expansion	41
	6.4 Independence of The Orientation	43
	6.5 Dual Notions for Cosheaves	43
7	No-Intersection Hypergraphs, Skeleton Expansion, Intersection Profiles	43
	7.1 No-Intersection Hypergraphs	43
	7.2 Skeleton Expansion	45
	7.3 Intersection Profiles	46

8 Mai	in Result:	Simple	Versions
-------	------------	--------	----------

9	2-Query LTCs from Sheaves on Square Complexes	53
	9.1 The Poset	53
	9.2 The Sheaf	54
	9.3 The Code and Its Tester	54
	9.4 Interpretation of Local Expansion Conditions	55
	9.5 Constructing 2-Query LTCs	57
	9.6 Realization as a Line Code	59
10	Local Testability of Two-Layer Lifted Codes	61
	10.1 Two-Laver Lifted Codes	61
	10.2 Subset-Labelled <i>d</i> -Posets	62
	10.3 A Criterion for a 2-Laver Lifted Code to be Locally Testable	63
	10.4 Proof of Theorem 10.8	65
11	Main Result: Technical Version	69
	11.1 Notation	69
	11.2 Main Theorem	70
	11.3 Bounds on Constants	71
	11.4 Proofs of Theorems 8.1 and 8.10	75
12	Proof of Theorem 11.2	79
	12.1 Mock Locally Minimal Cochains	79
	12.2 Reduction to Expansion of Small Locally Minimal Cochains	81
13	Proof of Theorem 12.8	83
10	12.1 Hopeyer Engog	00
	12.2 Exceptional Except	00 04
		04 00
	13.3 Ladders and Terminal Faces	80
	13.4 Completion of The Proof	88
A	Correction Algorithm	90
Bi	ibliography	90

1 Overview

1.1 General

We prove a local criterion for *cosystolic expansion* of sheaves on finite partially ordered sets, called *posets* for short. This extends the reach of known similar criteria established in [KKL16], [EK16] (see also [EK24]), [KM21], [KM22], [DD23] which (in our terminology) apply only to *constant* sheaves on *simplicial complexes*.

Criteria for establishing cosystolic expansion are motivated by applications to locally testable codes (LTCs). The relation between cosystolic expansion and property testing was first observed in [KL14], in the context of constant sheaves on simplicial complexes, and implicitly for general sheaves on posets in [PK22]. We use the expressive language of sheaves over posets and our criterion for cosystolic expansion of sheaves over posets to get good 2-query locally testable codes (while prior works provided good locally testable codes that use many queries) and to get a genuinely local criterion for testability of lifted codes (while prior works used a non-trivial mixture of local and global conditions to derive testability of lifted codes).

Our main result may also be seen a unifying mechanism through which one can recover many known results about cosystolic expansion and testability. For example, it recovers in part the main results of [KO21] and [DLV24].

1.2 Posets

A poset is a finite set X endowed with a transitive anti-reflexive relation $\langle \cdot, \operatorname{For} x, y \in X$, we write $x \leq y$ to denote that x < y or x = y. The posets that we consider in this work will always be equipped with a *dimension function* (also called a *rank function*) dim : $X \to \mathbb{Z}$ which is required to satisfy dim $x < \dim y$ whenever x < y and dim $x + 1 = \dim y$ if in addition no elements of X lie strictly between x and y. See §4.1 for further details.

Our main example of a poset will be the poset of faces of a regular cell complex¹ together with the dimension function assigning every face its usual dimension. (There is also a single empty face of dimension -1.) This includes simplicial complexes and cube complexes. Another example of a poset with a dimension function is the affine Grassmannian.

Following the notation for simplicial complexes, given a poset X, we write X(i) for the elements of X of dimension i and call such elements *i*-faces of X. A poset X is *pure of dimension* d if it is nonempty and each of its faces is contained in a d-face. It is called a d-poset if it moreover has a a unique (-1)-face, denoted \emptyset_X , which is a face of every other face in X. Examples of d-posets include pure d-dimensional simplicial and cubical complexes, and the poset of affine spaces in \mathbb{F}^n (\mathbb{F} is a finite field) of dimension d or less plus the empty set. The degree of a d-poset X is the largest possible number of faces containing a 0-face.

All our posets will carry a *weight function* and an *orientation*, which we suppress in this overview for the sake of simplicity. For details, see §4.2 and §4.6.

1.3 Sheaves

Broadly speaking, a sheaf is a layer of linear-algebra data put on top of a poset. When the poset comes from a geometric source, a sheaf on it may also be viewed as a generalization of the group of coefficients that is used in the definition of homology and cohomology. We shall survey the history of sheaves after we present them. For simplicity, we only consider here sheaves of \mathbb{F}_2 -vector spaces — called \mathbb{F}_2 -sheaves later in the paper — and call them sheaves for brevity.

¹Also called a regular CW complex.

A sheaf \mathcal{F} on a poset X consists of the following data:

- (1) an \mathbb{F}_2 -vector space $\mathcal{F}(x)$ for every face $x \in X$;
- (2) an \mathbb{F}_2 -linear map $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}} : \mathcal{F}(x) \to \mathcal{F}(y)$ for every $x, y \in X$ with x < y;

subject to the requirement $\operatorname{res}_{z \leftarrow y}^{\mathcal{F}} \circ \operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{res}_{z \leftarrow x}^{\mathcal{F}}$ whenever $x < y < z^2$. The maps $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ are called *restriction maps*; the superscript \mathcal{F} will be dropped when it is clear from context.

By reversing the direction of the restriction maps one gets the notion of a *cosheaf* (called a *local* system in some works). A cosheaf on X is essentially the same thing as a sheaf on the opposite poset of X.

Here is a simple example of a sheaf on X: Take an \mathbb{F}_2 -vector space V and define $\mathcal{F}(x) = V$ and $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{id}_X$ for all x and y. This sheaf is denoted V_X . Sheaves of this form (up to isomorphism) are called *constant*.

If X is a 1-poset, then one can get a sheaf \mathcal{F} on X by setting $\mathcal{F}(\emptyset_X) = 0$ and $\operatorname{res}_{x \leftarrow \emptyset}^{\mathcal{F}} = 0$ for every $x \in X - \{\emptyset_X\}$, and choosing the remaining spaces $\mathcal{F}(x)$ and restrictions map $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ arbitrarily. The condition $\operatorname{res}_{z \leftarrow y}^{\mathcal{F}} \circ \operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{res}_{z \leftarrow x}^{\mathcal{F}}$ holds automatically for x < y < z, because we must have $x = \emptyset_X$.

Further examples will be given later in this work.

Brief History and Related Notions. In topology and algebraic geometry, sheaves are defined over *topological spaces*, and this is the common definition in the literature. They were studied since the 1950s and their definition is more involved, e.g., see [Ive86], [MLM94, Chapter II].

The sheaves defined in the paper may be seen as discrete, elementary versions of sheaves on topological spaces. When X is the poset of a cell complex, they are known as *cellular sheaves*. They were first considered by Shepard [She85], and their theory was further developed by Curry [Cur14], who also considered cosheaves. A more concise treatment (over regular cell complexes) appears in [HG19]. The definition of sheaves on posets given here is a natural generalization of (cellular) sheaves on cell complexes, and was briefly considered in [Cur14, §4.2.2].

We remark that despite the differences between sheaves on posets and sheaves on topological spaces, the former is actually a special case of the latter [Cur14, §4.2].

The special case of sheaves on graphs was considered independently in several other works, e.g., the local systems of [JL97, §2]. The sheaves on graphs considered by Friedman [Fri15] are cosheaves in our notation.

1.4 Sheaf Cohomology.

Let X be a cell complex and let $i \in \mathbb{Z}$. Recall that the space of *i*-cochains on X (with coefficients in \mathbb{F}_2) is $C^i = C^i(X, \mathbb{F}_2) := \mathbb{F}_2^{X(i)}$. One then defines the coboundary maps $d_i : C^i \to C^{i+1}$ by

$$(d_i f)(y) = \sum_{x \in y(i)} f(x) \qquad \forall f \in C^i, \ y \in X(i+1),$$

$$(1.1)$$

where here, we wrote y(i) for the set of *i*-faces of the face y. It is well-known that $d_i \circ d_{i-1} = 0$, and so $Z^i = Z^i(X, \mathbb{F}_2) := \ker d_i$ contains $B^i = B^i(X, \mathbb{F}_2) := \operatorname{im} d_{i-1}$. The spaces Z^i , B^i and the quotient $\mathrm{H}^i(X, \mathbb{F}_2) := Z^i/B^i$ are the \mathbb{F}_2 -spaces of *i*-cocycles, *i*-coboundaries and *i*-th cohomology of X, respectively. They are all well-studied.

²In this case, it is convenient to define $\operatorname{res}_{x \leftarrow x}^{\mathcal{F}} = \operatorname{id}_{\mathcal{F}(x)}$ so that $\operatorname{res}_{z \leftarrow y}^{\mathcal{F}} \circ \operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{res}_{z \leftarrow x}^{\mathcal{F}}$ holds whenever $x \leq y \leq z$.

In the same manner, with every sheaf \mathcal{F} on a (graded, oriented) poset X, we can associate \mathbb{F}_2 spaces of cochains, cocycles, coboundaries and cohomology: First, set $C^i = C^i(X, \mathcal{F}) = \prod_{x \in X(i)} \mathcal{F}(x)$.
That is, elements $f \in C^i$ are collections $\{f(x)\}_{x \in X(i)}$ with $f(x) \in \mathcal{F}(x)$ for every $x \in X(i)$. The *i*-th coboundary map is defined as in (1.1), but using the restriction maps of \mathcal{F} :³

$$(d_i f)(y) = \sum_{x \in y(i)} \operatorname{res}_{y \leftarrow x} f(x) \in \mathcal{F}(y) \qquad \forall f \in C^i, \ y \in X(i+1).$$
(1.2)

The spaces $Z^i = Z^i(X, \mathcal{F})$, $B^i = B^i(X, \mathcal{F})$ and $H^i(X, \mathcal{F})$ are then defined to be ker d_i , im d_{i-1} and ker $d_i / \text{ im } d_{i-1}$, respectively.⁴

Observe that if \mathcal{F} is the constant sheaf $(\mathbb{F}_2)_X$, then $C^i(X, \mathcal{F})$ is just $C^i(X, \mathbb{F}_2)$ and the coboundary maps agree.

1.5 Locally Testable Codes From Sheaves on Posets

Our interest in sheaves on posets and their cohomology is motivated by the fact that they give rise to locally testable codes (LTCs).

Locally Testable Codes. Let Σ be a finite alphabet and let $C \subseteq \Sigma^n$ be a code with block length n. We write $\delta(C)$ and r(C) for the relative distance and rate of C, respectively. As usual, when $C \subseteq \Sigma^n$ ranges across a family $\{C_i \subseteq \Sigma^{n_i}\}_{i \ge 0}$, we say that C is good if its relative distance and rate are bounded away from 0.

Recall that a *tester* for $C \subseteq \Sigma^n$ is a randomized algorithm T which, given access to a word $f \in \Sigma^n$, can decide with high probability whether it is close to a codeword or not by querying just a few (i.e. O(1)) of its letters. Formally, a q-query tester T may probe at most q letters from the input f, and must accept all words $f \in C$. The tester T has soundness μ ($\mu \ge 0$) if for every $f \in \Sigma^n$,

$$\Pr(T \text{ rejects } f) \ge \mu \cdot \operatorname{dist}(f, C).$$

Here, dist (\cdot, \cdot) is the normalized Hamming distance in Σ^n . A locally testable code (LTC) is a family of codes-with-testers with block length tending to ∞ such that all the testers have the same query size q and the same positive soundness μ . See [Gol11] for a survey.

The question of whether there exists good LTCs was open until it was recently answered on the positive by Dinur, Evra, Livne, Lubotzky and Mozes [DEL⁺22] and Panteleev and Kalachev [PK22] (independently); see also [LZ22]. The many works on the subject predating this result are surveyed in [DEL⁺22, §1.2].

Cocycle Codes and Cosystolic Expansion. Let \mathcal{F} be a sheaf on a poset X and let $i \in \mathbb{Z}$. Suppose further that there exists an \mathbb{F}_2 -vector space Σ such that $\mathcal{F}(x) = \Sigma$ for every $x \in X(i)$. Then $C^i = C^i(X, \mathcal{F}) = \Sigma^{X(i)}$, and so we may view $Z^i = Z^i(X, \mathcal{F})$ as a code inside $\Sigma^{X(i)}$. We call Z^i the *i*-cocycle code of (X, \mathcal{F}) . The *i*-cocycle code $Z^i \subseteq \Sigma^{X(i)}$ also has a natural tester: Given $f \in \Sigma^{X(i)}$, choose $y \in X(i+1)$ uniformly at random, probe f(x) for every $x \in y(i)$ and accept f if and only if $d_i f(y) = \sum_{x \in y(i)} \operatorname{res}_{y \leftarrow x} f(x) = 0$ (cf. (1.2)). The query size of the this tester is the largest number of *i*-faces that an (i + 1)-face of X can have. For example, when i = 0 and X is a regular cell complex, the natural tester probes only 2 letters.

³When X is not a cell complex or \mathbb{F}_2 is replaced with a field of characteristic not 2, one needs to introduce signs into this formula, see §5.2.

⁴Caution: At this level of generality, one can have $\mathrm{H}^{i}(X, \mathcal{F}) \neq 0$ for i < 0. Also, for a general poset, $\{\mathrm{H}^{i}(X, -)\}_{i \geq 0}$ are in general not the *right derived functions* of $\mathrm{H}^{0}(X, -)$.

Write $\|\cdot\|$ for the normalized Hamming norm on $C^i = \Sigma^{X(i)}$ or $C^{i+1} = \prod_{y \in X(i+1)} \mathcal{F}(y)$. Then the natural tester of Z^i has soundness $\mu \ge 0$ if and only if

$$\|d_i f\| \ge \mu \operatorname{dist}(f, Z^i) \qquad \forall f \in C^i.$$
(1.3)

This condition may also be regarded as an expansion condition for *i*-cochains, and indeed, the *i*-cosystolic expansion of \mathcal{F} , denoted

 $\csc_i(X, \mathcal{F})$

is defined to be the supremum of the set of $\mu \ge 0$ for which (1.3) holds.⁵ Note that this makes sense even without requiring that $\mathcal{F}(x) = \Sigma$ for every $x \in X(i)$.

Observe further that the relative distance of Z^i is the largest $\delta \ge 0$ such that

$$||f|| \ge \delta \qquad \forall f \in Z^i - \{0\}$$

Again, this may be viewed as condition on the expansion of *i*-cochains. For a general sheaf \mathcal{F} on X, we define the *i*-cocycle distance of \mathcal{F} to be⁶

$$\operatorname{ccd}_i(X,\mathcal{F}) = \min\{\|f\| \mid f \in Z^i - B^i\}.$$

The reason why we let f range on $Z^i - B^i$ and not on $Z^i - \{0\}$ is because B^i typically contains vectors of small support (e.g. the coboundary of a small (i - 1)-cochain). However, when $B^i = 0$, we have $\delta(Z^i) = \operatorname{ccd}_i(X, \mathcal{F})$.

Following [KKL16], [EK24] and similar sources, we say that (X, \mathcal{F}) is an (μ, δ) -cosystolic expander in dimension i if $\operatorname{cse}_i(X, \mathcal{F}) \geq \mu$ and $\operatorname{ccd}_i(X, \mathcal{F}) \geq \delta$.

To conclude, provided that $B^i = 0$, the *i*-cocycle code $Z^i = Z^i(X, \mathcal{F}) \subseteq \Sigma^{X(i)}$ is locally testable and has linear distance if and only if (X, \mathcal{F}) is an (μ, δ) -cosystolic expansion in dimension *i* for $\mu, \delta > 0$.

Remark 1.1. Cosheaves on posets similarly give rise to *cycle* codes, and their distance and testability are governed by the *systolic expanion* and *cycle distance* of the cosheaf at hand. This is completely dual to the case of sheaves. For example, the famous *expander codes* of [SS96] may be realized as 1-cycle codes on graphs [Mes18].

Coboundary Expansion Coboundary expansion is a stronger version of cosystolic expansion that will be needed to state our main result. Given a sheaf \mathcal{F} on X, its *i*-coboundary expansion, denoted

$$\operatorname{cbe}_i(X,\mathcal{F})$$

is the supremum of the set of $\mu \geq 0$ such that

$$||d_i f|| \ge \mu \operatorname{dist}(f, B^i) \qquad \forall f \in C^i.$$

In the context of *i*-cocycle codes, $cbe_i(X, \mathcal{F})$ is the soundness of the natural tester of Z^i , but when used a tester for the smaller code B^i .

⁵Actually, the definition of cosystolic expansion involves the weight function on X, so $\operatorname{cse}_i(X, \mathcal{F})$ and the soundness of the tester of Z_i are the same only up to a constant. See Lemma 6.2.

⁶Again, the actual definition of $\operatorname{ccd}_i(X, \mathcal{F})$ involves the weights on X.

Brief History of Coboundary and Cosystolic Expansion. Coboundary expansion originated in the works of Linial–Meshulam [LM06] and Meshulam–Wallach [MW09] on the cohomology of random simplicial complexes, and the work of Gromov [Gro10] on the minimal amount of overlapping forced by mapping a simplicial complex to \mathbb{R}^n . These works did not mention sheaves explicitly, and in our terminology, only considered the case of constant sheaves on simplicial complexes. Within this restricted setting, cosystolic expansion was developed in [DKW18], [KKL16], [EK16] as a relaxed version of coboundary expansion meant to extend the reach of Gromov's methods. The first connections between cosystolic expansion and property testing were observed and studied in [KL14].

1.6 A Criterion For Cosystolic Expansion of Sheaves

Our main result is a criterion for bounding the *i*-cosystolic expansion and *i*-cocycle distance of a sheaf by means of mostly-local expansion conditions.

To state it, we need four more pieces of notation. Let \mathcal{F} be a sheaf on a *d*-poset X.

Lower Irregularity. For integers $-1 \leq i < j < k \leq d$, let $F_{i,j,k}^{\max}$ (resp. $F_{i,j,k}^{\min}$) denote the maximum (resp. minimum) possible number of *j*-faces lying between an *i*-face and a *k*-face that are incident. The ratio $L_{i,j,k}(X) := F_{i,j,k}^{\max}/F_{i,j,k}^{\min}$ is the (i, j, k)-lower irregularity of X. The maximum of $L_{i,j,k}(X)$ over all i, j, k is called the *lower irregularity* of X and denoted L(X). For example, simplicial complexes and cube complexes have the lowest possible irregularity, namely 1.

Links. Let $z \in X$. The link of X at z is $X_z := \{x \in X : x \ge z\}$ together with partial order inherited from X and the dimension function \dim_z given by $\dim_z(x) = \dim x - \dim z - 1$. Note that X_{\emptyset_X} is just X. The sheaf \mathcal{F} restricts to a sheaf \mathcal{F}_z on X_z defined by $\mathcal{F}_z(x) = \mathcal{F}(x)$ and $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}_z} = \operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ for all $x, y \in X$.

No-Intersection Graph. Let i, j, k be integers with $-1 \leq i, j \leq k$. The (i, j, k)-no-intersection graph of X, denoted NIG^{*i*,*j*,*k*}(X), is a graph with vertex set $X(i) \cup X(j)$ (the vertex set is just X(i) if i = j), where for every triple $(x, y, z) \in X(i) \times X(j) \times X(k)$ with $x \neq y, x, y \leq z$ and $\inf\{x, y\} = \emptyset_X$, one adds an edge between x and y.⁷

For example, if X is a regular cell complex, then $\text{NIG}^{0,0,1}(X)$ is just the underlying graph of X, denoted Gr(X). Also, if X is a cube complex, then $\text{NIG}^{1,1,2}(X)$ is the graph whose vertices are the edges of X and in which two edges (viewed as vertices of the graph) are connected when they are the opposite sides of a square in X.

Skeleton Expanders. Let $\alpha, \beta \ge 0$. A weighted graph (G, w) is called an (α, β) -skeleton expander if for every set of vertices $A \subseteq G(0)$, we have $w(E(A)) \le \alpha w(A) + \beta w(A)^2$.

For example, if G is a regular graph, w assigns every vertex (resp. edge) the weight $\frac{1}{|G(0)|}$ (resp. $\frac{1}{|G(1)|}$) and the second largest normalized eigenvalue of G is λ , then G is a $(\lambda, 1)$ -skeleton expander (Proposition 2.2).

Given $i, j, k \in \mathbb{Z}$ as above and $z \in X$ with $\ell := \dim z < i$, it will be convenient to write $\operatorname{NIG}_{z}^{i,j,k}(X)$ for the graph $\operatorname{NIG}^{i-\ell-1,j-\ell-1,k-\ell-1}(X_z)$. Our main result states the following:

Theorem 1.2 (Simplified; see Theorem 8.1). For every $k \in \mathbb{N}$, $F \in \mathbb{N}$, $L \in [1, \infty)$ and $B \in \mathbb{R}_+$ there are constants $K, K' \in (0, 1]$ such that the following hold: Let X be a d-poset $(d \ge k+2)$ such that $L(X) \le L$ and every (k+2)-face of X has at most F-subfaces, let \mathcal{F} be a sheaf on X and let $\varepsilon \in (0, 1]$. Suppose that:

⁷The implicit weight function on X induces a weight function on the no-intersection graph; see Section 7.

- (1a) $\operatorname{cbe}_{k-\dim z-1}(X_z, \mathcal{F}_z) \ge \varepsilon$ for every $z \in X(0) \cup \cdots \cup X(k)$;
- (1b) $\operatorname{cbe}_{k-\dim z}(X_z, \mathcal{F}_z) \ge \varepsilon$ for every $z \in X(0) \cup \cdots \cup X(k+1)$;
- (2) $\operatorname{NIG}_{z}^{i,j,t}(X)$ is a $((K\varepsilon)^{2^{k-i}}, B)$ -skeleton expander for every $z \in X(-1) \cup \cdots \cup X(k)$ and i, j, t with dim $z < i \leq j < t \leq k+2$.

Then $\operatorname{cse}_k(\mathcal{F}) \ge K'(K\varepsilon)^{2^{k+2}-1}$ and $\operatorname{ccd}_k(\mathcal{F}) \ge K'(K\varepsilon)^{2^{k+1}-1}$.

We encourage the reader to think of F, L, B and ε as constant, and of X and \mathcal{F} as varying. In typical applications of Theorem 1.2, as the degree of X grows, F, L, B and ε will remain constant while the skeleton expansion of the $\operatorname{NIG}_{z}^{i,j,t}(X)$ will tend to (0,c) for some constant c > 0. Thus, once the degree of X is large enough (but constant), all three conditions (1a), (1b) and (2) will be satisfied.

A few remarks are now in order.

First, if sheaf \mathcal{F} from the theorem also satisfies $\mathcal{F}(x) = 0$ for all $x \in X(k-1)$ (so that $B^k = 0$) and $\mathcal{F}(x) = \Sigma$ for all $x \in X(k)$, then the theorem says that (up to scaling caused by non-uniform weights) the k-cocycle code $Z^k(X, \mathcal{F}) \subseteq C^k = \Sigma^{X(k)}$ has relative distance $\Theta(\varepsilon^{2^{k+1}-1})$ and its natural tester has soundness $\Theta(\varepsilon^{2^{k+2}-1})$. Moreover, in this case Z^k has a linear-time decoding algorithm able to correct words that are $\Theta(\varepsilon^{2^{k+2}-1})$ -close to Z^k ; see Corollary 8.5.

Second, assumption (2) in the theorem can often be relaxed; it is in general not-necessary to bound the skeleton expansion of *all* the no-intersection graphs $\text{NIG}_z^{i,j,t}(X)$. Here are three such notable examples:

- When X is a simplicial complex, we only need (2) to hold for the graphs $\operatorname{NIG}_{z}^{i,i,i+1}(X)$ with $\dim z = i 1$, i.e., the underlying graph of every X_z with $z \in X(-1) \cup \cdots \cup X(k)$.
- When X is a cube complex, it is enough that (2) holds for every $\operatorname{NIG}_{z}^{i,i,i+1}(X) = \operatorname{Gr}(X_z)$ as in the simplicial case and also for the graphs $\operatorname{NIG}^{1,1,2}(X), \ldots, \operatorname{NIG}^{k+1,k+1,k+2}(X)$. The graph $\operatorname{NIG}^{i,i,i+1}(X)$ is obtained from X by taking the *i*-cubes as vertices and connecting two *i*-cubes by an edge whenever they are the opposite sides of an (i + 1)-cube.
- When k = 0, we need to consider in (2) only the graphs $\operatorname{NIG}^{0,0,1}(X) = \operatorname{Gr}(X)$, $\operatorname{NIG}^{1,1,2}(X)$ and $\operatorname{NIG}^{1,1,2}_v(X)$ for every $v \in X(0)$.

In the general case, the graphs that we need to consider in (2) are specified by an *intersection* profile for the poset X — a novel notion that we introduce in Section 7.

Third, assumptions (1a), (1b) and assumption (2) in the case $z \neq \emptyset_X$ are *local* in the sense that they care only about the structure of X_z and \mathcal{F}_z for $\emptyset \neq z \in X$ and not about the global structure of X and \mathcal{F} . Thus, Theorem 1.2 may be informally summarized as: If

- \mathcal{F}_z is a good coboundary expander for every $z \neq \emptyset$ (" \mathcal{F} has good local coboundary expansion"),
- $\operatorname{NIG}_{z}^{i,j,t}(X)$ is a good skeleton expander for all $z \neq \emptyset$, i, j, t ("X is locally expanding") and
- $\operatorname{NIG}^{i,j,t}(X)$ is a good skeleton expander for all i, j, t ("X is globally expanding"),

then \mathcal{F} is a good cosystolic expander in dimension k (a global condition). For special X, we can make Theorem 1.2 into a purely local criterion for cosystolic expansion. For example, if X is a simplicial complex, then by our previous remark, the only global expansion condition that X needs to satisfy is that Gr(X) is good skeleton expander, and this can be deduced from expansion of the proper links of X by Oppenheim's Trickling Down Theorem [Opp15, Thm. 1.4]. The Trickling Down Theorem was extended to some non-simplicial posets in [KT23], so with more work, there may likely be a purely local criteria for cosystolic expansion for sheaves on such posets.

Finally, we actually prove a more flexible and technical version of Theorem 1.2 — Theorem 11.2. In the special case of 0-cosystolic expansion, this stronger version admits a neat and accessible formulation which we find useful to state here explicitly.

Theorem 1.3 (Criterion for 0-Cosystolic Expansion; see Theorem 8.10). For every $F \in \mathbb{N}$ and $L \in [1, \infty)$, there are constants E, E', E'', E''' > 0 such that the following hold: Let X be a d-poset⁸ $(d \geq 2)$ with $L(X) \leq L$ and such that every 2-face of X contains at most F subfaces, and let \mathcal{F} be a sheaf on X. Let $\varepsilon, \alpha_0, \beta_0, \alpha_{-1}, \beta_{-1}, \alpha_{||}, \beta_{||} \in \mathbb{R}_+$ and suppose that:

- (1a) $\operatorname{cbe}_{-1}(X_z, \mathcal{F}_z) \ge \varepsilon$ for every $z \in X(0) \cup X(1)$;
- (1b) $\operatorname{cbe}_0(X_v, \mathcal{F}_v) \geq \varepsilon$ for every $v \in X(0)$;
- (2a) NIG^{0,0,1}(X_v) is an (α_0, β_0)-skeleton expander for all $v \in X(0)$;
- (2b) NIG^{0,0,1}(X) is an $(\alpha_{-1}, \beta_{-1})$ -skeleton expander;
- (2c) NIG^{1,1,2}(X) is an $(\alpha_{||}, \beta_{||})$ -skeleton expander.

Suppose further that

$$\alpha_{-1} < E\varepsilon$$

and one can find $h_{-1}, h_0, h_{||} \in (0, 1]$ satisfying the inequality

$$(\alpha_0 + \beta_0 h_0) + (\alpha_{||} + \beta_{||} h_{||}) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} \le E' \varepsilon$$

 $\begin{aligned} Then \ \mathrm{ccd}_0(X,\mathcal{F}) &\geq \frac{E'''(E\varepsilon - \alpha_{-1})}{\beta_{-1}} \ and \ \mathrm{cse}_0(X,\mathcal{F}) \geq \frac{E''}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}. \ When \ X \ is \ a \ simplicial \ (resp. \ cube) \\ complex, \ we \ can \ take \ E &= E''' = 1, \ E' = \frac{1}{12} \ (resp. \ E' = \frac{1}{16}) \ and \ E'' = \frac{1}{2}. \end{aligned}$

The constants E, E', E'', E''' are explicit and may be found in Table 1. Theorem 1.3 is more general than Theorem 1.2, because it can be applied with $\beta_0, \beta_{-1}, \beta_{||}$ arbitrarily large.

Relation to Other Works. Local criteria for establishing cosystolic expansion of constant sheaves on simplicial complexes appeared in [KKL16] (dim $X \leq 3$, $\mathcal{F} = (\mathbb{F}_2)_X$), [EK16] ($\mathcal{F} = (\mathbb{F}_2)_X$, see also [EK24]), [KM21] (any constant sheaf), [DD23] (same). Our Theorem 1.2 is an improvement of these results in two ways. First, it applies to general posets, and second, it applies to all sheaves. It further improves [KKL16], [EK16], [KM21] by providing a lower bound on the cosystolic expansion which is independent of the dimension and the degree of X. On the other hand, in the simplicial case, [DD23] gives a better lower bound on the k-cosystolic expansion and [KM22, Thm. 7] establishes a better lower bound on the k-cocycle distance in the case $\mathcal{F} = (\mathbb{F}_2)_X$; both bounds are $\Theta(\varepsilon^{k+1})$. The reason for this difference seems to stem from the fact that the arguments in [KM22], [DD23] make critical use of the fact that the sheaf is constant and the poset is simplicial. A new feature of our result that did not appear in previous works on cosystolic expansion is the use of no-intersection graphs, which turned out to be necessary in treating the non-simplicial case.

 $^{^{8}}$ Recall that all our posets carry a weight function and an orientation and those should be taken into account; see §4.2, §4.6 and Definition 7.3.

No-intersection graphs were already studied in [KO21] in the context of amplified testability, but not as a way to get cosystolic expansion.

A local criterion for establishing cosystolic expansion of *general* sheaves on *simplicial complexes* first appeared in an earlier work of the authors [FK23b, §8] now superseded by this work.

The main result of [DLV24] gives sufficient conditions for certain (non-constant) sheaves on certain cube complexes to have good cosystolic and *systolic* expansion. It is likely that the cosystolic part of this result follows from our Theorem 1.2, possibly with different constants. Indeed, the *two-way robustness* requirement in [DLV24] is essentially the same as saying that the proper links of the sheaved cube complex at hand are good coboundary and boundary expanders, and therefore satisfy assumptions (1a) and (1b) of Theorem 1.2, and the expansion conditions in [DLV24] imply the necessary expansion condition (2) in Theorem 1.2 when X is a cube complex [DLV24, Claim 5.11, Lemma 5.12].

Our main theorem is also related to the main result of [KO21] (see also [KO22]). There, the authors consider codes modelled over 2-layer expanding systems, and show that these codes are locally testable if the underlying system satisfies some global and local expansion conditions. In our language, the 2-layer system is a 2-poset X, and the code modelled over it, while not strictly being a sheaf, is very similar to an \mathbb{F} -sheaf \mathcal{F} on X in which $\mathcal{F}(x) = \mathbb{F}$ for all $x \in X$ and all the restriction maps are isomorphisms. The conditions under which [KO21, Thm. 1.17] applies resemble assumptions (1a)–(2c) of Theorem 1.3 and suggest that conditions of this flavor may be necessary in general. In fact, the codes modelled on 2-layer systems of [KO21] are examples of constraint systems on a poset — a variation on the definition of a sheaf to which our main result still applies, see 5.5 —, and so our Theorem 1.2 recovers the testability part of [KO21, Thm. 1.17]. On the other hand, [KO21] establishes a stronger kind of testability called *amplified local testability*.

Finally, we note that Kaufman and Tessler [KT23] extended Garland's Method and Oppenheim's Trickling Down Theorem, which are examples of local criteria for other types of expansion, from simplicial complexes to general posets.

About The Proof. The proof of Theorem 1.2 is loosely based on the *fat machinery* method of [KKL16], [EK24], called *heavy machinery* here. Broadly speaking, the idea is to first reduce the problem into showing that a *locally minimal k*-cochain $f \in C^k := C^k(X, \mathcal{F})$ with small support must expand under $d_k : C^k \to C^{k+1}$. Being locally minimal means that for every $z \in X$ of dimension $i \geq 0$, the restriction $f_z := f|_{X_z(k-i-1)} \in C^{k-i-1}(X_z, \mathcal{F}_z)$ satisfies $||f_z|| = \text{dist}(f_z, B^{k-i-1}(X_z, \mathcal{F}_z))$. Thus, if we assume that $\text{cbe}_{k-i-1}(X_z, \mathcal{F}_z) \geq \varepsilon$, then we would know that $||d_{k-i-1}f_z|| \geq \varepsilon ||f_z||$. One would like to take advantage of this to show that $||d_kf||$ is at least proportional to ||f||, but in general, $d_{k-i-1}f_z$ and $(d_kf)_z$ may differ. The heavy machinery is a method of keeping track of faces z such that $d_{k-i-1}f_z = (d_kf)_z$, ultimately showing that the contribution of faces for which this equality fails is negligible.

We follow this general strategy, but introduce many new ingredients. For example, we use information from no-intersection graphs (that are not underlying hypergraphs of links), which is necessary to make the argument work for general posets, and introduce *intersection profiles* to keep track of the types of no-intersection graphs that we need. We also introduce *terminal faces* and use them to make the delicate summation process over the faces z above more efficient and streamlined. Furthermore, following ideas from [KO21] and [DD23], we replace locally minimal cochains with a variation — mock q-locally minimal cochains —, which allows us to make the lower bound on $\operatorname{cse}_k(X, \mathcal{F})$ independent of the dimension and the degree of X. Instead, the bound depends on the lower regularity of X (in all dimensions); this dependence was transparent in works concerning with simplicial and cube complexes, because they have lower irregularity 1.

1.7 First Application: Good 2-Query LTCs

As an application of Theorem 1.2 and its finer version Theorem 1.3, we give an example of good 2query LTCs arising from sheaves on square complexes. These codes are in fact the *line codes* of the of good LTCs of [DEL⁺22]. By interpreting these codes as 0-cocycle codes of sheaves, we can apply Theorem 1.3 to neatly deduce that they form a 2-query LTC. This offers a new perspective on the LTCs [DEL⁺22], showing that their testability may be seen as a consequence of cosystolic expansion. It also shows that the *agreement testability* requirement appearing in [DEL⁺22, Theorem 4.5] is actually a manifestation of coboundary expansion. (We remark that while the good LTCs of [DEL⁺22] are related to the good LTCs of [PK22], we do not know how to directly relate our LTCs to those of [PK22].)

We shall first describe our good 2-query LTCs, and after that explain their relation to the LTCs of [DEL⁺22].

The Example. We take our base poset X to be a *left-right Cayley complex*. Let G be a finite group and let $A, B \subseteq G$ be two sets of generators for G such that $A = A^{-1}$, $B = B^{-1}$, $1 \notin A \cup B$ and no element of A is a conjugate of an element of B. Recall that X = Cay(A, G, B) is a square complex with faces determined as follows:

- $X(0) = \{\{g\} \mid g \in G\},\$
- $X(1) = \{\{g, ag\} \mid g \in G, a \in A\} \cup \{\{g, bg\} \mid g \in G, b \in B\},\$
- $X(2) = \{\{g, ag, gb, agb\} \mid g \in G, a \in A, b \in B\}.$

(We also have $X(-1) = \{\emptyset\}$.) Our assumptions imply that for every $\{g\} \in X(0)$ and $e \in X(1)$ containing $\{g\}$, there is a unique $x \in A \cup B$ such that $e = \{g, xg\}$ if $x \in A$ and $e = \{g, gx\}$ if $x \in B$. A similar claim applies to edges and squares.

Let $C_A \subseteq \mathbb{F}_2^A$ and $C_B \subseteq \mathbb{F}_2^B$ be linear codes. It will be convenient to view $\mathbb{F}_2^A \otimes \mathbb{F}_2^B$ as the space $M_{A \times B}(\mathbb{F}_2)$ of matrices with rows indexed by A and columns index by B. Given a $m \in M_{A \times B}(\mathbb{F}_2)$, we write $r_a(m)$ for the *a*-th row of m and $c_b(m)$ for the *b*-th column of m. The tensor code $C_A \otimes C_B \subseteq \mathbb{F}_2^A \otimes \mathbb{F}_2^B = M_{A \times B}(\mathbb{F}_2)$ may now be view as the space of $A \times B$ -matrices m with $r_a(m) \in C_B$ and $c_b(m) \in C_A$ for all $a \in A$ and $b \in B$.

We define a sheaf \mathcal{F} on $X = \operatorname{Cay}(A, G, B)$ as follows:

- $\mathcal{F}(\emptyset) = 0$,
- $\mathcal{F}(\{g\}) = C_A \otimes_{\mathbb{F}} C_B,$
- $\mathcal{F}(\{g, ag\}) = C_B,$
- $\mathcal{F}(\{g,gb\}) = C_A,$
- $\mathcal{F}(\{g, ag, gb, agb\}) = \mathbb{F},$
- $\operatorname{res}_{\{g,ag\}\leftarrow\{g\}} = r_a : C_A \otimes C_B \to C_B,$
- $\operatorname{res}_{\{q,qb\}\leftarrow\{q\}} = c_b : C_A \otimes C_B \to C_A,$
- $\operatorname{res}_{\{g,ag,gb,agb\}\leftarrow\{g,ag\}}: C_B \to \mathbb{F}$ is projection on the *b*-component,
- $\operatorname{res}_{\{q,aq,qb,aqb\}\leftarrow\{q,qb\}}: C_A \to \mathbb{F}$ is projection on the *a*-component,

where here, $g \in G$, $a \in A$, $b \in B$. It is straightforward to check that this is well-defined. Observe further that if we put $\Sigma = C_A \otimes C_B$, then $\mathcal{F}(v) = \Sigma$ for every $v \in X(0)$. We may therefore form the 0-cocycle code $Z^0 = Z^0(X, \mathcal{F}) \subseteq C^0(X, \mathcal{F}) = \Sigma^{X(0)} = \Sigma^G$.

Upon unfolding the definition, one sees that the natural 2-query tester of Z^0 works as follows: Given $f \in \Sigma^G$, it chooses an edge $\{g, h\} \in X(1)$ and probes $f(\{g\})$ and $f(\{h\})$. If h = ag for some $a \in A$, then f is accepted if and only if $r_a(f(g)) = r_{a^{-1}}(f(h))$, and if h = gb for some $b \in B$, then f is accepted if and only if $c_b(f(g)) = c_{b^{-1}}(f(h))$.

We use Theorem 1.3 to give sufficient conditions on $Z^0(X, \mathcal{F})$ to be locally testable and have linear distance. To phrase them, recall [DEL⁺22, Definition 2.8] that the tensor code $C_A \otimes C_B \subseteq$ $M_{A \times B}(\mathbb{F}_2)$ is said to be κ -agreement testable if for all $m_1 \in C_A \otimes \mathbb{F}_2^B$ and $m_2 \in \mathbb{F}_2^A \otimes C_B$, there is $m \in C_1 \otimes C_B$ such that

$$\kappa \cdot \left[\frac{\#\{a \in A : r_a(m_2) \neq r_a(m)\}}{2|A|} + \frac{\#\{b \in B : c_b(m_1) \neq c_b(m)\}}{2|B|} \right] \\ \leq \frac{\#\{(a,b) \in A \times B : (m_1)_{a,b} \neq (m_2)_{a,b}\}}{|A||B|}.$$

Informally, this means that if m_1 and m_2 agree on nearly all entries, there is $m \in C_A \otimes C_B$ which agrees with m_1 and m_2 on nearly all columns and rows, respectively. See [DEL⁺22] for more information and examples.

Theorem 1.4. For every $\varepsilon > 0$ there are (small) real constants $\lambda, \mu, \delta_0, \eta > 0$ such that the following hold: Let G, A, B, X, \mathcal{F} and $Z^0 \subseteq \Sigma^G$ be as above and suppose the following conditions are met:

 $(1a') \ \delta(C_A) \ge \varepsilon,$

 $(1b') \ \delta(C_B) \geq \varepsilon,$

- (1c') $C_A \otimes C_B$ is ε -agreement testable,
- (2') the Cayley graphs $\operatorname{Cay}(A, G)$ and $\operatorname{Cay}(G, B)$ are λ -expanders, i.e., the second largest eigenvalue of their normalized adjacency operator is at most λ .

Then $\delta(Z^0) \geq \delta_0$ and the natural 2-query tester of Z_0 has soundness $\frac{\mu}{|A|+|B|+1}$. Moreover, $r(Z^0) \geq \frac{4r(C_A)r(C_B)-3}{4r(C_A)r(C_B)}$ and Z^0 admits a linear-time decoding algorithm for words that are η -close to Z_0 .

We derive Theorem 1.4 by applying Theorem 1.3 to the X and \mathcal{F} we constructed. The full details are given in Section 9. Briefly, conditions (1a') and (1b') are equivalent to saying that $\operatorname{cse}_{-1}(X_e, \mathcal{F}_e) \geq \varepsilon$ for every $e \in X(1)$, and condition (1c') is equivalent to having $\operatorname{cse}_0(X_e, \mathcal{F}_e) \geq \varepsilon$. ⁹ With a little more work, one further derives from (1a') and (1b') that $\operatorname{cse}_{-1}(\mathcal{F}_v) \geq \varepsilon$, so (1a')–(1c') imply conditions (1a) and (1b) of Theorem 1.3. One can further show that (2') implies that $\operatorname{Gr}(X) = \operatorname{NIG}^{0,0,1}(X)$ is a $(\lambda, 1)$ -skeleton expander and $\operatorname{NIG}^{1,1,2}(X)$ is a $(2\lambda, 4(|A| + |B|))$ -skeleton expander. Moreover, $\operatorname{Gr}(X_v)$ is a (0, 1)-skeleton expander for every $v \in X(0)$, being a complete bipartite graph. Now, one can readily check that the inequalities in Theorem 1.3 are solvable if λ is small enough and deduce Theorem 1.4.

It was observed in [DEL⁺22] that there is $\varepsilon > 0$ and $n_0 \in \mathbb{N}$ such that whenever $|A|, |B| \ge n_0$, there exist codes $C_A \subseteq \mathbb{F}_2^A$ and $C_B \subseteq \mathbb{F}_2^B$ satisfying conditions (1a')-(1c') and also $r(C_A), r(C_B) > \frac{3}{4}$. Let λ be the constant obtained by applying Theorem 1.4 to that ε . It is further known that there is

⁹Checking that this follows readily from the definitions is a recommended exercise.

 $n_1 \ge n_0$ for which there are infinitely many examples $(G_i, A_i, B_i)_{i \in \mathbb{N}}$ of G, A, B as above such that $|A_i| = |B_i| = n_1$ and both $\operatorname{Cay}(A_i, G_i)$ and $\operatorname{Cay}(G_i, B_i)$ are λ -expanders. By applying Theorem 1.4 to the family $(G_i, A_i, B_i)_{i \in \mathbb{N}}$ and suitable codes $C_A, C_B \subseteq \mathbb{F}_2^{n_1}$, we obtain a family of good 2-query LTCs.

Relation Between Lifted Codes and Line Codes. To better describe the relation between our 2-query LTCs and the LTCs of [DEL⁺22], we need to a briefly digress and discuss the relation between lifted codes and their so-called line codes. See Section 3 for an extensive discussion.

Recall that a *lifted code* or a *Tanner code* $C \subseteq \Sigma^n$ is determined by a family of subsets S of $[n] := \{1, \ldots, n\}$ covering [n] and, for every $s \in S$, a code $C_s \subseteq \Sigma^s$. The lifted code that the family $\{C_s\}_{s \in S}$ determines is

$$C = C(\{C_s\}_{s \in S}) := \{f \in \Sigma^n : f|_s \in C_s \text{ for all } s \in S\}.$$

When all the C_s are the same code $D \subseteq \Sigma^m$ (or, more generally, whenever $|C_s| = |D|$ for all $s \in S$), we may further associate with C a code $L \subseteq D^S$ with alphabet D known as its *line code*; it is defined by

$$L = \{ f = (f_s)_{s \in S} \in D^S : f_s|_{s \cap s'} = f_{s'}|_{s \cap s'} \text{ for all } s, s' \in S \}.$$

There is a bijection $f \mapsto (f|_s)_{s \in S} : C \to L$, so both C and L have proportional rates, and under mild assumptions, $\delta(C)$ and $\delta(L)$ are also proportional (Proposition 3.6).

The presentation of C as a lifted code $C = C(\{C_s\}_{s \in S})$ gives rise to a natural tester: Given $f \in \Sigma^n$, choose $s \in S$ uniformly at random and accept f if and only if $f|_s \in C_s$. This tester usually has poor soundness, which is why LTCs are considered difficult to construct. However, we show in Theorem 3.10 that if the line code L of a lifted code C (varying in a family) is 2-query locally testable, then same holds for the original code C with its natural tester. Also, if L has a linear-time decoding algorithm, then the same holds for C (Proposition 3.7).

Relation to [DEL⁺22]. Let G, A, B, X, C_A and C_B be as before. For $g \in G$, we write s(g) for the set of squares containing the 0-face $\{g\}$ of X. There is a bijection from $A \times B$ to s(g) given by sending (a, b) to $\{g, ag, gb, agb\}$, and we use it to identify $\mathbb{F}_2^{s(g)}$ with $\mathbb{F}_2^{A \times B} = \mathbb{F}_2^A \otimes \mathbb{F}_2^B$. Now let C(A, G, B) denote the lifted code $C \subseteq \mathbb{F}_2^{X(2)}$, determined by the sets $\{s(v) \mid v \in X(0)\}$ and the codes

$$C_{s(v)} = C_A \otimes C_B \subseteq \mathbb{F}_2^A \otimes \mathbb{F}_2^B = \mathbb{F}_2^{s(v)}.$$

We endow C(A, G, B) with its natural $|A \times B|$ -query tester. In [DEL⁺22], it was shown that codes C(A, G, B) form an LTC if conditions (1a')–(2') of Theorem 1.4 hold (with different contants) and C_A and C_B have sufficiently large rate.

It is straightforward to see that code $Z^0(X, \mathcal{F})$ considered in Theorem 1.4 is the line code of C(A, G, B). Thus, our earlier discussion implies that we can derive the fact that C(A, G, B) is a good LTC from the fact that $Z^0(X, \mathcal{F})$ is a good LTC.

In fact, the fact that the line code of C(A, G, B) is locally testable is already proved implicity in [DEL⁺22], and the testability of C(A, G, B) is derived from it (look at Algorithm 1 in [DEL⁺22], which is also a correction algorithm for the line code of C(A, G, B)). Therefore, a variant of Theorem 1.4 is already implicit in [DEL⁺22]. Our discussion here is meant to highlight the role of the line code $Z^0(X, \mathcal{F})$ in the proof that C(A, G, B) is locally testable, the fact that the testability of $Z^0(X, \mathcal{F})$ is a consequence of \mathcal{F} being a good cosystolic expander in dimension 0, and that this can be shown using our main theorem.

1.8 Second Application: A Local Criterion for Local Testablity of 2-Layer Lifted Codes

In [DDHRZ20], the authors give a criterion for a lifted code with additional structure to be locally testable. When working in Σ^n , the system of sets used to define the lifted code is required to be embedded in an auxiliary 3-layer system of subsets of $[n] = \{1, \ldots, n\}$, which is required to satisfy some expansion conditions and a global condition on agreement testability.

We apply Theorem 1.2 to give a simpler, purely *local* criterion for establishing the local testability of a lifted code. The "small" codes defining our lifted code are required to be lifted codes themselves; we call this 2-layered structure, defined below, a 2-*layer lifted code*. A third layer is needed to apply our criterion, but not to define the code; it is required in order to be able to talk about agreement testability for the "small" lifted codes defining our global code.

Agreement Testability. The notion of agreement expansion was first considered in [DK17] and studied further in [DD19]. Informally, an agreement expander consists of a collection of subsets S of [n] such that for any finite set Σ and any ensemble of functions $\{f_s : s \to \Sigma\}_{s \in S}$ such that $f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}$ for almost all $s, s' \in S$, there is $g : [n] \to \Sigma$ such that $g|_s = f_s$ for almost all $s \in S$. Here, we will consider a more refined version of this notion where each f_s is required to be in a code $C_s \subseteq \Sigma^s$ and g comes from $C = C(\{C_s\}_{s \in S})$. In the special case of tensor codes, realized as lifted codes (Example 2.6), this already appeared in [DEL⁺22, Dfn. 2.8] under the name agreement testability, which we also use here. We give here a simplified version of the definition and refer to $\S2.4$ for the general definition.

Let S be a collection of subsets of [n] and let $C = C(\{C_s\}_{s \in S}) \subseteq \Sigma^n$ be a lifted code. Suppose further that we are given a collection T of 2-element subsets of S such that $s \cap s' \neq \emptyset$ for every $\{s, s'\} \in T$. The agreement testability of the lifted code C w.r.t. T measures how far is an ensemble of local views $\{f_s \in C_s\}_{s \in S}$ such that $f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}$ for almost all $\{s, s'\} \in T$ from being induced by a single global $g \in C$. Formally, we say that $C = C(\{C_s\}_{s \in S})$ is κ -agreement testable w.r.t. T if for every $(f_s)_{s \in S} \in \prod_{s \in S} C_s$, there is $g \in C$ such that

$$\kappa \cdot \frac{\#\{s \in S : g|_s = f_s\}}{|S|} \le \frac{\#\{\{s, s'\} \in T : f_s|_{s \cap s'} \neq f_{s'}|_{s \cap s'}\}}{|T|}.$$

Two-Layer Lifted Codes. Again, for the sake of simplicity, we give a special case of the general definition, which can be found in §10.1.

Let $n \in \mathbb{N}$ and let Σ be a finite alphabet. A 2-layer lifted code inside Σ^n is a triple $(S, T, \{C_{s,s'}\}_{s,s'})$ consisting of:

- a collection S of subsets of [n];
- a collection T of two-element subsets of S;
- a code $C_{s,s'} \subseteq \Sigma^{s \cap s'}$ for every $\{s, s'\} \in T$;

such that:

- (1) S covers [n],
- (2) for every $s \in S$, the collection $S_s := \{s \cap s' \in S : s \in S \text{ and } \{s, s'\} \in T\}$ covers s, and
- (3) $s \cap s' \neq \emptyset$ for every $\{s, s'\} \in T$.

We then associate with $(S, T, \{C_{s,s'}\}_{s,s'})$ a local lifted code

$$C_s = C(\{C_{s,s'}\}_{s' \in S_s}) \subseteq \Sigma^s$$

for every $s \in S$, and a global lifted code

$$C = C(\{C_s\}_{s \in S}) = C(\{C_{s,s'}\}_{\{s,s'\} \in T}) \subseteq \Sigma^n.$$

The natural tester of C is its natural tester when realized as a lifted code using the codes $\{C_s\}_{s\in S}$.

Local Testability of Two-Layer Lifted Codes. Let $(S, T, \{C_{s,s'}\}_{s,s'})$ by a two-layer lifted code in Σ^n . Our local criterion for local testability of the lifted code $C = C(\{C_s\}_{s \in S})$ requires an additional layer of subsets of [n]. Specifically, suppose that we are further given a collection U of 3-element subsets of S such that:

- (1) $\{s, s'\}, \{s, s''\}, \{s', s''\} \in T$ for every $\{s, s', s''\} \in U$;
- (2) $s \cap s' \cap s'' \neq \emptyset$ for every $\{s, s', s''\} \in U$;
- (3) for every $\{s, s'\} \in T$, the sets of the form $s \cap s' \cap s''$ with $\{s, s', s''\} \in U$ cover $s \cap s'$.

In particular, the union of T, U and $\{\{s\} | s \in S\} \cup \{\emptyset\}$ forms a 2-dimensional simplicial complex denoted X.

Given $s \in S$, let $S_s := \{s \cap s' \in S : s \in S \text{ and } \{s,s'\} \in T\}$ as before, and let T_s denote the set of pairs $\{s \cap s', s \cap s''\}$ where $\{s,s',s''\} \in U$. We say that the system (S,T,U) is *lower regular* if for every $i \in [n]$, the number of $s \in S$ (resp. $\{s,s'\} \in T$, $\{s,s',s''\} \in U$) with $i \in S$ (resp. $i \in s \cap s'$, $i \in s \cap s' \cap s''$) is independent of i. We say that (S,T,U) is upper regular if for every $s \in S$ (resp. $\{s,s'\} \in T$, $\{s,s',s''\} \in U$), the number #s (resp. $\#(s \cap s'), \#(s \cap s' \cap s'')$) is independent of s(resp. $\{s,s'\}, \{s,s',s''\}$).

Theorem 1.5 (Simplified; see Theorem 10.10). There are constants K, K' > 0 such that the following hold: Let $(S, T, \{C_{s,s'}\}_{s,s'})$ be a two-layer lifted code whose alphabet Σ is an \mathbb{F}_2 -vector space and such that every $C_{s,s'}$ is a subspace of $\Sigma^{s \cap s'}$. Let U and X be as above. Suppose that (S, T, U) is both lower and upper regular and satisfies:

(0) For every $s, s' \in S$ and $i \in s \cap s'$, there are $s_0, s_1, \ldots, s_m \in S$ such that $s = s_0, s' = s_m, \{s_0, s_1\}, \ldots, \{s_{m-1}, s_m\} \in T$ and $i \in s_0 \cap \cdots \cap s_m$.

Let $\varepsilon > 0$ and suppose in addition that:

- (1a) $\delta(C_{s,s'}) \ge \varepsilon$ for every $\{s, s'\} \in T$;
- (1b) the lifted code $C_s = C(\{C_{s,s'}\})_{s' \in S_s})$ is ε -agreement testable w.r.t. the set T_s ;
- (2) $\operatorname{Gr}(X_v)$ is a (spectral) $K\varepsilon^2$ -expander for every $v \in X(0)$.

Then the natural tester of the lifted code $C = C(\{C_s\}_{s\in S}) \subseteq \Sigma^n$ has soundness $K'\varepsilon^3$. Moreover, writing $D = \max_{s\in S} |s|$, we have $\delta(C) > \frac{K'\varepsilon}{D}$ and C admits a linear-time decoding algorithm for words that are $\frac{K'\varepsilon^3}{D}$ -close to C.

Observe that assumptions (1a)-(2) are all *local* in the sense that they care only about the local structure of X and about the small codes C_s and $C_{s,s'}$. We actually prove a more general variant of this theorem where no regularity assumptions are necessary and (S, T, U) may be replaced with a general three-layered system of subsets of [n] organized into a pure 2-dimensional regular cell complex; see §10.2 and Theorem 10.8. In this more general setting, one needs to require that the underlying graph of X and its (1, 1, 2)-no-intersection graph are sufficiently good skeleton expanders.

In order to prove Theorem 1.5, we define a sheaf \mathcal{F} on X as follows:

- $\mathcal{F}(\emptyset) = 0$,
- $\mathcal{F}(\{s\}) = C_s$ for all $s \in S$,
- $\mathcal{F}(\{s, s'\}) = C_{s,s'}$ for all $\{s, s'\} \in T$,
- $\mathcal{F}(\{s, s', s''\}) = \Sigma^{s \cap s' \cap s''},$
- $\operatorname{res}_{\{s,s'\} \leftarrow \{s\}}(f) = f|_{s \cap s'},$
- $\operatorname{res}_{\{s,s',s''\}\leftarrow\{s,s'\}}(f) = f|_{s\cap s'\cap s''}.$

Condition (0) of Theorem 1.5 implies that the 0-cocycle code $Z^0 = Z^0(X, \mathcal{F})$ is precisely the line code of the lifted code $C = C(\{C_s\}_{s \in S})$. Thus, as noted earlier in §1.7, in order to prove that Cis locally testable w.r.t. its natural tester, it is enough to show that Z^0 is locally testable. To that end, we apply Theorem 1.2 or Theorem 1.3. The prerequisites of those theorems can be derived from conditions (1a)–(2) and Oppenheim's Trickling Down Theorem [Opp15, Thm. 4.1], thanks to the fact that X is simplicial and (S, T, U) is upper and lower regular.

Comparison with [DDHRZ20] As we noted earlier, [DDHRZ20] also provides a criterion for a lifted code to be locally testable. Both the main result of [DDHRZ20] and our Theorem 1.5 assume that a three-layered system of subsets of [n] is provided, but otherwise, they differ both in the setting and the assumptions. To state these differences, let S, T, U be as above and put $\tilde{T} = \{s \cap s' \mid s, s' \in T\}$ and $\tilde{U} = \{s \cap s' \cap s'' \mid \{s, s', s''\} \in U\}$.¹⁰

In [DDHRZ20], one starts from small codes $C_u \subseteq \Sigma^u$ for every $u \in \tilde{U}$ using which one constructs bigger lifted codes $C_t \in \Sigma^t$ and $C_s \in \Sigma^s$ for every $t \in \tilde{T}$ and $s \in S$. The main result of [DDHRZ20] may now be loosely summarized as saying that $C = C(\{C_u\}_{u \in \tilde{U}}) \subseteq \Sigma^n$ is locally testable when the following conditions are met: (1) each lifted code $C_t \subseteq \Sigma^t$ on the layer \tilde{T} has linear distance, (2) each lifted code $C_s \subseteq \Sigma^s$ on the layer S is locally testable w.r.t. its natural tester, (3) the incidence graph of (\tilde{T}, \tilde{U}) satisfies an expansion condition, and (4) the pair (S, \tilde{T}) satisfies an agreement expansion condition (for δ -ensembles). Note that conditions (3) and (4) concern with the global structure of the collections S, \tilde{T}, \tilde{U} . By contrast, our criterion for local testability (Theorem 1.5) starts with "bigger" small codes $C_t \subseteq \Sigma^t$ on the layer \tilde{T} and replaces requirements (2),(3),(4) with two local requirements which may loosely be summarized as saying that for every $s \in S$, the incidence graph of the $t \in \tilde{T}$ and $u \in \tilde{U}$ contained in s is a good expander, and the lifted code $C_s = C(\{C_t | t \in \tilde{T}, t \subseteq s\}) \subseteq \Sigma^s$ is agreement testable.

That said, Theorem 1.5 applies only when the sets S, \tilde{T}, \tilde{U} may be organized into a 2-dimensional simplicial complex and satisfy some regularity assumptions. In addition, it requires that the alphabet Σ is an \mathbb{F}_2 -vector space and the small codes $C_{s,s'}$ are \mathbb{F}_2 -linear. No such requirements are

¹⁰In [DDHRZ20], the collections $S, \tilde{T}, \tilde{U}, [n]$ are denoted S, K, T, V.

imposed in [DDHRZ20]. As we noted earlier, our approach gives a more general version of Theorem 1.5 applying to more general 3-layer collections of subsets of [n]. It requires some additional global expansion assumptions, but no global agreement testablity as in [DDHRZ20].

1.9 Conclusion

By using the relation between lifted codes and their line codes, we can translate questions about local testability to statements about cosystolic expansion of sheaves. Our main theorem (Theorem 1.2) serves as a powerful tool to establish the desired cosystolic expansion.

1.10 Structure of This Paper

The remainder of this paper is structured as follows: Section 2 is preliminary and recalls relevant facts about expander graphs and error correcting codes, setting some notation along the way. Section 3 concerns with line codes of lifted codes; in this section, we relate the rate, distance, testability and decodability of a lifted code and its line code. In Section 4, we recall posets and introduce additional structure on them that will be needed for this work, e.g., weight functions and orientation. Sheaves on posets and their cohomology are then discussed in Section 5. The subject matter of Section 6 is cocycle codes of sheaves and their relation to cosystolic expansion. Section 7 concerns with no-intersection (hyper)graphs and their skeleton expansion, and introduces the notion of an intersection profile. We then give simplified versions of our main result in Section 8. The results of Section 8 are applied in Section 9 to give examples of good 2-query LTCs, and in Section 10 to give a local criterion for a two-layer lifted code to be locally testable. In Section 11, we formulate our main result in its general form (Theorem 11.2) and derive the simpler versions of Section 8 from it. The remaining Sections 12 and 13 are dedicated to proving the main result — Section 12 reduces it to a result about the expansion of (mock) locally minimal cochains (Theorem 12.8), which is then proved in Section 13.

2 Preliminaries

We begin by recalling relevant definitions and facts concerning expander graphs, locally testable codes, lifted codes and agreement testability.

General Conventions

The set of natural numbers \mathbb{N} does not include 0. A ring means a commutative (unital, associative) ring, and a module means a left module. The group of invertible elements in a ring R is denoted R^{\times} .

A cell complex means a CW complex, or more precisely, its underlying partially ordered set, which we assume to include a unique empty cell.

2.1 Expander Graphs

Throughout this paper, graphs are finite and allowed to have double edges, but no loops. A *simple* graph is a graph with no double edges and a *pure graph* is a nonempty graph in which every vertex belongs to some edge.

Given a graph G, we let G(0) denote its set of vertices and G(1) denote it set of edges. We also use G to denote the set $G(0) \cup G(1)$. We write v < e to indicate that v is a vertex of the edge e. The the set of (two) vertices of an edge $e \in G(1)$ is denoted e(0), and the set of edges having $v \in G(1)$ as a vertex is denoted $G(1)_v$. We will sometimes abuse the notation and write $e = \{u, v\}$ to say that e connects the vertices u and v, even though there may be other edges with that property.

A weight function on a graph G is a function $w: G(0) \cup G(1) \to \mathbb{R}_+$; we call (G, w) a weighted graph and, given $A \subseteq G$, write $w(A) = \sum_{a \in A} w(a)$. We make no assumptions on w. However, we will say that w is normalized if w(G(0)) = w(G(1)) = 1 and proper if we moreover have $w(v) = \frac{1}{2} \sum_{e \in G(1)_v} w(e)$ for every $v \in G(0)$ (which forces G to be pure). A normalized weight function defines probability measures on G(0) and G(1). It is normalized precisely when the probability of sampling a vertex v according to w is equal to the probability of getting v by choosing an edge according to w and then choosing one of its vertices uniformly at random.

Example 2.1. Let G be a graph.

(i) The uniform weight function $w_{\text{uni}}: G \to \mathbb{R}_+$ assigns every $v \in G(0)$ the weight $\frac{1}{|G(0)|}$ and every $e \in G(1)$ the weight $\frac{1}{|G(1)|}$. It is defined when G has at least one vertex and one edge, and is normalized.

(ii) Suppose that G is pure. The natural weight function of G is $w_{nat}: G \to \mathbb{R}_+$ defined by

$$w_{\text{nat}}(e) = \frac{1}{|G(1)|}$$
 and $w_{\text{nat}}(v) = \frac{|G(1)_v|}{2|G(1)|}$

for all $e \in G(1)$ and $v \in G(0)$. This weight function is proper.

The uniform weight function is not proper in general. When G is a regular graph (i.e. every vertex belongs to the same number of edges), the natural and uniform weight functions of G coincide.

All graphs (and also hypergraphs) in this work will carry a weight function, which by default will be the natural weight function.

Suppose henceforth that (G, w) is a properly weighted graph. The proofs of the following facts can be found in [FK23a, §2C], for instance.

Let $C^0(G, \mathbb{R})$ denote the space of functions $f : G(0) \to \mathbb{R}$, and let $C^0_{\circ}(G, \mathbb{R})$ denote its subspace of functions satisfying $\sum_{v \in G(0)} f(v) = 0$. As usual, the *weighted adjacency operator* of (G, w) is $\mathcal{A} = \mathcal{A}_{G,w} : C^0(G, \mathbb{R}) \to C^0(G, \mathbb{R})$ given by

$$(\mathcal{A}f)(v) = \sum_{e \in G(0)_v} \frac{w(e)}{2w(v)} f(e-v) \qquad \forall v \in G(0),$$

where e - v denotes the vertex of e which is different from v. For example, if G is k-regular and w is its natural weight function, then \mathcal{A} is just the usual vertex adjacency operator scaled by $\frac{1}{k}$.

The operator $\mathcal{A} : C^0_{\circ}(G, \mathbb{R}) \to C^0_{\circ}(G, \mathbb{R})$ is diagonalizable. The constant function $1_{G(0)}$ is an eigenfunction of \mathcal{A} with eigenvalue 1 and all other eigenvalues lie in the interval [-1, 1]. The subspace $C^0_{\circ}(G, \mathbb{R})$ is invariant under \mathcal{A} and, given $\lambda \in [-1, 1]$, we call (G, w) a λ -expander if all eigenvalues of \mathcal{A} on $C^0_{\circ}(G, \mathbb{R})$ lie in the interval $[-1, \lambda]$.

We will need the following special case of the Expander Mixing Lemma for weighted graphs.

Proposition 2.2 ([FK23a, Theorem 3.2(ii)]). Let (G, w) be a properly weighted graph, let $A \subseteq G(0)$ and let E(A) denote the set of edges $e \in G(0)$ with $e(0) \subseteq A$. If (G, w) is a λ -expander, then

$$w(E(A)) \le w(A)^2 + \lambda w(A).$$

Weighted graphs satisfying the condition $w(E(A)) \leq w(A)^2 + \lambda w(A)$ for every $A \subseteq G(0)$ are known as λ -skeleton expanders. Thus, every λ -expander weighted graph is also a λ -skeleton expander.

2.2 Conventions about Codes

Let Σ be a finite alphabet and $n \in \mathbb{N}$. In this work, an *error correcting code*, or a *code* for short, with alphabet Σ and block length n is a nonempty subset $C \subseteq \Sigma^n$. We also say that C is a code inside Σ^n . As usual, the normalized Hamming distance function on Σ^n is denoted dist and is given by $\operatorname{dist}(f,g) = \frac{1}{n} \cdot \#\{i \in \{1,\ldots,n\} : f_i \neq g_i\}$. When Σ is an abelian group, the normalized Hamming norm of $f \in \Sigma^n$ is $\|f\| = \frac{1}{n} \#\{i \in \{1,\ldots,n\} : f_i \neq 0\}$, so that $\operatorname{dist}(f,g) = \|f - g\|$. The *relative distance* of the code $C\sigma\Sigma^n$ is

$$\delta(C) := \frac{1}{n} \Delta(C) = \max\{\operatorname{dist}(f,g) \mid f, g \in C, f \neq g\}$$

and its rate is

$$r(C) = \log_{|\Sigma^n|} |C|.$$

The distance of f is $n \cdot \delta(C)$. Given $\eta \in [0,1]$ and $f \in \Sigma^n$, We say that f is η -close to C if $\operatorname{dist}(f,C) < \eta$ and η -far from C if $\operatorname{dist}(f,C) \ge \eta$ is smaller than η .

We will often think of $C \subseteq \Sigma^n$ as being part of a family of codes $\{C_i \subseteq \Sigma^{n_i}\}_{i \in \mathbb{N}}$ with block length tending to ∞ , and (abusing the notation) sometimes ascribe properties of the entire family $\{C_i \subseteq \Sigma^{n_i}\}_{i \in \mathbb{N}}$ to C. In this case, we will say that C (or the family $\{C_i \subseteq \Sigma^{n_i}\}_{i \in \mathbb{N}}$) is good if there are $\rho, \delta > 0$ such that each $r(C_i) \ge \rho$ and $\delta(C_i) \ge \delta$ for all i. When the latter holds, we also say that C has linear distance (as a function of the block length n).

Let $\eta \in [0, 1]$. A decoding algorithm for words that are η -close to C is an algorithm which takes as input some $f \in \Sigma^n$ with $\operatorname{dist}(f, C) < \eta$ and outputs some $f' \in C$ with $\operatorname{dist}(f, f') < \eta$; this f'is unique when $\eta \leq \frac{1}{2}\delta(C)$. The time complexity of a decoding algorithm will always be measure w.r.t. the block length n; ideally, it should be linear.

Remark 2.3 (Codes with Varying Alphabets). We can relax the definition of an error correcting code by considering words in which each letter comes from a different alphabet, i.e., the *i*-th letter of a word would come from an alphabet Σ_i depending on the position *i*. A code would then be a nonempty subset C of $\prod_{i=1}^{n} \Sigma_i$. All the notions just defined extend verbatim to such generalized codes.

The notation of a *locally testable code* (LTC) was recalled in §1.5.

2.3 Lifted Codes

Let Σ be a finite alphabet and $n \in \mathbb{N}$. Recall that a *lifted code*, or a *Tanner code*, is determined by specifying a collection S of subsets of $[n] := \{1, \ldots, n\}$ with $[n] = \bigcup_{s \in S} s$ and a code $C_s \subseteq \Sigma^s$ for every $s \in S$. Typically, all the sets in S will have the same size $k = \Theta(1)$ (as n grows) and every $i \in [n]$ will be contained in $D = \Theta(1)$ sets from S; the number of sets in S will therefore be $\frac{D}{k}n = \Theta(n)$. However, these extra assumptions are not necessary. The lifted code defined by the $\{C_s\}_{s \in S}$ is

$$C = C(\{C_s\}_{s \in S}) := \{g \in \Sigma^n : g|_s \in C_s \text{ for all } s \in S\} \subseteq \Sigma^n.$$

The codes $\{C_s\}_{s\in S}$ are often called the *small codes* defining C.

The description of $C \subseteq \Sigma^n$ as a lifted code also gives rise to a natural tester: Given $g \in \Sigma^n$, choose $s \in S$ uniformly at random, probe g(i) for every $i \in s$, and accept g if and only if $g|_s \in C_s$.

By replacing [n] with an arbitrary set M, we can study lifted codes inside Σ^M , rather than Σ^n .

2.4 Agreement Testability

Informally, an agreement expander consists of a collection S of subsets of [n] such that for any finite set Σ and any ensemble of functions $\{f_s : s \to \Sigma\}_{s \in S}$ such that $f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}$ for almost all $s, s' \in S$, there is $g : [n] \to \Sigma$ such that $g|_s = f_s$ for almost all $s \in S$; see [DK17] and [DD19]. For this work, we need to consider a refinement of this notion where each f_s is required to be in a code $C_s \subseteq \Sigma^s$ and the globally defined function $g : [n] \to \Sigma$ is required to be in the associated lifted code $C = C(\{C_s\}_{s \in S}) \subseteq \Sigma^n$. This notion already appeared in [DEL⁺22, Dfn. 2.8] in the special case of tensor codes, realized as lifted codes as in Example 2.6 below, and under the name agreement testability, which use as well.

The formal definition of agreement testability requires us to state which pairs $(s, s') \in S \times S$ are considered and in what probability. It is convenient to encode this information in a normalized weighted graph whose vertices are in bijection with S and whose edges are labelled by subsets of [n].

Definition 2.4 (Agreement Tester). Let $C = C(\{C_s\}_{s \in S}) \subseteq \Sigma^n$ be a lifted code as in §2.3. An agreement tester for the lifted code C consists of a normalized weighted graph (G, w) and a function $\ell: G \to P([n])$ assigning every vertex and edge a subset of [n] such that the following hold:

- (1) ℓ restricts to a bijection between G(0) and S;
- (2) for every edge $e \in G(1)$ and $u \in e(0)$, we have $\ell(e) \subseteq \ell(u)$.

In this case, we also say that $({C_s}_{s\in S}, G, w, \ell)$ is an agreement tester. This agreement tester is said to have soundness $\kappa \ge 0$ if for every ensemble $(f_s)_{s\in S} \in \prod_{s\in S} C_s$, there is $g \in C$ such that

$$\kappa \cdot w(\{v \in G(0) : g|_{\ell(v)} \neq f_{\ell(v)}\}) \le w(\{e = \{u, v\} \in G(1) : f_{\ell(u)}|_{\ell(e)} \neq f_{\ell(v)}|_{\ell(e)}\}).$$

We will also say that $({C_s}_{s\in S})$ is κ -agreement testable w.r.t. the labelled weighted graph (G, w, ℓ) .

Example 2.5. Any lifted code $\{C_s\}_{s\in S}$ can be naively enriched into an agreement tester as follows. Construct G by taking the vertex set to be S, and then connect a pair $s, s' \in S$ by an edge if $s \cap s' \neq \emptyset$, or if $s \cap s'$ has some desired cardinality. The labelling ℓ then maps every $s \in G(0)$ to itself and every edge $\{s, s'\}$ to $s \cap s'$. The weight function w can be taken to be the uniform one, for instance.

Example 2.6 (Agreement Testability of Tensor Codes). Let \mathbb{F} be a finite field. Let $C_1 \subseteq \mathbb{F}^{[n_1]}$ and $C_2 \subseteq \mathbb{F}^{[n_2]}$ be linear codes. The tensor code $C_1 \otimes C_2$ (all tensors are over \mathbb{F}) is the code $C \subseteq M_{n_1 \times n_2}(\mathbb{F}) = \mathbb{F}^{[n_1] \times [n_2]}$ consisting of the matrices $m \in M_{n_1 \times n_2}(\mathbb{F})$ such that every row of m lies in C_2 and every column of m lies in C_1 . In [DEL⁺22, Definition 2.8], the tensor code $C = C_1 \otimes C_2 \subseteq \mathbb{F}^{[n_1] \times [n_2]}$ is said to be a κ -agreement testable¹¹ if for every choice of codewords $\{f_j \in C_1\}_{j \in [n_2]}$ and $\{f'_i \in C_2\}_{i \in [n_1]}$, there is a matrix $m \in C$ such that

$$\kappa \cdot \left[\frac{\#\{i \in [n_1] : f'_i \neq r_i(m)\}}{2n_1} + \frac{\#\{j \in [n_2] : f'_j \neq c_j(m)\}}{2n_2}\right] \le \frac{\#\{(i,j) \in [n_1] \times [n_2] : f_{i,j} \neq f'_{j,i}\}}{n_1 n_2}.$$

(Here, $r_i(m)$ is the *i*-th row of m and $c_j(m)$ is the *j*-column of m.) Informally, this means that if the matrix whose rows are the $\{f'_i\}_i$ and the matrix whose columns are the $\{f_j\}_j$ agree in almost all entries, then some matrix in $C_1 \otimes C_2$ agrees almost everywhere with both of these matrices.

We can recover κ -agreement testablity for $C = C_1 \otimes C_2$ as a special case of Definition 2.4. First we realize $C = C_1 \otimes C_2 \subseteq \mathbb{F}^{[n_1] \times [n_2]}$ as a lifted code by taking $S = \{s_1, \ldots, s_{n_1}, s'_1, \cdots, s'_{n_2}\}$, where

¹¹Actually, what we define here is $\frac{\kappa}{2}$ -agreement testability in the setting of [DEL⁺22, Definition 2.8].

 $s_i = \{i\} \times [n_2] \text{ and } s'_j = [n_1] \times \{j\}, \text{ and putting } C_{s_i} = \{(f_j)_{(i,j) \in \{i\} \times [n_2]} | f \in C_2\} \subseteq \mathbb{F}^{\{i\} \times [n_2]}$ and $C_{s'_j} = \{(f_i)_{(i,j) \in [n_1] \times \{j\}} | f \in C_1\} \subseteq \mathbb{F}^{[n_1] \times \{j\}}$ for all i and j. Now choose the graph G to be the complete biparatite graph on $\{s_1, \ldots, s_{n_1}\}$ and $\{s'_1, \cdots, s'_{n_2}\}$ endowed with its natural weight function. The labelling ℓ maps every vertex to itself, and every edge $\{s_i, s'_j\}$ to $s_i \cap s'_j = \{(i, j)\}$. It is routine (and a recommend exercise for newcomers) to check that $(\{C_s\}_{s \in S}, G, \ell)$ has soundness κ if and only if $C_1 \otimes C_2$ is κ -agreement testable.

3 Lifted Codes and Their Line Codes

In this section, we recall the construction of the (so-called) line code of a lifted code. We then establish relations between the rate, distance, testability and decodability of these codes. The results of this section will be important to the applications our main result.

Let $C \subseteq \Sigma^n$ be a lifted code determined by small codes $\{C_s\}_{s \in S}$ (§2.3). Suppose moreover that all the small codes C_s have the same cardinality σ and choose a set Σ' of that cardinality. The *line* code of $C = C(\{C_s\}_{s \in S})$ is a code $L = L(\{C_s\}_{s \in S}) \subseteq \Sigma'^S$ with alphabet Σ' constructed as follows: For every $s \in S$, choose a bijection $C_s \cong \Sigma'$. We use these bijections to freely identify $\prod_{s \in S} C_s$ with Σ'^S . We then define $L \subseteq \Sigma'^S$ to be the code consisting of the (words in Σ'^S corresponding to) ensembles $f = (f_s)_{s \in S} \in \prod_{s \in S} C_s$ satisfying $f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}$ for all $s, s' \in S$. That is,

$$L = L(\{C_s\}_{s \in S}) = \{(f_s)_{s \in S} \in \prod_{s \in S} C_s : f_s|_{s \cap s'} = f_{s'}|_{s \cap s'} \text{ for all } s, s' \in S\}$$

Since the sets in S cover $[n] = \{1, \ldots, n\}$, we have a bijection $C \to L$ given by $g \mapsto (g|_s)_{s \in S}$.

Remark 3.1. By allowing codes with a varying alphabet, see Remark 2.3, we may define the line code of $C(\{C_s\}_{s\in S})$ even when the C_s have varying cardinalities — simply let $\Sigma_s = C_s$ and define L as a subset of $\prod_{s\in S} \Sigma_s$. With the exception of Proposition 3.6(i), the results of this section can be adapted in a straightforward manner to this more general setting.

Example 3.2 (Line Codes of Reed-Müller Codes). Let \mathbb{F} be a finite field of cardinality q and characteristic p and let $0 \leq d \leq n$. Let V_n denote an n-dimension \mathbb{F} -vector space, e.g., \mathbb{F}^n . Recall that the Reed-Müller code of degree-d functions on V_n is the set $C = \operatorname{RM}(n, d, q) \subseteq \mathbb{F}^{V_n}$ for functions $g: V_n \to \mathbb{F}$ having degree at most d. It is known [KR06, Thm. 2] that when $d \leq q(1 - \frac{1}{p}) - 1$, a function $g: V_n \to \mathbb{F}$ has degree d or less if and only if its restriction to every 1-dimensional affine subspace of V_n — called a *line* for short — is also of degree d or less. Assuming this holds, we can describe $\operatorname{RM}(n, d, q)$ as a lifted code $C(\{C_s\}_{s \in S})$: Let S be the set of lines in V_n , and for every $s \in S$, let C_s be the Reed-Müller code of degree-d functions on the line s. By identifying each line s in V_n with $V_1 \cong \mathbb{F}$, we can identify each C_s with the Reed-Müller code $\operatorname{RM}(1, d, q)$. Thus, the line code L of $\operatorname{RM}(n, d, q)$, realized as a lifted code as just explained, has alphabet $\Sigma' = \operatorname{RM}(1, d, q)$ and it consists of the ensembles $(f_s)_{s \in S} \in \operatorname{RM}(1, d, q)^S$ such that $f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}$ for any two lines s, s' in V_n . This example is the reason why $L(\{C_s\}_{s \in S})$ is called a line code in general.

Notation 3.3. For the remainder of this section, fix a lifted code $C = C(\{C_s\}_{s \in S}) \subseteq \Sigma^n$ with block length n and alphabet Σ . Suppose moreover that each C_s is identified with another alphabet Σ' , and let $L = L(\{C_s\}_{s \in S}) \subseteq \Sigma'^S$ be the the associated line code. We will make repeated use of the following quantities associated to the family $S \subseteq P([n])$:

- $k_{\min} = \min_{s \in S} |s|,$
- $k_{\max} = \max_{s \in S} |s|,$

- $D_{\min} = \min_{i \in [n]} \# \{ s \in S : i \in s \},\$
- $D_{\max} = \max_{i \in [n]} \# \{ s \in S : i \in s \}.$

Thus, every $s \in S$ contains between k_{\min} and k_{\max} elements, and every $i \in [n]$ is contained in at least D_{\min} and at most D_{\max} sets from S. We encourage the reader think of $k_{\min}, k_{\max}, D_{\min}, D_{\max}$ as being $\Theta(1)$ as n grows. This implies that $|S| = \Theta(n)$, as the following lemma shows.

Lemma 3.4. With notation as above, $\frac{D_{\min}}{k_{\max}}n \leq |S| \leq \frac{D_{\max}}{k_{\min}}n$.

Proof. The right inequality holds because $k_{\min}|S| \le \#\{(i,s) \in [n] \times S : i \in s\} \le nD_{\max}$. The left inequality is shown similarly.

The following lemma will be used repeatedly.

Lemma 3.5. With notation as in Notation 3.3, let $g_0 \in C$ correspond to $f_0 = (g_0|_s)_{s \in S} \in L$. Let $g \in C$, put $S_g = \{s \in S : g|_s \in C_s\}$ and define $f \in \prod_{s \in S} C_s$ by letting $f_s = g|_s$ if $g|_s \in C_s$ and otherwise choosing $f_s \in C_s$ arbitrarily. Then

$$\frac{D_{\min}k_{\min}}{D_{\max}k_{\max}}\operatorname{dist}(g,g_0) - \frac{|S_g|}{|S|} \le \operatorname{dist}(f,f_0) \le \frac{D_{\max}k_{\max}}{D_{\min}}\operatorname{dist}(g,g_0)$$

Proof. Write $A = \{i \in [n] : g_i \neq g_{0,i}\}, B = \{s \in S : g|_s \neq g_0|_s\}$ and $I = \{(i, s) \in A \times B : i \in s\}$. Every $i \in A$ has between D_{\min} and D_{\max} preimages under the first projection $\operatorname{pr}_1 : I \to A$, so $\frac{1}{D_{\max}}|I| \leq |A| \leq \frac{1}{D_{\min}}|I|$. Every $s \in B$ has between 1 and k_{\max} preimages under the second projection $\operatorname{pr}_2 : I \to B$, so $\frac{1}{k_{\max}}|I| \leq |B| \leq |I|$. Together, both inequalities imply that

$$\frac{D_{\min}}{k_{\max}}|A| \le |B| \le D_{\max}|A|.$$

Observe that if $f_s \neq f_{0,s}$, then we must have $g|_s \neq g_0|_s$ (otherwise $g|_s \in C_s$ and then $f_s = g|_s = g_0|_s = f_{0,s}$). Thus,

$$\operatorname{dist}(f, f_0) = \frac{\#\{s \in S : f_s \neq f_{0,s}\}}{|S|} \le \frac{|B|}{|S|} \le \frac{D_{\max}|A|}{\frac{D_{\min}}{k_{\max}}n} = \frac{D_{\max}k_{\max}}{D_{\min}}\operatorname{dist}(g, g_0),$$

where in the second inequality we used Lemma 3.4. Next, observe that if $g|_s \neq g_0|_s$ and $s \notin S_g$, then $f_s \neq f_{0,s}$. (Indeed, $g|_s = f_s$ because $s \notin S_g$, so $f_s = g|_s \neq g_0|_s = f_{0,s}$.) Thus,

$$\operatorname{dist}(f, f_0) = \frac{\#\{s \in S : f_s \neq f_{0,s}\}}{|S|} \ge \frac{|B| - |S_g|}{|S|} \ge \frac{\frac{D_{\min}}{k_{\max}}|A|}{\frac{D_{\max}}{k_{\min}}n} - \frac{|S_g|}{|S|} = \frac{D_{\min}k_{\min}}{D_{\max}k_{\max}}\operatorname{dist}(g, g_0) - \frac{|S_g|}{|S|}.$$

This proves the lemma.

The following proposition relates the rate and distance of C and L; this is fairly standard.

Proposition 3.6. Using Notation 3.3, we have:

(i)
$$r(C) = \frac{\gamma \log |\Sigma'|}{\log |\Sigma|} r(L)$$
, where $\gamma := \frac{|S|}{n} \in [\frac{D_{\min}}{k_{\max}}, \frac{D_{\max}}{k_{\min}}]$.
(ii) $\frac{D_{\min}}{D_{\max}k_{\max}} \delta(L) \le \delta(C) \le \frac{D_{\max}k_{\max}}{D_{\min}k_{\min}} \delta(L)$.

Proof. (i) Recall that we have a bijection $g \mapsto (g|_s)_{s \in S} : C \to L$. Thus,

$$r(L) = \frac{\log |L|}{|S| \log |\Sigma'|} = \frac{\log |C|}{\gamma n \log |\Sigma'|} = \frac{\log |\Sigma|}{\gamma \log |\Sigma'|} r(C).$$

That $\gamma \in [\frac{D_{\min}}{k_{\max}}, \frac{D_{\max}}{k_{\min}}]$ follows from Lemma 3.4. (ii) Let $f, f' \in L$ and let g, g' be the corresponding codewords in C. Applying Lemma 3.5 with our g and $g_0 = g'$ (note that $S_g = \emptyset$) gives

$$\frac{D_{\min}k_{\min}}{D_{\max}k_{\max}}\operatorname{dist}(g,g') \le \operatorname{dist}(f,f') \le \frac{D_{\max}k_{\max}}{D_{\min}}\operatorname{dist}(g,g').$$

This implies readily that $\frac{D_{\min}}{D_{\max}k_{\max}}\delta(L) \leq \delta(C) \leq \frac{D_{\max}k_{\max}}{D_{\min}k_{\min}}\delta(L)$.

The next proposition says that if the line code L has a decoding algorithm, then C also has a decoding algorithm of a similar complexity. We do not know whether the converse holds in general, but a partial converse will be given in Theorem 3.11(ii) below.

Proposition 3.7. With notation as in Notation 3.3, let $\eta \in [0, \frac{1}{2}\delta(L)]$ and suppose that the line code L has a decoding algorithm for words that are η -close to L. Then C has a decoding algorithm for words that are $\frac{D_{\min}}{D_{\max}k_{\max}}\eta$ -close to C. Provided that $k_{\max} = O(1)$ (as a function of n), its time complexity is O(n + |S|) plus the time complexity of the decoding algorithm for L.

Proof. Consider the following algorithm, which takes $q \in \Sigma^n$ and outputs $q' \in C$.

- (1) For every $s \in S$: If $g|_s \in C_s$, set $f_s = g|_s$; otherwise, let f_s be some element of C_s .
- (2) Apply the decoding algorithm of L to $f = (f_s)_{s \in S} \in \prod_{s \in S} C_s$. Let f' be the output.
- (3) The ensemble $f' = (f'_s)_{s \in S} \in L$ determines an element $g' \in C$. Output g'.

We claim that this algorithm has the required properties.

The time complexity is clearly the one stated in the proposition.

Suppose now that the input g of the algorithm satisfies $\operatorname{dist}(g, C) < \frac{D_{\min}}{D_{\max}k_{\max}}\eta$ and choose $g_0 \in C$ such that $\operatorname{dist}(g, g_0) = \operatorname{dist}(g, C)$. We need to show that the output g' of the algorithm is g_0 . Let $f_0 \in L$ correspond to g_0 . By Lemma 3.5,

$$\operatorname{dist}(f, f_0) \le \frac{D_{\max} k_{\max}}{D_{\min}} \operatorname{dist}(g, g_0) < \eta.$$

Thus, applying the decoding algorithm of L to f returns f_0 . Consequently $f' = f_0$ and the algorithm outputs g_0 , as required.

We now turn to consider the testability of the line code L by a 2-query tester. To that end, let G be a graph equipped with a labelling $\ell: G \to P([n])$ such that ℓ restrict to a bijection $G(0) \to S$ and $\ell(v) \supseteq \ell(e)$ for every $e \in G(1)$ with vertex v.

Example 3.8. We can take G to be the *intersection graph* of S: The vertex set of G is S and $s, s' \in S = G(0)$ are connected an edge precisely when $s \cap s' \neq \emptyset$. The labelling $\ell: G \to P([n])$ then maps every $s \in G(0)$ to itself and every edge $\{s, s'\}$ to $s \cap s'$.

Having fixed a labelled graph (G, ℓ) as above, we define a 2-query tester $T_{G,\ell}$ for $L \subseteq \Sigma'^S$ as follows: Given $f = (f_s)_{s \in S} \in \Sigma'^S$, choose an edge $e = \{u, v\}$ in G uniformly at random, probe $f_{\ell(u)}$ and $f_{\ell(v)}$, and accept f if and only if $f_{\ell(u)}|_{\ell(e)} = f_{\ell(v)}|_{\ell(e)}$.

Remark 3.9. Give G the uniform weight function w_{uni} . Then $(\{C_s\}_{s\in S}, G, w_{\text{uni}}, \ell)$ is an agreement tester (§2.4) and it has soundness $\mu \geq 0$ if and only if the tester $T_{G,\ell}$ for the code $L \subseteq \Sigma'^S$ has soundness μ . This continues to hold if we give G any normalized weight function w which is uniform on G(0), provided that in $T_{G,\ell}$ we choose $e \in G(1)$ according to w (rather than uniformly).

We now show that if the tester $T_{G,\ell}$ of L has soundness μ , then the natural tester of the lifted code $C = C(\{C_s\})_{s \in S}$ has soundness $\Omega(\mu)$. We will apply this key observation to some particular lifted codes and their line codes later on.

Theorem 3.10. Keep Notation 3.3, let (G, ℓ) be a labelled graph as above, and suppose that every vertex in G belongs to at least d_{\min} and at most d_{\max} edges. If the tester $T_{G,\ell}$ for $L \subseteq \Sigma'^S$ has soundness μ ($\mu \ge 0$), then the natural tester of C has soundness $\frac{k_{\min}D_{\min}}{k_{\max}D_{\max}} \cdot \frac{\mu}{\mu + 2d_{\max}d_{\min}^{-1}}$.

Proof. We give G the uniform weight function $w := w_{\text{uni}}$ and identify G(0) with S via ℓ .

Let $g \in \Sigma^n$ and $S_g = \{s \in S : g|_s \in C_s\}$. The probability that the natural tester of C rejects g is $\frac{|S_g|}{|S|}$, so we need to show that

$$\frac{|S_g|}{|S|} \ge \frac{k_{\min}D_{\min}}{k_{\max}D_{\max}} \cdot \frac{\mu}{\mu + 2d_{\max}d_{\min}^{-1}} \cdot \operatorname{dist}(g,C).$$
(3.1)

Define f as in Lemma 3.5, choose $f_0 \in L$ such that $\operatorname{dist}(f, f_0) = \operatorname{dist}(f, L)$ and let $g_0 \in C$ be the codeword corresponding to f_0 . By Lemma 3.5, we have

$$\operatorname{dist}(g, g_0) \le \frac{D_{\max}k_{\max}}{D_{\min}k_{\min}} \left(\frac{|S_g|}{|S|} + \operatorname{dist}(f, f_0)\right) = \frac{D_{\max}k_{\max}}{D_{\min}k_{\min}} \left(\frac{|S_g|}{|S|} + \operatorname{dist}(f, L)\right).$$
(3.2)

Next, observe that the probability that $T_{G,\ell}$ rejects f is at most

$$w(\bigcup_{s\in S_g} G(1)_s) \le \sum_{s\in S_g} \frac{|G(1)_s|}{|G(1)|} \le \sum_{s\in S_g} \frac{d_{\max}}{\frac{1}{2}d_{\min}|S|} = \frac{2d_{\max}}{d_{\min}} \frac{|S_g|}{|S|}.$$

(Recall that we identified G(0) with S and that $G(1)_s$ is the set of edges having s as a vertex.) Since $T_{G,\ell}$ has soundness μ , it follows that

$$\operatorname{dist}(f, L) \le \frac{2d_{\max}}{\mu d_{\min}} \frac{|S_g|}{|S|}.$$

Plugging this into (3.2) gives

$$\operatorname{dist}(g, C) \le \operatorname{dist}(g, g_0) \le \frac{D_{\max}k_{\max}}{D_{\min}k_{\min}} (1 + \frac{2d_{\max}}{\mu d_{\min}}) \frac{|S_g|}{|S|}$$

Rearranging gives the desired conclusion (3.1).

We finish this section with giving a converse to Theorem 3.10 when G is the intersection graph of S, as well as a partial converse to Proposition 3.7. This will not be needed in the sequel. We do not know if there is a converse to Theorem 3.10 which holds in general.

Theorem 3.11. Keep Notation 3.3, let (G, ℓ) be the intersection graph of the set S (Example 3.8) and let d_{\max} (resp. d_{\min}) denote the maximal (resp. minimal) degree of a vertex in G. Suppose further that the natural tester of the lifted code C has a soundness $\mu \geq 0$. Then:

- (i) The tester $T_{G,\ell}$ for $L \subseteq \Sigma'^S$ has soundness $\frac{D_{\min}k_{\min}\mu}{d_{\max}D_{\max}^2k_{\max}^2(\mu+D_{\min}^{-1}D_{\max}k_{\max})}$.
- (ii) If C has a decoding algorithm for words that are η -close to C ($\eta > 0$), then L has a decoding algorithm for words that are η' -close to L, where

$$\eta' = \min\left\{\frac{d_{\min}D_{\min}k_{\min}\mu}{2d_{\max}^2 D_{\max}^2 k_{\max}^2} \cdot \eta, \left(\frac{2d_{\max}^2 D_{\max}^2 k_{\max}^2 (\mu + D_{\min}^{-1} D_{\max} k_{\max})}{d_{\min} D_{\min} k_{\min} \mu} + 1\right)^{-1} \delta(L)\right\}$$

Its time complexity is O(|G(1)| + n) plus the time complexity of the decoding algorithm of C.

Proof. Again, we identify G(0) with S via the labelling ℓ . We also write E = G(1).

(i) Let $f \in \prod_{s \in S} C_s$, and let $E_f = \{e = \{s, s'\} \in E : f_s|_{s \cap s'} = f_{s'}|_{s \cap s'}\}$. We need to show that

$$\frac{|E_f|}{|E|} \ge \frac{D_{\min}k_{\min}\mu}{d_{\max}D_{\max}^2k_{\max}^2(\mu + D_{\min}^{-1}D_{\max}k_{\max})} \operatorname{dist}(f, L)$$

Let $S_f = \bigcup_{e \in E_f} e(0)$ (e(0) is the set of vertices of e). Then

$$\frac{|S_f|}{|S|} \le \frac{2|E_f|}{|S|} = \frac{2|E|}{|S|} \cdot \frac{|E_f|}{|E|} \le d_{\max} \frac{|E_f|}{|E|}.$$

Next, put $M = \bigcup_{s \in S_f} s$ (so that $M \subseteq [n]$). Then

$$\frac{|M|}{n} \le \frac{k_{\max}|S_f|}{n} = k_{\max} \frac{|S_f|}{|S|} \frac{|S|}{n} \le \frac{d_{\max}D_{\max}k_{\max}}{k_{\min}} \frac{|E_f|}{|E|},$$

where in the second inequality we used Lemma 3.4.

Let $i \in [n] - M$ and suppose that $s, s' \in S$ satisfy $i \in s \cap s'$. Then $\{s, s'\} \in E - E_f$ (here we need G to be the intersection graph of S), and thus $(f_s)_i = (f_{s'})_i$. This allows us to define $g \in \Sigma^n$ as follows: For $i \in [n] - M$, choose some $s \in S$ with $i \in s$ and define $g_i = (f_s)_i$; this is independent of s by what we just showed. For $i \in M$, choose $g_i \in \Sigma$ arbitrarily.

Let $S_g = \{s \in S : g | s \notin C_s\}$. Observe that every $s \in S$ with $s \cap M = \emptyset$ satisfies $s \notin S_g$, because $g|_s = f_s \in C_s$. Otherwise stated, $S_g \subseteq \{s \in S : s \cap M \neq \emptyset\}$. Thus,

$$\frac{|S_g|}{|S|} \le \frac{D_{\max}|M|}{|S|} = D_{\max}\frac{|M|}{n}\frac{n}{|S|} \le \frac{d_{\max}D_{\max}^2k_{\max}^2}{D_{\min}k_{\min}}\frac{|E_f|}{|E|}$$
(3.3)

(we used Lemma 3.4 again). The number $\frac{|S_g|}{|S|}$ is also the probability that the natural tester of C rejects g. Thus, there exists $g_0 \in C$ such that

dist
$$(g, g_0) \le \mu^{-1} \cdot \frac{|S_g|}{|S|} \le \frac{d_{\max} D_{\max}^2 k_{\max}^2}{D_{\min} k_{\min} \mu} \frac{|E_f|}{|E|}.$$

Let $f_0 \in L$ denote the codeword corresponding to g_0 .

Define $f' \in \Sigma'^S$ by letting $f'_s = g|_s$ if $g|_s \in C_s$, and choosing f'_s arbitrarily otherwise. By Lemma 3.5,

$$\operatorname{dist}(f', f_0) \leq \frac{D_{\max}k_{\max}}{D_{\min}}\operatorname{dist}(g, g_0) \leq \frac{D_{\max}k_{\max}}{\mu D_{\min}}\frac{d_{\max}D_{\max}^2k_{\max}^2}{D_{\min}k_{\min}}\frac{|E_f|}{|E|}.$$

Note also that if $s \in S$ satisfies $s \cap M = \emptyset$, then $g|_s = f_s \in C_s$, so $f_s = f'_s$. This means that $\{s \in S : f_s \neq f'_s\} \subseteq \{s \in S : s \cap M \neq \emptyset\}$, and together with (3.3), we get

$$\operatorname{dist}(f, f') \le \frac{D_{\max}|M|}{|S|} \le \frac{d_{\max}D_{\max}^2k_{\max}^2}{D_{\min}k_{\min}}\frac{|E_f|}{|E|}.$$

It follows that

$$\operatorname{dist}(f,L) \le \operatorname{dist}(f,f_0) \le \operatorname{dist}(f,f') + \operatorname{dist}(f',f_0) \le \frac{d_{\max}D_{\max}^2k_{\max}^2}{D^{\min}k_{\min}}(1 + \frac{D_{\max}k_{\max}}{\mu D_{\min}})\frac{|E_f|}{|E|}.$$

Rearranging gives the desired conclusion.

- (ii) Consider the following algorithm taking $f \in \prod_{s \in S} C_s$ and outputing $f_0 \in L$:
- (1) Define E_f , M, g as in the proof of (i).
- (2) Apply the decoding algorithm of C to g; let g_0 denote the output.
- (3) Return the codeword $f_0 \in L$ corresponding to g_0 .

It clearly has the time complexity claimed in the theorem. It remains to show that it decodes f if $dist(f, L) < \eta'$.

Let $f_1 \in L$ such that $\operatorname{dist}(f, f_1) = \operatorname{dist}(f, L) < \eta'$. Let $T = \{s \in S : f_s \neq f_{1,s}\}$. Then $|T| < \eta' |S|$. Note that any $e \in \{s, s'\} \in E$ with $s, s' \notin T$ does not lie in E_f because $f_s|_{s \cap s'} = f_{1,s}|_{s \cap s'} = f_{1,s'}|_{s \cap s'} = f_{s'}|_{s \cap s'}$. It follows that every $e \in E_f$ has a vertex in T. Thus,

$$\frac{|E_f|}{|E|} \le \frac{|T|d_{\max}}{|E|} = d_{\max}\frac{|S|}{|E|}\frac{|T|}{|S|} < \frac{2d_{\max}\eta'}{d_{\min}}.$$

As shown in the proof of (i),

$$\operatorname{dist}(g,g_0) \le \frac{d_{\max}D_{\max}^2k_{\max}^2}{D_{\min}k_{\min}\mu} \frac{|E_f|}{|E|} < \frac{2d_{\max}^2D_{\max}^2k_{\max}^2\eta'}{d_{\min}D_{\min}k_{\min}\mu} \le \eta.$$

This means that the decoding algorithm of C will work for g and return g_0 . We also observed in the proof of (i) that

$$dist(f, f_0) \leq \frac{d_{\max} D_{\max}^2 k_{\max}^2 (\mu + D_{\min}^{-1} D_{\max} k_{\max})}{D_{\min} k_{\min} \mu} \frac{|E_f|}{|E|} \\ < \frac{2d_{\max}^2 D_{\max}^2 k_{\max}^2 (\mu + D_{\min}^{-1} D_{\max} k_{\max})}{d_{\min} D_{\min} k_{\min} \mu} \eta' \leq \delta(L) - \eta'.$$

Thus, $dist(f_0, f_1) < \eta' + \delta(L) - \eta' = \delta(L)$, so $f_0 = f_1$, and the algorithm returns f_1 .

4 Graded Partially Ordered Sets

Recall that a partially ordered set, or a *poset* for short, is a set P equipped with a transitive antireflexive relation <. We then write $a \leq b$ to denote that a < b or a = b, and $a \triangleleft b$ to denote that a < b and there no $c \in X$ with a < c < b. Every subset of a poset X will also be viewed as a poset by giving it the partial order inherited from X. If not indicated otherwise, all posets are finite.

4.1 Graded Posets

Definition 4.1 (Graded Poset). A graded poset is a poset X together with a dimension function¹² dim = dim_X : $X \to \mathbb{Z}$ such that $x \leq y$ implies dim $x \leq \dim y$ and $x \triangleleft y$ implies dim $x + 1 = \dim y$ for all $x, y \in X$.¹³ In this case, we write

$$X(i) = \{x \in X \mid \dim x = i\}$$

for all $i \in \mathbb{Z}$ and define the dimension of X to be dim $X := \sup\{i \in \mathbb{Z} : X(i) \neq \emptyset\}.$

Beware that a subset of a graded poset is not a graded poset in general.

Motivated by examples of geometric nature, we call the elements of X(i) the *i*-faces of X. Given a face $x \in X$, a subface of x is a $y \in X$ satisfying $y \leq x$. We further write

$$x(i) := \{ y \in X(i) : y \le x \}$$

and call elements of x(i) *i*-faces of x. The set of faces $y \in X$ having x as a face is denoted $X_x := \{y \in X : y \ge x\}$. More generally, for every $A \subseteq X$, we write

$$A_x = \{a \in A : a \ge x\}.$$

In particular, $X(i)_x$ is the set of *i*-faces of X having x as a subface. Finally, we write $X(\leq i)$ for the graded subposet $\bigcup_{j \leq i} X(j)$.

Example 4.2. (i) Finite simplicial complexes and cube complexes are naturally graded posets. Their dimension function assigns every face its geometric dimension with the convention that the empty face has dimension -1.

(ii) Generalizing (i), the (closed) faces of a regular cell complex (also called a regular CW complex) form a graded poset w.r.t. inclusion of faces; see [AB08, Apx. A.2] for the definition. We follow the convention that a cell complex must includes a unique empty face of dimension -1. The posets of regular cell complexes can be characterized combinatorially [Bjö84, Prp. 3.1], so henceforth, a regular cell complex, we will mean the poset of faces of a regular cell complex (including the empty face).

(iii) Let \mathbb{F} be a finite field and $n, d \in \mathbb{N}$. Let $\operatorname{AG}_{d,n}(\mathbb{F})$ denote all affine subspaces of \mathbb{F}^n of dimension d or less together with the set \emptyset . Then $\operatorname{AG}_{d,n}(\mathbb{F})$ together with the containment relation is a poset known as the *affine Grassmannian* of d-spaces in \mathbb{F}^n . It can be made into a graded poset by setting the dimension of $V \in \operatorname{AG}_{d,n}(\mathbb{F})$ to be its ordinary \mathbb{F} -dimension if $V \neq \emptyset$ and -1 otherwise.

Example 4.3 (Viewing Hypergraphs as Graded Posets). A (finite) hypergraph X (possibly with multiple hyperedges) is nothing but a graded poset X concentrated in degrees 0 and 1, i.e., a poset such that $X(i) = \emptyset$ for all $i \neq 0, 1$. Indeed, think of the 0-faces are the vertices of X, the 1-faces as the hyperedges of X and the relation < as the incidence relation between vertices and hyperedges. In particular, we shall freely view graphs are graded posets concentrated in degrees 0 and 1.¹⁴

 $^{^{12}}$ Also called a rank function.

¹³Some texts impose additional assumptions, e.g., the requirement that X admits an element \emptyset_X (necessarily unique) satisfying dim $\emptyset_X = -1$ and $\emptyset_X \leq x$ for every $x \in X$. This forces dim $x \geq 0$ for every $x \in X - \{\emptyset_X\}$.

¹⁴According to our conventions, simple graphs and 1-dimensional simplicial complexes are not exactly the same thing, the difference being that a 1-dimensionsal simplicial complex must include an empty face of dimension -1 while a graph cannot include such a face.

Example 4.4 (Opposite Graded Poset). Let X be a graded poset. The opposite graded poset of X is the set $X^{\text{op}} = \{x^{\text{op}} | x \in X\}$ endowed with the relation $x^{\text{op}} < y^{\text{op}} \iff y < x$ and the dimension function $\dim(x^{\text{op}}) = -\dim x$.

Definition 4.5 (Pure Graded Poset, d-Poset). Let $d \in \mathbb{N} \cup \{0\}$. A graded poset X is said to be pure of dimension d if its nonempty and every face of x is a subface of a d-face; it is said to be pure if it is pure for some $d \in \mathbb{N} \cup \{0\}$. We say that X is a d-poset if it is pure of dimension d and in addition, there is an element $\emptyset_X \in X$ satisfying dim $\emptyset_X = -1$ and $\emptyset_X \leq x$ for all $x \in X$.

When X is a d-poset, the face \emptyset_X is unique. We call it the *empty face* of X and denote it by \emptyset when X is clear from the context.

The posets in Example 4.2 are examples *d*-posets when they are pure. A graph G is pure in the sense of §2.1 if and only if it is a pure poset of dimension 1 (but it is never a 1-poset; Example 4.3).

If X is a d-poset, then every subset $A \subseteq X$ has a lower bound. Let L be the set of lower bounds of A. As usual, an *infimum* of A is a maximal member L. The set of infima of A is denoted

$\operatorname{Inf} A.$

This set is often a singleton, e.g., when X is a regular cell complex.

4.2 Weighted Posets

Definition 4.6 (Weighted Poset). A weighted poset is a pair (X, w) where X is a poset and $w: X \to \mathbb{R}_+$. In this case, for any $A \subseteq X$, we let $w(A) = \sum_{a \in A} w(a)$. We say that w or (X, w) is normalized if w(X(i)) = 1 for all $i \in \mathbb{Z}$ with $X(i) \neq \emptyset$.

Definition 4.7 (Properly Weighted Poset). A properly weighted poset is a weighted graded poset (X, w) such that

- (1) X is pure of dimension d for some (necessarily unique) $d \ge 0$,
- (2) w(X(d)) = 1, and
- (3) $w(x) = \sum_{y \in X(d): y \ge x} \frac{w(y)}{|y(i)|}$ for all $i \in \mathbb{Z}$ and $x \in X(i)$.

In this case, we also say that w is a proper weight function on X.

It follows readily from the definition that if X is a properly weighted poset of dimension d and x is an *i*-face of X, then w(x) is the probability of getting x by choosing a d-face y of X at random according to $w|_{X(d)}$ and then choose an *i*-face of y uniformly at random. Thus, a properly weighted poset is also normalized.

Following Example 4.3, a (properly) weighted hypergraph means a (properly) weighted graded poset (X, w) concentrated in degrees 0 and 1. In the case of graphs, this agrees with the notion of a properly weighted graph from §2.1.

Example 4.8. (i) Let X be a pure poset of dimension d. The *natural weight function* on X is the weight function $w_{\text{nat}}: X \to \mathbb{R}_+$ defined by

$$w_{\text{nat}}(x) = \frac{1}{|X(d)|} \sum_{y \in X(d): y \ge x} \frac{w(y)}{|y(i)|}.$$

The natural weight function is always proper.

(ii) Let X be an poset. The uniform weight function on X is weight function $w_{\text{uni}} : X \to \mathbb{R}_+$ defined by

$$w_{\mathrm{uni}}(x) = \frac{1}{|X(i)|}.$$

The uniform weight function is normalized, but not always proper.

4.3 Links

Definition 4.9 (Link in Graded Poset). Let X be a graded poset and let $z \in X$. The link of X at z is

$$X_z = \{ x \in X : x \ge z \},$$

viewed as a subposet of X, and endowed with the dimension function given by $\dim_{X_z}(x) = \dim_X x - \dim_X z - 1$.

We will abbreviate \dim_{X_z} to \dim_z when there is no risk of confusion.

Example 4.10. Let X be a simplicial complex and let $z \in X$. The link of X in z is usually defined to be the poset $X'_z := \{y \in X : y \cup z \in X \text{ and } y \cap z = \emptyset\}$ which is also a simplicial complex; see [AB08, Dfn. A.19], for instance. While our X_z is different from X'_z in general, we have a graded poset isomorphism $X'_z \to X_z$ given by $y \mapsto y \cup z$, so X_z may be though of the usual link of X at z.

When the graded poset X is pure of dimension d and z is an *i*-face of X, the link X_z is a graded $(d - \dim z - 1)$ -poset with $\emptyset_{X_z} = z$. Moreover, every proper weight function $w : X \to \mathbb{R}_+$ induces a proper weight function on $w_z : X_z \to \mathbb{R}_+$ defined by

$$w_z(x) = \frac{1}{|w(X_z)|} \sum_{y \in X(d)_z} \frac{w(y)}{|y(\dim_X x)|}.$$

For details about the ratio between w and w_z , see Lemma 4.17 below.

When X is a *d*-poset, a proper link of X means a link X_z with $z \neq \emptyset_X$. For a proper link X_z , we have dim $X_z < \dim X$, whereas $X_{\emptyset_X} = X$.

4.4 Face Counting Constants and Lower-Regular Posets

Throughout, let X be a graded poset.

Given integers $i \leq j \leq k$, we let $F_{i,j,k}^{\max}(X)$ (resp. $F_{i,j,k}^{\min}(X)$) denote the maximal (resp. minimal) possible number of *j*-faces living between an *i*-face and a *k*-face that are incident in *X*. Formally, if there exist $x \in X(i)$ and $z \in X(k)$ with $x \leq z$, define

$$F_{i,j,k}^{\max}(X) = \max\{\#\{y \in X(j) : x \le y \le z\} \mid x \in X(i), z \in X(k), x \le z\},\$$

$$F_{i,j,k}^{\min}(X) = \min\{\#\{y \in X(j) : x \le y \le z\} \mid x \in X(i), z \in X(k), x \le z\}.$$

Otherwise, set $F_{i,j,k}^{\max} = F_{i,j,k}^{\min} = 0$. When X is a d-poset and i = -1, the number $F_{i,j,k}^{\max}(X)$ (resp. $F_{i,j,k}^{\min}(X)$) is the maximal (resp. minimal) possible number of j-faces contained in a k-face of X, and we abbreviate

$$F_{j,k}^{\max}(X) = F_{-1,j,k}^{\max}(X) \quad \text{ and } \quad F_{j,k}^{\min}(X) = F_{-1,j,k}^{\max}(X).$$

Once X is clear from the context, we will drop it from the notation, writing just $F_{i,j,k}^{\max}$ and $F_{i,j,k}^{\min}$

Lemma 4.11. Let X be a graded poset. Suppose that $i \leq j \leq k \leq \ell$ are integers. Then

$$F_{i,j,\ell}^{\min}F_{j,k,\ell}^{\min} \leq F_{i,k,\ell}^{\max}F_{i,j,k}^{\max} \quad and \quad F_{i,j,\ell}^{\max}F_{j,k,\ell}^{\max} \geq F_{i,k,\ell}^{\min}F_{i,j,k}^{\min}$$

Proof. Let $u \in X(i)$ and $z \in X(\ell)$ be incident; if there are no such u and z then both sides of both inequalities evaluate to 0. Write [u, z](j) for the set of j-faces of X lying between u and z. Then

$$F_{i,j,\ell}^{\min}F_{j,k,\ell}^{\min} \le \sum_{x \in [u,z](j)} \sum_{y \in [x,z](k)} 1 = \sum_{y \in [u,z](k)} \sum_{x \in [u,y](j)} 1 \le F_{i,k,\ell}^{\max}F_{i,j,k}^{\max}.$$

This proves the first inequality. The second inequality is shown similarly.

Definition 4.12 (Lower-Regular Graded Poset). A graded poset X is called lower-regular if for all integers $i \leq j \leq k$, we have $F_{i,j,k}^{\max} = F_{i,j,k}^{\min}$. In this case, we write $F_{i,j,k}$ for both quantities (and $F_{j,k} = F_{-1,j,k}$).¹⁵

Lower-regular graded posets are both common and better behaved than general graded posets. For such posets, the inequalities of Lemma 4.11 become an equality: $F_{i,j,\ell}F_{j,k,\ell} = F_{i,k,\ell}F_{i,j,k}$

Example 4.13. Simplicial complexes, cube complexes and the affine Grassmannian $AG_{d,n}(\mathbb{F})$ are lower regular graded posets. For a simplicial complex of dimension d, we have $F_{i,j,k} = \binom{k-i}{j-i}$ if $-1 \leq i \leq j \leq k \leq d$ and $F_{i,j,k} = 0$ otherwise.

Definition 4.14 (Lower-Irregularity of a *d*-Poset). Let X be a *d*-poset and let $-1 \le i \le j \le k \le d$ be integers. The the (i, j, k)-lower regularity of X is

$$L_{i,j,k} = L_{i,j,k}(X) = \frac{F_{i,j,k}^{\max}(X)}{F_{i,j,k}^{\min}(X)}$$

and we abbreviate $L_{-1,j,k}$ to $L_{j,k}$. The lower-irregularity of X is

$$L(X) = \max_{-1 \le i \le j \le k \le d} L_{i,j,k}(X).$$

Note that $L_{i,j,k}(X)$ is well-defined because the assumption that X is a *d*-poset guarantees that $F_{i,j,k}^{\min} \ge 1$. The reason is that every face of X contains in \emptyset_X (of dimension -1) and is contained some *d*-face, so every pair of incident faces in X is a part of a chain of faces $\emptyset = x_{-1} < x_0 < x_1 < \cdots < x_d$ with dim $x_{\ell} = \ell$ for all ℓ .

The lower-irregularity of X measures how far X is from being lower-regular. We always have $L(X) \ge 1$ and equality holds if and only if X is lower-regular.

We now consider properly weighted d-posets (X, w). The following fundamental lemmas use the constants $F_{i,j,k}^{\max}$ and $F_{i,j,k}^{\min}$ to relate the weights of subsets of X and its links. They will be used repeatedly later on. Note that the inequalities in the lemmas become equalities when X is lower-regular.

Lemma 4.15. Let (X, w) a properly weighted d-poset, let $-1 \le i \le j \le d$ and let $z \in X(i)$. Then

$$\frac{F_{i,j,d}^{\min}F_{i,d}^{\min}}{F_{j,d}^{\max}} \le \frac{w(X(j)_z)}{w(z)} \le \frac{F_{i,j,d}^{\max}F_{i,d}^{\max}}{F_{j,d}^{\min}}$$

¹⁵Caution: This condition is stronger than the lower regularity considered in [KT23].

Proof. We have

$$w(X(j)_z) = \sum_{x \in X(j)_z} w(x) = \sum_{x \in X(j)_z} \sum_{y \in X(d)_x} \frac{w(y)}{|y(j)|} = \sum_{y \in X(d)_z} \sum_{x \in X(j): z \le x \le y} \frac{w(y)}{|y(j)|}$$
$$\leq \sum_{y \in X(d)_z} \frac{F_{i,j,d}^{\max} w(y)}{|y(j)|} \le \sum_{y \in X(d)_z} \frac{F_{i,j,d}^{\max} F_{i,d}^{\max}}{F_{j,d}^{\min}} \cdot \frac{w(y)}{|y(i)|} = \frac{F_{i,j,d}^{\max} F_{i,d}^{\max}}{F_{j,d}^{\min}} w(z).$$

This gives the right inequality. The left inequality is similar.

Lemma 4.16. Let (X, w) a weighted d-poset, let $-1 \le i \le j \le d$ and let $\emptyset \ne A \subseteq X(j)$. Then

$$F_{i,j}^{\min} \le \frac{\sum_{z \in X(i)} w(A_z)}{w(A)} \le F_{i,j}^{\max}.$$

Proof. We have

$$\sum_{z \in X(i)} w(A_z) = \sum_{z \in X(i)} \sum_{x \in A_z} w(x) = \sum_{x \in A} \sum_{z \in x(i)} w(x) \le \sum_{x \in A} F_{i,j}^{\max} w(x) = F_{i,j}^{\max} w(A).$$

This gives the right inequality. The left inequality is similar.

Lemma 4.17. Let (X, w) a properly weighted d-poset, let $-1 \le i \le j \le d$, let $z \in X(i)$ and let $x \in X(j)_z$. Then

$$w(X(d)_z)^{-1} \cdot \frac{F_{j,d}^{\min}}{F_{i,j,d}^{\max}} \le \frac{w_z(x)}{w(x)} \le w(X(d)_z)^{-1} \cdot \frac{F_{j,d}^{\max}}{F_{i,j,d}^{\min}}.$$

Proof. We have

$$w_{z}(x) = \sum_{y \in X(d)_{x}} \frac{w(y)}{w(X(d)_{z})|y(j)_{z}|} \le \sum_{y \in X(d)_{x}} \frac{w(y)}{w(X(d)_{z})F_{i,j,d}^{\min}} \cdot \frac{F_{j,d}^{\max}}{|y(j)|}$$
$$= w(X(d)_{z})^{-1} \cdot \frac{F_{j,d}^{\max}}{F_{i,j,d}^{\min}} \sum_{y \in X(d)_{x}} \frac{w(y)}{|y(j)|} = w(X(d)_{z})^{-1} \cdot \frac{F_{j,d}^{\max}}{F_{i,j,d}^{\min}} w(x).$$

This gives the right inequality. The left inequality is similar.

Corollary 4.18. Let (X, w) a properly weighted d-poset, let $-1 \le i \le d$ and let $z \in X(i)$. Then

$$\frac{v(X(d)_z)}{F_{i,d}^{\max}} \le w(z) \le \frac{w(X(d)_z)}{F_{i,d}^{\min}}$$

Proof. Apply Lemma 4.17 with x = z and observe that $w_z(z) = 1$ and $F_{i,i,d}^{\min} = F_{i,i,d}^{\max} = 1$.

Corollary 4.19. Let (X, w) be a properly weighted lower-regular d-poset and let $0 \le j \le d$. Then $(X(\le j), w|_{W(\le j)})$ is a properly weighted lower regular j-poset. In particular, for every $-1 \le i \le j$ and $x \in X(i)$, we have

$$w(x) = \sum_{y \in X(j)_x} \frac{w(y)}{|y(i)|}.$$

Proof. Let $i \in \{-1, \ldots, j\}$ and $x \in X(i)$. By Lemma 4.15 and the lower-regularity of X, we have $w(x) = \frac{F_{j,d}}{F_{i,j,d}F_{i,d}}w(X(j)_x)$. By Lemma 4.11, the right hand size equals $\frac{w(X(j)_x)}{F_{i,j}} = \sum_{y \in X(j)_x} \frac{w(y)}{|y(i)|}$, so we proved the equality in the corollary. All other assertions are now straightforward.

Remark 4.20. Corollary 4.19 implies that if (X, w) is a properly weighted lower-regular poset, then w is a standard weight function in the sense of [KT23].

4.5 Degree and Upper-Regular Posets

Again, let X be a graded poset. Given integers $i \leq j$ such that X has an *i*-face incident to a *j*-face, the maximal (resp. minimal) (i, j)-degree of X is largest (respect. smallest) possible number of *j*-faces containing an *i*-face in X. The maximal (i, j)-degree and minimal (i, j)-degree of X are denoted

$$D_{i,j}^{\max}(X)$$
 and $D_{i,j}^{\min}(X)$,

respectively. If no *i*-face in X is incident to a *j*-face, we set $D_{i,j}^{\max}(X) = D_{i,j}^{\min}(X) = 0$. When X is clear from the context, we shall simply write $D_{i,j}^{\max}$ and $D_{i,j}^{\min}$.

Definition 4.21 (Upper-Regular Graded Poset). A graded poset X is called upper regular if $D_{i,j}^{\max}(X) = D_{i,j}^{\min}(X)$ for all integers $i \leq j$.

Definition 4.22 (Upper-Irregularity of a *d*-Poset). Let X be a *d*-poset. For integers $-1 \le i \le j \le d$, then (i, j)-upper irregularity of X is

$$U_{i,j} = U_{i,j}(X) = \frac{D_{i,j}^{\max}}{D_{i,j}^{\min}}.$$

The upper-irregularity of X is

$$U(X) = \max_{-1 \le i \le j \le d} U_{i,j}(X).$$

As with lower-regularity, the upper irregularity of a *d*-poset X measure how far it is from being upper-regular — we have $U(X) \ge 1$ and equality holds and if and only X is upper-regular. Unfortunately, upper regular *d*-posets are not so common for $d \ge 2$.

Example 4.23. (i) A graph (viewed as a poset) is upper-regular if and only if it is a regular graph in the usual sense, i.e., there is $k \in \mathbb{N}$ such that every vertex belongs to exactly k edges.

(ii) The explicit Ramanujan complexes of [LSV05] (see also [Li04]) are famous examples of simplicial complexes that are *high dimensional expanders*. They are upper-regular in dimension 2 but are not upper-regular in dimensions 3 and above.

(iii) The double Cayley complex $\operatorname{Cay}(A, G, B)$ associated to a group G and two symmetric generating sets $A, B \subseteq G$ — see §1.7 or §9.1 — is a square complexes that is upper-regular if and only if |A| = |B|. Its upper-irregularity is $\max\{\frac{|A|}{|B|}, \frac{|B|}{|A|}\}$.

Proposition 4.24. Let X be a d-poset, let w be the natural weight function of X and let $-1 \le i \le d$. Then for every $x, x' \in X(i)$, we have $w(x) \le U_{i,d}L_{i,d}w(x')$.

Proof. We have

$$w(x) = \sum_{y \in X(d)_x} \frac{w(y)}{|y(i)|} = \frac{1}{|X(d)|} \sum_{y \in X(d)_x} \frac{1}{|y(i)|} \le \frac{1}{|X(d)|} \frac{D_{i,d}^{\max}}{F_{i,d}^{\min}} = \frac{1}{|X(d)|} \frac{D_{i,d}^{\min}}{F_{i,d}^{\max}} U_{i,d} L_{i,d}.$$

Similarly, $w(x') \ge \frac{1}{|X(d)|} \frac{D_{i,d}^{\min}}{F_{i,d}^{\max}}$ and the proposition follows.

4.6 Orientation

Recall that, given a poset X, we write $x \triangleleft y$ to denote that x < y and there is no $z \in X$ with x < z < y. Recall also that rings are assumed to be commutative and R^{\times} denotes the group of invertible elements in a ring R.

Definition 4.25 (Oriented Poset). Let X be a graded poset. Let R be a commutative ring, e.g. \mathbb{Z} . An R-orientation on X is a function

$$(x,y) \mapsto [y:x] : \{(x,y) \in X \times X : x \triangleleft y\} \to R^{\times}$$

such that whenever $x, z \in X$ satisfy $x \leq z$ and dim $z = \dim x + 2$, we have

$$\sum_{y:x < y < z} [z:y][y:x] = 0$$

in R. An R-oriented poset is a graded poset X endowed with an R-orientation [:].

We will often be agnostic about which *R*-orientation is chosen and only care that an *R*-orientation exists. In this case, we will say that our poset is *R*-orientable. If X admits a \mathbb{Z} -orientation [:], then X admits an *R*-orientation for any commutative ring *R* defined by $(y, x) \mapsto [y : x]\mathbf{1}_R$.

Example 4.26 (Regular Cell Complexes are \mathbb{Z} -Orientable). Every regular cell complex X admits a \mathbb{Z} -orientation, and therefore an *R*-orientation for every commutative ring *R*, such that $[v: \emptyset] = 1$ for every $v \in X(0)$. In particular, simplicial complexes and cube complexes are *R*-orientable.

In more detail, let \mathcal{X} be a topological realization of X. Then every $x \in X$ with dim $x \geq 0$ corresponds to a topological embedding $j_x : D^n \to \mathcal{X}$ of an *n*-dimensional disc in \mathcal{X} . Choose an orientation for every cell $j_x : D^n \to \mathcal{X}$ (i.e., a generator of $a_x \in \pi_n(D^n, \partial D^n) \cong \mathbb{Z}$). Then, given nonempty faces x, y with $x \triangleleft y$, take [y : x] be 1 if the orientations of the discs of y and x agree and -1 otherwise. When x is empty, just set [y : x] = 1.

In practice, choosing an orientation for the faces of X means choosing a sign (+ or -) for every vertex $v \in X(0)$, a direction for every edge $e \in X(1)$ (i.e. labelling one its vertices with a + and the other with a -), a spin for every 2-dimensional $x \in X(2)$ (i.e. a direction for every edge of x such that two edges sharing a vertex give opposite signs to that vertex), and so on. In general, choosing an orientation to an *i*-face x ($i \ge 0$) amounts to choosing an orientation for every (i - 1)-face of x such that every two (i - 1)-faces which share an (i - 2)-face restrict to opposite orientations on that face.

Example 4.27. (i) Let \mathbb{F}_q be a finite field with q elements and let $X = \operatorname{AG}_{d,n}(\mathbb{F}_q)$ (Example 4.2(ii)). Then X admits a $\mathbb{Z}/(q+1)\mathbb{Z}$ -orientation given by [y:x] = 1 for every $x, y \in X$ with $x \triangleleft y$. This is an orientation because for every $x, z \in X$ with $x \leq z$ and dim $z = \dim x + 2$, we have $q+1 \mid \#\{y \in X : x < y < z\}$, so $\sum_{y:x < y < z} [z:y][y:x] = \sum_{y:x < y < z} 1 = 0$ in $\mathbb{Z}/(q+1)\mathbb{Z}$. On the other hand, parity considerations show that $\operatorname{AG}_{d,n}(\mathbb{F}_q)$ has no $(\mathbb{Z}/2\mathbb{Z})$ -orientation when q is even, and hence no \mathbb{Z} -orientation.

(ii) A *linear* graded poset with at least 3 elements has no R-orientation for every nonzero commutative ring R.

Example 4.28. Let X be a graded poset and $[:]_X$ and R-orientation on X. Let $z \in X$. Then the restriction of $[:]_X$ to the link X_z is an orientation of X_z . We will always give X_z the orientation it inherits from from X.

5 Sheaves on Partially Ordered Sets

Sheaves on (certain) cell complexes, also called *cellular sheaves* were first considered by Shepard [She85]. The theory was further developed by Curry [Cur14], who also considered the dual notion

of cellular *cosheaves*. A more concise treatment appears in [HG19] (for regular cell complexes). The definition of sheaves on cell complexes extends naturally to general posets; this is briefly considered in [Cur14, §4.2.2] and [PK22]. We recall it here, and then define sheaf cohomology when the underlying poset is graded and oriented.

Recall our standing assumption that rings are commutative and all modules are left modules. Throughout, R is a ring.

5.1 Sheaves on Posets

Definition 5.1 (Sheaf on a Poset). Let R be a ring, e.g., \mathbb{Z} or a field \mathbb{F} , and let X be a poset. An R-sheaf \mathcal{F} on X consists of:

- an *R*-module $\mathcal{F}(x)$ for every $x \in X$;
- an R-linear map $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}} : \mathcal{F}(x) \to \mathcal{F}(y)$ for every $x, y \in X$ with x < y;

such that whenever x < y < z, we have

$$\operatorname{res}_{z \leftarrow y}^{\mathcal{F}} \circ \operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{res}_{z \leftarrow x}^{\mathcal{F}}.$$
(5.1)

In this case, we also define $\operatorname{res}_{x \leftarrow x}^{\mathcal{F}} = \operatorname{id}_{\mathcal{F}(x)}$, so that (5.1) also holds when $x \leq y \leq z$.

One can similarly define sheaves of abelian groups, but they are the same thing as \mathbb{Z} -sheaves.

The maps $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ are called the *restriction maps* of \mathcal{F} . We will write $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}}(f)$ as $\operatorname{res}_{y \leftarrow x}(f)$ or just $f|_y$ when there there is no risk of confusion. One can similarly define *cosheaves* by reversing the direction of the restriction maps. We spell out the definition explicitly.

Definition 5.2 (Cosheaf on a Poset). Let R be a ring and let X be a poset. An R-cosheaf \mathcal{G} on X consists of:

- an *R*-module $\mathcal{G}(x)$ for every $x \in X$;
- an R-linear map $\operatorname{res}_{x \leftarrow y}^{\mathcal{G}} : \mathcal{G}(y) \to \mathcal{G}(x)$ for every $x, y \in X$ with x < y;

such that whenever x < y < z, we have

$$\operatorname{res}_{x \leftarrow y}^{\mathcal{G}} \circ \operatorname{res}_{y \leftarrow z}^{\mathcal{G}} = \operatorname{res}_{x \leftarrow z}^{\mathcal{G}}.$$
(5.2)

In this case, we also define $\operatorname{res}_{x\leftarrow x}^{\mathcal{G}} = \operatorname{id}_{\mathcal{F}(x)}$, so that (5.2) also holds when $x \leq y \leq z$.

The maps $\operatorname{res}_{x \leftarrow y}^{\mathcal{G}}$ are the *corestriction maps* of \mathcal{G} . A cosheaf on X is essentially the same thing as a sheaf on the opposite poset X^{op} (Example 4.4). It is beneficial to differ between these two notions because there are times where one needs to consider sheaves and cosheaves on the same poset X.

Example 5.3. Let M be an R-module. The constant R-sheaf on X associated to M is defined by setting $\mathcal{F}(x) = M$ for every M and $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}} = \operatorname{id}_M$ for every $x \leq y$. One can similarly define a constant cosheaf associated to M.

More sophisticated examples will be considered later.

5.2 Sheaf Cohomology

Suppose that X is an R-oriented graded poset, where R is a commutative ring. In this case, one can associate cohomology groups to R-sheaves on X and homology groups to R-cosheaves on X as follows.

Let \mathcal{F} be an *R*-sheaf on *X*. For every $i \in \mathbb{Z}$, define

$$C^{i} = C^{i}(X, \mathcal{F}) = \prod_{x \in X(i)} \mathcal{F}(x).$$

We call C^i the space of *i*-cochains on X with coefficients in \mathcal{F} , and given an *i*-cochain $f \in C^i = \prod_{x \in X(i)} \mathcal{F}(x)$, we write the x-component of f as f(x). As usual, the *i*-th coboundary map $d_i : C^i \to C^{i+1}$ is determined by

$$(d_i f)(y) = \sum_{x \in y(i)} [y : x] \operatorname{res}_{y \leftarrow x} f(x)$$

for all $f \in C^i$ and $y \in X(i+1)$. The subscript *i* in $d_i f$ will often be clear from the context, so we shall sometimes just write df. It is a recommended standard exercise to check that $d_{i+1} \circ d_i = 0$. Thus, we get a cochain complex of *R*-modules

$$C^{\bullet} = C^{\bullet}(X, \mathcal{F}) := [\dots \to C_{-1} \xrightarrow{d_{-1}} C_0 \xrightarrow{d_0} C_1 \xrightarrow{d_1} C_2 \to \dots].$$

Its R-modules of *i*-boundaries and *i*-cocycles are

$$B^i = B^i(X, \mathcal{F}) := \operatorname{im} d_{i-1}$$
 and $Z^i = Z^i(X, \mathcal{F}) := \ker d_i$.

Clearly, $B^i \subseteq Z^i$. The quotient module Z^i/B^i is known as $H^i(X, \mathcal{F})$ and called the *i*-th cohomology group of X with coefficients in \mathcal{F} , but this will not be needed in the sequel.

Remark 5.4. Beware that at this level of generality, $\mathrm{H}^{i}(X, \mathcal{F})$ may be nonzero for negative values of *i*. (To experts, we also caution that $\{\mathrm{H}^{i}(X, -)\}_{i\geq 0}$ are may not be the *right derived functors* of $\mathrm{H}^{0}(X, -)$, even when $X(i) = \emptyset$ for all i < 0.) In addition, the isomorphism classes of Z^{i} , B^{i} and $\mathrm{H}^{i}(X, \mathcal{F})$ may depend on the *R*-orientation of *X*. However, when *X* is a regular cell complex and $\mathcal{F}(\emptyset) = 0$, everything behaves as expected: $\mathrm{H}^{i}(X, \mathcal{F}) = 0$ for i < 0 and, as we shall see in §5.4, changing the *R*-orientation has no effect on the isomorphism class of C^{\bullet} . (Moreover, $\{\mathrm{H}^{i}(X, -)\}_{i\geq 0}$ are indeed the right derived functors of $\mathrm{H}^{0}(X, -)$, but that will not be needed here.)

The homology groups of a cosheaf \mathcal{G} on X are defined similarly, but with the following differences. One replaces $d_i: C^i \to C^{i+1}$ with the *i*-th boundary map $\partial_i: C^i \to C^{i-1}$ defined by

$$(\partial_i f)(y) = \sum_{x \in X(i)_y} \operatorname{res}_{y \leftarrow x} f(x),$$

and we get a chain complex:

$$\cdots \leftarrow C_{-1} \xleftarrow{\partial_0} C_0 \xleftarrow{\partial_1} C_1 \xleftarrow{\partial_2} C_2 \leftarrow \cdots$$

The *i*-boundaries and *i*-cycles are $B_i = B_i(X, \mathcal{F}) := \operatorname{im} \partial_{i+1}$ and $Z_i = Z_i(X, \mathcal{F}) := \operatorname{ker} \partial_i$ and the *i*-th homology is $\operatorname{H}^i(X, \mathcal{F}) = Z_i/B_i$.¹⁶

¹⁶If X is allowed to be infinite, then one should also replace $C^i = \prod_{x \in X(i)} \mathcal{F}(x)$ with the *R*-module of *i*-chains $C_i = C_i(X, \mathcal{F}) = \bigoplus_{x \in X(i)} \mathcal{F}(x)$. Otherwise, the summation in the definition of ∂_i is not always well-defined.
Remark 5.5. Let \mathcal{G} be a cosheaf on graded poset X. Define a sheaf \mathcal{G}^{op} on X^{op} (Example 4.4) by setting $\mathcal{G}^{\text{op}}(x^{\text{op}}) = \mathcal{G}(x)$ and $\operatorname{res}_{y^{\text{op}} \leftarrow x^{\text{op}}}^{\mathcal{G}^{\text{op}}} = \operatorname{res}_{x \leftarrow y}^{\mathcal{G}} (x < y)$. We call \mathcal{G}^{op} the *opposite sheaf* of the cosheaf \mathcal{G} . Clearly, $C^{i}(X, \mathcal{G}) = C^{-i}(X^{\text{op}}, \mathcal{G}^{\text{op}}) = \text{and } \partial_{i}^{\mathcal{G}} = d_{-i}^{\mathcal{G}^{\text{op}}}$. Thus, cosheaf homology may be realized as sheaf cohomology of the opposite sheaf.

Example 5.6. Let \mathcal{F} be an R-sheaf on an R-oriented hypergraph X (viewed as poset; Example 4.3). Then $C^0 = C^0(X, \mathcal{F}) = \prod_{v \in X(0)} \mathcal{F}(v)$, and after unfolding the definitions, one finds that Z^0 is the set of $f = (f(v))_{v \in X(0)} \in \prod_{v \in X(0)} \mathcal{F}(v)$ such that for every hyperedge $e \in X(1)$,

$$\sum_{v \in e(0)} [e:v] \operatorname{res}_{e \leftarrow v} f(v) = 0$$

5.3 Restricting Sheaves to The Links

Definition 5.7 (Sheaf Restricted to a Link). Let \mathcal{F} be an *R*-sheaf on a graded poset *X* and let $z \in X$. The restriction of \mathcal{F} to X_z is the *R*-sheaf \mathcal{F}_z obtained by restricting \mathcal{F} to X_z . That is, $\mathcal{F}_z(x) = \mathcal{F}(x)$ and $\operatorname{res}_{y \leftarrow x}^{\mathcal{F}_z} = \operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ for all $x, y \in X_z$ with $x \leq y$.

Suppose now that X is an R-oriented graded poset and \mathcal{F} is an R-sheaf on X. For every $z \in X$ and $i \in \mathbb{Z}$, we define maps

$$f \mapsto f_z : C^i(X, \mathcal{F}) \to C^{i-\dim z - 1}(X_z, \mathcal{F}_z),$$

$$f' \mapsto f'^z : C^{i-\dim z - 1}(X_z, \mathcal{F}_z) \to C^i(X, \mathcal{F}),$$

by

$$f_z(x') = f(x') \qquad \forall x' \in X_z,$$

$$f'^z(x) = \begin{cases} f'(x) & x \in X_z \\ 0 & x' \notin X_z \end{cases} \qquad \forall x \in X.$$

That is, f_z is the restriction of f to X_z , and f'^z is obtained by extending f' from X_z to X by setting it to be 0 on *i*-faces not in X_z .

A straightforward computation (and a recommended exercise) gives the following lemma.

Lemma 5.8. Let \mathcal{F} be an R-sheaf on an R-oriented graded poset X, let $z \in X$ and let $i \in \mathbb{Z}$. Write d_z for the coboundary map of \mathcal{F}_z . Then, for every $f \in C^{i-\dim z-1}(X_z, \mathcal{F}_z)$,

$$(d_z f)^z = d(f^z).$$

5.4 Independence of The Orientation for Regular Cell Complexes

Let X be an R-oriented graded poset and let \mathcal{F} be an R-sheaf on X. We noted earlier that the definition of the cochain complex $C^{\bullet}(X, \mathcal{F})$ depends on the R-orientation of X. We now show that when X is a regular cell complex (e.g. a simplicial complex or a cube complex) this choice has essentially no effect. We shall first need two lemmas. Given $x \in X$, an x-flag is a sequence $f = (x_{-1}, x_0, x_1, \ldots, x_i)$ of faces in X such that $\emptyset = x_{-1} < x_0 < \cdots < x_i = x$.

Lemma 5.9. Let X be a regular cell complex and let $x \in X$. Then, for any two x-flags f and f', there is a sequence of x-flags $f = f_0, f_1, \ldots, f_n = f'$ in which every two consecutive x-flags differ by at most one term.

Proof. This is well-known, but we sketch the proof for the sake of completeness. Given x-flags f and f', we write $f \sim f'$ to denote that there is a sequence of x-flags as in the lemma. The proof is by induction on $i := \dim x$. The cases i = -1 and i = 0 are clear, so assume that i > 0 and write $f := (x_{-1}, x_0, \ldots, x_i)$ and $f' = (x'_{-1}, x'_0, \ldots, x'_i)$. Suppose that x_{i-1} and x'_{i-1} share an (i-2)-face y, and choose a y-flag $(y_{-1}, \ldots, y_{i-2})$. Then, by the induction hypothesis, $(x_{-1}, \ldots, x_{i-2}, x_{i-1}) \sim (y_{-1}, \ldots, y_{i-2}, x_{i-1})$ and $(y_{-1}, \ldots, y_{i-2}, x'_{i-1}) \sim (x'_{-1}, \ldots, x'_{i-1})$. This means that $f \sim (y_{-1}, \ldots, y_{i-2}, x_{i-1}, x_i)$ and $f' \sim (y_{-1}, \ldots, y_{i-2}, x'_{i-1}, x_i)$, so $f \sim f'$. In general, since the topological realization of $X_{< x} = \{y \in X : y < x\}$ is a sphere, we can find a sequence $x_{i-1} = z^{(0)}, \ldots, z^{(\ell)} = x'_{i-1}$ of (i-1)-faces of X such that $z^{(k-1)}$ and $z^{(k)}$ share an (i-2) face for every k. By choosing a flag $f^{(k)} = (\cdots, z^{(k)}, x)$ for every $0 < k < \ell$ and using what we have shown, we see that $f \sim f^{(1)} \sim \cdots \sim f^{(\ell-1)} \sim f'$.

Lemma 5.10. Let X be a regular cell complex, let [:] and (:) be two R-orientations on X and let $x \in X$. Choose an x-flag $f = (x_{-1}, \ldots, x_i)$ and define $c_x = \prod_{n=0}^{i} [x_i : x_{i-1}]^{-1}(x_i : x_{i-1}) \in \mathbb{R}^{\times}$. Then c_x does not depend on the choice of f.

Proof. Thanks to Lemma 5.9, it is enough to show that c_x does not change if we replace x_j (-1 < j < i) with a different face x'_j lying between x_{j-1} and x_{j+1} . Indeed, since X is a regular cell complex, x_j and x'_j are the only faces between x_{j-1} and x_{j+1} , which means that $[x_{j+1} : x_j][x_j : x_{j-1}] + [x_{j+1} : x'_j][x'_j : x_{j-1}] = 0$, or rather, $[x_{j+1} : x_j][x_j : x_{j-1}] = -[x_{j+1} : x'_j][x'_j : x_{j-1}]$. Similarly, $(x_{j+1} : x_j)(x_j : x_{j-1}) = -(x_{j+1} : x'_j)(x'_j : x_{j-1})$, and it follows that

$$[x_{j+1}:x_j]^{-1}(x_{j+1}:x_j)[x_j:x_{j-1}]^{-1}(x_j:x_{j-1}) = [x_{j+1}:x_j']^{-1}(x_{j+1}:x_j')[x_j':x_{j-1}]^{-1}(x_j':x_{j-1}).$$

As a result, c_x does not change when we replace x_j with x'_j .

Proposition 5.11. Let X be a regular cell complex, let \mathcal{F} be an R-sheaf on X, let [:] and (:) be two R-orientations on X. Denote by C^{\bullet} and C'^{\bullet} the cochain complexes associated to X and \mathcal{F} using the orientations [:] and (:), respectively. (The *i*-coboundary map of C'^{\bullet} is denoted d'_i .) For every $x \in X$, define $c_x \in R^{\times}$ as in Lemma 5.10, and for every $i \in \mathbb{Z}$, define $T_i = \prod_{x \in X(i)} c_x \operatorname{id}_{\mathcal{F}(x)} : C^i \to C'^i$, that is, $T_i(f) = (c_x f(x))_{x \in X(i)}$. Then $T = (T_i)_{i \in \mathbb{Z}}$, defines an isomorphism of cochain complexes from C^{\bullet} to C'^{\bullet} , *i.e.*, each T_i is an R-module isomorphism and the following diagram commutes.

$$C^{-1} \xrightarrow{d_{-1}} C_0 \xrightarrow{d_0} C_1 \xrightarrow{d_1} \cdots$$

$$\downarrow T_{-1} \qquad \downarrow T_0 \qquad \downarrow T_1$$

$$C'^{-1} \xrightarrow{d'_{-1}} C'_0 \xrightarrow{d'_0} C'_1 \xrightarrow{d'_1} \cdots$$

In particular, the map T_i induces isomorphisms $Z^i \to Z'^i$ and $B^i \to B'^i$, where Z'_i and B'_i denote the *i*-cocycles and *i*-cochains of C'^{\bullet} .

Proof. It is clear that T_i is bijective. It remains to check that $T_{i+1} \circ d_i = d'_i \circ T_i$ for every $i \ge -1$. Let $f \in C^i = \prod_{x \in X(i)} \mathcal{F}(x)$. Then for every $y \in X(i+1)$,

$$(T_{i+1}d_if)(y) = c_y \sum_{x \in y(i)} [y:x] \operatorname{res}_{y \leftarrow x} f(x) = \sum_{x \in y(i)} (y:x) \operatorname{res}_{y \leftarrow x} (c_x f(x)) = (d'_i T_i f)(y),$$

where the second equality holds because $c_y[y:x] = (x:y)c_x$.

An analogue of Proposition 5.11 holds for cosheaves. We omit the details.

5.5 Aside: Constraint Systems

Our discussion of sheaf cohomology and our main results (Theorem 11.2 and its simpler versions in Section 8) actually apply to a slightly more general structure which we call a constraint system.

Let R be a commutative ring. An R-constraint system S on a graded poset X consists of

- an *R*-module $\mathcal{S}(x)$ for every $x \in X$ and
- an R-homomorphism $r_{y\leftarrow x}^{\mathcal{S}}: \mathcal{S}(x) \to \mathcal{S}(y)$ for every $x, y \in X$ with $x \triangleleft y$

such that for every $x, z \in X$ with dim $z = \dim x + 2$, we have

$$\sum_{y:x < y < z} r_{z \leftarrow y}^{\mathcal{S}} \circ r_{y \leftarrow x}^{\mathcal{S}} = 0.$$

Note that, unlike the case of sheaves, $r_{y\leftarrow x}^{\mathcal{S}}$ is only defined when $x \triangleleft y$. Also, it need not be the case that $r_{z\leftarrow y}^{\mathcal{S}} \circ r_{y\leftarrow x}^{\mathcal{S}} = r_{z\leftarrow x}^{\mathcal{S}}$. One can dually define *R*-constraint cosystems on X by reversing the direction of the maps $r_{y\leftarrow x}^{\mathcal{S}}$; we omit the details.

Given an *R*-constraint system on S on X, we let $C^i = C^i(X, S) := \prod_{x \in X(i)} S(x)$ and define a coboundary map $d^i : C^i \to C^{i+1}$ exactly as we did for sheaves, but without the invoking the factor [y:x]. This gives rise to a cochain complex and a cohomology theory of *R*-constraint systems, and the proof of our main result applies verbatim in this context. Note, however, that X is not required to have an *R*-orientation as in the case of sheaves.

Examples of constraint systems include the *codes modeled over 2-layer systems* of [KO21]; they can be realized as constraint systems on 2-posets. In addition *based chain complexes* of [PK22] can realized as the chain complexes of constraint cosystems on graded posets.

Every *R*-sheaf \mathcal{F} on an *R*-oriented poset *X* gives rise to an *R*-constraint system \mathcal{S} on *X* by setting $\mathcal{S}(x) = \mathcal{F}(x)$ and $r_{y \leftarrow x}^{\mathcal{S}} = [y : x] \operatorname{res}_{y \leftarrow x}^{\mathcal{F}}$ when $x \triangleleft y$. In this case, $C^i(X, \mathcal{F}) = C^i(X, \mathcal{S})$ and $d_i^{\mathcal{F}} = d_i^{\mathcal{S}}$, so \mathcal{F} and \mathcal{S} give the same modules of cocycles and coboundaries. By constant, not all *R*-constraint systems are induced from sheaves and orientations in this manner. (Indeed, X may not be *R*-orientable.)

We use sheaves and not constraint systems in this work because the theory of sheaves is more developed and more intuitive, and all the examples of constraint systems that are of interest to us naturally arise from sheaves.

6 Locally Testable Codes from Sheaves and Cosheaves

In this section, we explain in detail how sheaves on graded posets give rise to error correcting codes. These codes come equipped with a natural tester, the soundness of which is governed by the cosystolic expansion of the sheaf at hand. The entire discussion dualizes to cosheaves thanks to Remark 5.5.

Throughout, R is a (commutative) ring, e.g. a finite field, and X is an R-oriented graded poset.

6.1 Cocycle Codes

Let \mathcal{F} be an *R*-sheaf on *X* and let $i \in \mathbb{Z}$ be an integer such that X(i) and X(i+1) are nonempty. Suppose further that there is an *R*-module Σ such that $\mathcal{F}(x) = \Sigma$ for every $x \in X(i)$. Then $C^i = C^i(X, \mathcal{F}) = \Sigma^{X(i)}$ and so we may view $Z^i = Z^i(X, \mathcal{F})$ as a code inside $C^i = \Sigma^{X(i)}$; it is called the *i*-cocycle code of (X, \mathcal{F}) . The constraints defining Z^i inside $C^i = \Sigma^{X(i)}$ give rise to a natural tester for this code: Given $f \in \Sigma^{X(i)}$, choose $y \in X(i+1)$ uniformly at random, read f(x) for every $x \in y(i)$, and accept f if and only if $df(y) = \sum_{x \in y(i)} \operatorname{res}_{y \leftarrow x} f(x) = 0$.

Let $w_{\text{uni}} : X \to \mathbb{R}_+$ denote the uniform weight function (Example 4.8(ii)). Given $f, g \in C^j$, let $\text{supp } f = \{x \in X(j) : f(x) \neq 0 \text{ in } \mathcal{F}(x)\}$ and define

$$||f||_{\text{uni}} = w_{\text{uni}}(\operatorname{supp} f)$$
 and $\operatorname{dist}_{\text{uni}}(f,g) = ||f-g||_{\text{uni}}$

For j = i, these are just the normalized Hamming norm and distance in $C^i = \Sigma^{X(i)}$. Now, by definition, the natural tester of $Z^i \subseteq \Sigma^{X(i)}$ has soundness $\geq \mu$ if and only if

$$\|df\|_{\text{uni}} \ge \mu \operatorname{dist}_{\text{uni}}(f, Z^i) \qquad \forall f \in C^i \tag{6.1}$$

and the distance of $Z^i \subseteq \Sigma^{X(i)}$ is at least δ if and only if

$$\|f\|_{\mathrm{uni}} \ge \delta \qquad \forall f \in Z^i - \{0\}.$$
(6.2)

Usually, $B^i = im(d_{i-1})$ contains short vectors, so we can expect the distance of Z^i to be large only if $B^i = 0$.

Conditions (6.1) and (6.2) may be seen as conditions on the expansion of the *i*-th coboundary map $d_i: C^i \to C^{i+1}$, and indeed, upon changing the weight function w_{uni} and replacing $Z^i - \{0\}$ with $Z^i - B^i$, we recover the definition of a (μ, δ) -cosystolic expander in dimension *i*, which we now discuss in detail.

6.2 Cosystolic Expansion

Suppose henceforth that $w : X \to \mathbb{R}_+$ be a normalized weight function; typically, we would like w to be proper. Given an R-sheaf \mathcal{F} on X and $f, g \in C^i(X, \mathcal{F})$, set $||f||_w = w(\operatorname{supp}(f))$ and $\operatorname{dist}_w(f,g) = ||f-g||_w$. We will drop the subscript w when it is clear from the context.

Definition 6.1 (Cosystolic Expansion of Sheaves). Let (X, w) and \mathcal{F} be as above and let $i \in \mathbb{Z}$. The *i*-cosystolic expansion of \mathcal{F} (w.r.t. w), denoted $\csc_i(X, w, \mathcal{F})$ is the supremum of the set of $\varepsilon \in [0, \infty)$ such that

$$\|df\|_w \ge \varepsilon \operatorname{dist}_w(f, Z^i) \qquad \forall f \in C^i.$$

The *i*-cocycle distance of \mathcal{F} (w.r.t. w) is

$$\operatorname{ccd}_i(X, w, \mathcal{F}) := \inf\{\|f\|_w \mid f \in Z^i - B^i\}.$$

Given $\varepsilon, \delta > 0$, we say that (X, w, \mathcal{F}) is an (ε, δ) -cosystolic expander in dimension i if $\operatorname{cse}_i(\mathcal{F}) \ge \varepsilon$ and $\operatorname{ccd}_i(\mathcal{F}) \ge \delta$.

When X is a d-poset, the *i*-cosystolic expansion of \mathcal{F} and *i*-cocycle distance of \mathcal{F} are defined to be $\operatorname{cse}_i(\mathcal{F}) := \operatorname{cse}_i(X, w_{\operatorname{nat}}, \mathcal{F})$ and $\operatorname{ccd}_i(\mathcal{F}) := \operatorname{ccd}(X, w_{\operatorname{nat}}, \mathcal{F})$, where w_{nat} is the natural weight function of X.

Suppose that there is a finite *R*-module Σ such that $\mathcal{F}(x) = \Sigma$ for all $x \in X(i)$. Then $\operatorname{cse}_i(X, w_{\operatorname{uni}}, \mathcal{F})$ is precisely the soundness of the natural tester of the cocycle code $Z^i = Z^i(X, \mathcal{F}) \subseteq C^i(X, \mathcal{F}) = \Sigma^{X(i)}$, and provided that $B^i(X, \mathcal{F}) = 0$ (e.g., if $\mathcal{F}(y) = 0$ for all $y \in X(i-1)$), $\operatorname{ccd}_i(X, w_{\operatorname{uni}}, \mathcal{F}) = \delta(Z^i)$.

Unfortunately, we cannot use these observations directly because our results about cosystolic expansion and cocycle distance will only apply when w is proper, and that is often not the case for w_{uni} . Nevertheless, when our weight function $w: X \to \mathbb{R}_+$ is not too far from being uniform, we can effectively relate $\csc_i(X, w, \mathcal{F})$ and $\operatorname{ccd}_i(X, w, \mathcal{F})$ to the soundness and distance of $Z^i \subseteq \Sigma^{X(i)}$ by means of the following lemma.

Lemma 6.2. Let \mathcal{F} , X, i and Σ be as in §6.1 and let $w : X \to \mathbb{R}_+$ be a normalized weight function. Assume further that there are $M, M' \in [1, \infty)$ such that $w(x) \leq Mw(y)$ for every $x, y \in X(i)$ and $w(x) \leq M'w(y)$ for every $x, y \in X(i+1)$. Then:

- (i) If $B^i = 0$, then the relative distance of $Z^i \subseteq \Sigma^{X(i)}$ is at least $\frac{1}{M} \operatorname{ccd}_i(X, w, \mathcal{F})$.
- (ii) The soundness of the natural tester of $Z^i \subseteq \Sigma^{X(i)}$ is at least $\frac{1}{MM'} \operatorname{cse}_i(X, w, \mathcal{F})$.

Moreover, if we would modify the natural tester to choose an (i+1)-face according to the distribution $w|_{X(i+1)}$, then its soundness would be at least $\frac{1}{M} \operatorname{cse}_i(\mathcal{F})$.

Proof. We first observe that $M^{-1}w_{\text{uni}}w(x) \leq w(x) \leq Mw_{\text{uni}}w(x)$ for every $x \in X(i)$, and thus $M^{-1}||f||_{\text{uni}} \leq ||f||_w \leq M||f||_{\text{uni}}$ for all $f \in C^i = C^i(X, \mathcal{F})$. To see the right inequality, observe that

$$|X(i)|w(x) = \sum_{y \in X(i)} w(x) \le \sum_{y \in X(i)} Mw(y) = M = M|X(i)|w_{\text{uni}}(x).$$

The left inequality is shown similarly. In the same way, $M'^{-1} ||g||_{\text{uni}} \leq ||g||_w \leq M' ||g||_{\text{uni}}$ for all $g \in C^{i+1}$. We now prove (i) and (ii).

(i) Let $0 \neq f \in Z^i$. Then $f \in Z^i - B^i$ because $B^i = 0$, and thus, $||f||_{\text{uni}} \geq M^{-1}||f|| \geq M^{-1} \operatorname{ccd}_i(X, w, \mathcal{F})$.

(ii) For every $f \in C^i$, we have

$$\|d_i f\|_{\mathrm{uni}} \ge M'^{-1} \|d_i f\|_w \ge M'^{-1} \operatorname{cse}_i(X, w, \mathcal{F}) \operatorname{dist}_w(f, Z^i) \ge \frac{\operatorname{cse}_i(X, w, \mathcal{F})}{MM'} \operatorname{dist}_{\mathrm{uni}}(f, Z^i).$$

A similar computation shows that $||d_i f||_w \ge \frac{1}{M} \operatorname{cse}_i(\mathcal{F}) \operatorname{dist}_{\operatorname{uni}}(f, Z^i)$ and this gives the last assertion of the lemma.

6.3 Coboundary Expansion

We will also need to consider a stronger variant of cosystolic expansion, called *coboundary expansion*.

Definition 6.3 (Coboundary Expansion of Sheaves). With notation as in Definition 6.1, the icoboundary expansion of \mathcal{F} (w.r.t. w), denoted $\operatorname{cbe}_i(X, w, \mathcal{F})$ is the smallest $\varepsilon \in [0, \infty)$ such that

$$\|df\|_w \ge \varepsilon \operatorname{dist}_w(f, B^i(X, \mathcal{F})).$$

We say that (X, w, \mathcal{F}) is an ε -cosystolic expander in dimension i if $\operatorname{cbe}_i(X, w, \mathcal{F}) \geq \varepsilon$.

When X is a d-poset, the i-coboundary expansion of \mathcal{F} is $\operatorname{cbe}_i(\mathcal{F}) := \operatorname{cbe}_i(X, w_{\operatorname{nat}}, \mathcal{F})$, where w_{nat} is the natural weight function of X.

If $\operatorname{cbe}_i(X, w, \mathcal{F}) > 0$, then we must have $Z^i = B^i$, or rather, $\operatorname{H}^i(X, \mathcal{F}) = 0$. In the context of §6.1, the *i*-coboundary expansion may be thought of as measuring the soundness of the natural tester of $Z^i \subseteq \Sigma^{X(i)}$, but for the code $B^i \subseteq \Sigma^{X(i)}$ (which usually has poor distance if $B^i \neq 0$).

Coboundary expansion of d-posets in dimensions -1 and 0 appears in some of our main results, so it is worthwhile to unfold the definition in these cases. This can be informally summarized as follows:

- Coboundary expansion in dimension -1 is similar to the relative distance of a code.
- Coboundary expansion in dimension 0 is similar to agreement testability (§2.4).

Example 6.4 (Coboundary Expansion in Dimension -1). Let (X, w) be a normalized weighted *R*-oriented *d*-poset $(d \ge 0)$ and let \mathcal{F} be an *R*-sheaf on *X*. Then $C^{-1} = C^{-1}(X, \mathcal{F}) = \mathcal{F}(\emptyset)$ and $d_{-1} : \mathcal{F}(\emptyset) \to C^0 = \prod_{v \in X(0)} \mathcal{F}(v)$ is given by $df = ([v : \emptyset] \operatorname{res}_{v \leftarrow \emptyset}(f))_{v \in X(0)}$. Moreover, we have $\operatorname{dist}(f, B^{-1}) = 1$ for every $f \in C^{-1} - B^{-1} = \mathcal{F}(\emptyset) - \{0\}$. This means that

$$\operatorname{cbe}_{-1}(X,\mathcal{F}) = \begin{cases} 0 & \ker d_{-1} \neq 0\\ \inf\{\|g\| \,|\, g \in B^0 - \{0\}\} & \ker d_{-1} = 0. \end{cases}$$

Since $\|([v:\emptyset]\operatorname{res}_{v\leftarrow\emptyset}(f))_{v\in X(0)}\| = \|(\operatorname{res}_{v\leftarrow\emptyset}(f))_{v\in X(0)}\|$, we can disregard the orientation and get that

$$\operatorname{cbe}_{-1}(X,\mathcal{F}) = \begin{cases} 0 & \ker \tilde{d}_{-1} \neq 0\\ \min\{\|g\| \,|\, g \in (\operatorname{im} \tilde{d}_{-1}) - \{0\}\} & \ker \tilde{d}_{-1} = 0. \end{cases}$$

where, $\tilde{d}_{-1} : \mathcal{F}(\emptyset) \to \prod_{v \in X(0)} \mathcal{F}(v)$ is given by $\tilde{d}_{-1}f = (\operatorname{res}_{v \leftarrow \emptyset}(f))_{v \in X(0)}$. In particular, if \tilde{d}_{-1} is injective, w is the uniform weight function and there is an R-module Σ such that $\mathcal{F}(v) = \Sigma$ for every $v \in X(0)$, then $\operatorname{cbe}_{-1}(X, \mathcal{F})$ is the relative distance of the code $\operatorname{im}(\tilde{d}_{-1})$ inside $\prod_{v \in X(0)} \mathcal{F}(v) = \Sigma^{X(0)}$.

Example 6.5 (Coboundary Expansion in Dimension 0). Let X be a regular cell complex and let $w: X \to \mathbb{R}_+$ be a normalized weight function. We choose a \mathbb{Z} -orientation on X such that $[v: \emptyset] = 1$ for every $v \in X(0)$; see Example 4.26. This implies that every edge $e \in X(1)$ has a unique vertex u with [e:u] = 1 and a unique vertex v with [e:v] = -1; denote the former by e^+ and the latter by e^- .

Let \mathcal{F} be an *R*-sheaf on *X*. Then for every $f \in C^0(X, \mathcal{F}) = \prod_{v \in X(0)} \mathcal{F}(v)$, we have $df = (\operatorname{res}_{e \leftarrow e^+} f(e^+) - \operatorname{res}_{e \leftarrow e^-} f(e^-))_{e \in X(1)}$. This means that $\operatorname{cbe}_0(\mathcal{F})$ is the smallest $\kappa \geq 0$ such that

$$w(\{e \in X(1) : \operatorname{res}_{e \leftarrow e^+} f(e^+) \neq \operatorname{res}_{e \leftarrow e^-} f(e^-)\}) \ge \kappa \operatorname{dist}_w(f, B^0)$$

This is reminiscent of the soundness of an agreement tester (§2.4), and in fact, agreement testability may be realized as the 0-coboundary of a sheaf.

Indeed, let $(\{C_s\}_{s\in S}, G, w, \ell)$ be an agreement tester for a lifted code $C = C(\{C_s\}_{s\in S}) \subseteq \Sigma^n$. Suppose moreover that Σ is an abelian group and every C_s is a subgroup of Σ^s . Let X be the 1-dimensional simplicial complex obtained from G by adding a single face of dimension -1. We extend w from G to X by setting $w(\emptyset_X) = 1$ and endow X with a \mathbb{Z} -orientation as above. Define a \mathbb{Z} -sheaf \mathcal{F} on X by setting

- $\mathcal{F}(\emptyset) = C := C(\{C_s\}) \subseteq \Sigma^n$,
- $\mathcal{F}(v) = C_{\ell(s)}$ for all $v \in X(0)$,
- $\mathcal{F}(e) \subseteq \Sigma^{\ell(e)}$ for all $e \in X(1)$,
- $\operatorname{res}_{v \leftarrow \emptyset} : C \to C_{\ell(v)}$ is given by $f \mapsto f|_{\ell(v)}$ for all $v \in X(0)$,
- $\operatorname{res}_{e\leftarrow v}: C_{\ell(v)} \to \Sigma^{\ell(e)}$ is given by $f \mapsto f|_{\ell(e)}$ for all $e \in X(1)$ and $v \in e(0)$.

Since $B^0(X, \mathcal{F}) = C$, our earlier observations imply readily that $(\{C_s\}_{s \in S}, G, w, \ell)$ has soundness κ if and only if $\operatorname{cbe}_0(X, w, \mathcal{F}) \geq \kappa$.

6.4 Independence of The Orientation

We continue to use the notation of §6.2. Recall that $Z^i = Z^i(X, \mathcal{F})$ and $B^i = B^i(X, \mathcal{F})$ depends on the implicit *R*-orientation we gave *X*, so a priori, $\operatorname{cse}_i(X, w, \mathcal{F})$, $\operatorname{ccd}_i(X, w, \mathcal{F})$ and $\operatorname{cbe}_i(X, w, \mathcal{F})$ depend on that choice as well. However, when *X* is a regular cell complex, Proposition 5.11 shows that the effect of changing the *R*-orientation on Z^i and B^i is just a coordinate-dependent scaling, which does not effect the norm $\|\cdot\|_w$ on C^i . As a result, in this special case, $\operatorname{cse}_i(X, w, \mathcal{F})$, $\operatorname{ccd}_i(X, w, \mathcal{F})$ and $\operatorname{cbe}_i(X, \mathcal{F})$ do not depend on the *R*-orientation of *X*. We record this observation:

Proposition 6.6. Let (X, w) be a normalized weighted regular cell complex and let \mathcal{F} be an R-sheaf on X. Then $\operatorname{cse}_i(X, w, \mathcal{F})$, $\operatorname{ccd}_i(X, w, \mathcal{F})$ and $\operatorname{cbe}_i(X, w, \mathcal{F})$ do not depend on the R-orientation chosen for X.

6.5 Dual Notions for Cosheaves

Let (X, w) be a normalized weighted *R*-oriented graded poset and let \mathcal{G} be an *R*-cosheaf on *X* such that $\mathcal{G}(x) = \Sigma$ for every $x \in X(i)$. Here, Σ is some *R*-module and $X(i) \neq \emptyset$. Then we may consider $Z_i = Z_i(X, \mathcal{G})$ as a code inside $C_i = \Sigma^{X(i)}$; such codes are called *i*-cycle codes. As with cocycle codes, the code $Z_i \subseteq \Sigma^{X(i)}$ has a natural tester: Given $f \in C^i$, choose $y \in X(i-1)$ uniformly at random, read f(x) for every $x \in X(i)_y$ and accept f if $\partial_i f(y) = 0$.

For a general cosheaf \mathcal{G} , the *i*-systolic expansion and *i*-cycle distance of (X, w, \mathcal{G}) , denoted $\operatorname{se}_i(X, w, \mathcal{G})$ and $\operatorname{cd}_i(X, w, \mathcal{G})$, respectively, are defined exactly as their counterparts in §6.2 by replacing d_i, Z^i, B^i with ∂_i, Z_i, B_i . In the setting of the last paragraph, an analogue of Lemma 6.2 holds and may be used to relate $\operatorname{cd}_i(X, w, \mathcal{G})$ and $\operatorname{se}_i(X, w, \mathcal{G})$ to the distance of $Z^i \subseteq \Sigma^{X(i)}$ and the soundness of its natural tester. This may also be derived directly from our discussion of cocycle codes using Remark 5.5.

7 No-Intersection Hypergraphs, Skeleton Expansion, Intersection Profiles

7.1 No-Intersection Hypergraphs

Recall that a weighted hypergraph means a weighted poset (X, w) such that X is concentrated in degrees 0 and 1 (Example 4.3). We define the underlying hypergraph of a graded poset X to be the subposet $Gr(X) := X(0) \cup X(1)$. For example, when X is regular cell complex, Gr(X) is the underlying graph of X in the usual sense. The weighted hypergraph underlying a weighted graded poset is

$$Gr(X, w) := (X(0) \cup X(1), w|_{X(0) \cup X(1)}).$$

Definition 7.1 (Related Weighted Hypergraph). A related weighted hypergraph is a triple (X, w, \sim) consisting of a weighted hypergraph (X, w) and a binary relation \sim on the set of vertices X(0), subject to the requirement that $u \sim v$ implies that $u, v \in e(0)$ for some hyperedge in $e \in X(1)$.

We shall make every graph G into a related hypergraph by setting $u \sim v$ if and only if there is an edge $e \in G(1)$ with $e = \{u, v\}$.

Our main example of a related weighted hypergraph is the following:

Definition 7.2 (No-Intersection Related Hypergraph of a Poset). Let (X, w) be a normalized weighted d-poset and let $i, j, k \in \{0, ..., d\}$ with i, j < k. The (i, j, k)-no-intersection hypergraph of (X, w) is a related weighted hypergraph NIH^{*i*,*j*,*k*} $(X, w) = (\text{NIH}^{i,j,k}(X), w_{\text{NIH}^{i,j,k}(X)}, \sim)$ defined as follows:

- the vertices of $\text{NIH}^{i,j,k}(X)$ are $X(i) \cup X(j)$ (if i = j, the vertex set is just X(i));
- the hyperedges of $NIH^{i,j,k}(X)$ are X(k);
- the vertices of a hyperedge $z \in X(k)$ are the $x \in X(i) \cup X(j)$ with $x \leq z$.

We give $\operatorname{NIH}^{i,j,k}(X)$ the weight function $w_{i,j,k} := w_{\operatorname{NIH}^{i,j,k}(X)}$ determined as follows:

- for a hyperedge $z \in X(k)$, set $w_{i,j,k}(x) = w(x)$;
- for a vertex $x \in X(i) \cup X(j)$, set $w_{i,j,k}(x) = w(x)$ if i = j and $w_{i,j,k}(x) = \frac{1}{2}w(x)$ if $i \neq j$.

Finally, we endow the vertices of $\text{NIH}^{i,j,k}(X)$ with a binary relation \sim :

• for $x, y \in \text{NIH}^{i,j,k}(X)(0)$, let $x \sim y$ if (1) there is $z \in X(k)$ with $x, y \leq z$, (2) $\text{Inf}\{x, y\} = \{\emptyset_X\}$ in X and (3) $(\dim x, \dim y) \in \{(i, j), (j, i)\}.$

Since (X, w) is normalized, NIH^{*i*,*j*,*k*}(X, w) is normalized as well. If (X, w) is properly weighted, X is lower-regular and i = j, then NIH^{*i*,*j*,*k*}(X, w) is also properly weighted (Corollary 4.19), but this is false in general.

We will also consider the following simplified version of the no-intersection hypergraph. Recall that graphs are allowed to have multiple edges, but no loops.

Definition 7.3 (No-Intersection Graph of a Poset). Let (X, w) be a normalized weighted d-poset and let $i, j, k \in \{0, ..., d\}$ with $i, j \leq k$. The (i, j, k)-no-intersection graph of X is a weighted graph NIG^{*i*,*j*,*k*}(X, w) = (NIG^{*i*,*j*,*k* $}<math>(X), w_{NIG^{$ *i*,*j*,*k* $}(X))$ defined as follows:}

- the vertices of $\operatorname{NIG}^{i,j,k}(X)$ are $X(i) \cup X(j)$ (if i = j, the vertex set is just X(i));
- the edges of $\operatorname{NIG}^{i,j,k}(X)$ are pairs $(z, \{x, y\})$ such that $z \in X(k), x \in X(i), y \in X(j), x \neq y$, $\operatorname{Inf}\{x, y\} = \{\emptyset_X\}$ and $x, y \leq z$;
- the vertices of the edge $(z, \{x, y\})$ are x and y.

We endow $\operatorname{NIG}^{i,j,k}(X)$ with the weight function $w_{i,j,k} := w_{\operatorname{NIG}^{i,j,k}(X)}$ determined as follows:

- for an edge $(z, \{x, y\}) \in G(1)$, set $w_{i,j,k}(z, \{x, y\}) = w(z)$;
- for a vertex $x \in G(0) = X(i) \cup X(j)$, set $w_{i,j,k}(x) = w(x)$ if i = j and $w_{i,j,k}(x) = \frac{1}{2}w(x)$ if $i \neq j$.

Recall that we also view $\operatorname{NIG}^{i,j,k}(X,w)$ as a related hypergraph by setting $u \sim v$ if there is $e \in \operatorname{NIG}^{i,j,k}(X)(1)$ with $e = \{u, v\}$.

Note that while the total weight of the vertices in $\text{NIG}^{i,j,k}(X)$ is 1, the total weight of the edges may exceed 1. Thus, in contrast to $\text{NIH}^{i,j,k}(X,w)$, the graph $\text{NIG}^{i,j,k}(X,w)$ may not be normalized.

The no-intersection graph and the no-intersection hypergraph are related by surjective map

$$\varphi_X : \mathrm{NIG}^{i,j,k}(X) \to \mathrm{NIH}^{i,j,k}(X)$$

given by the identity on 0-faces and by $(z, \{x, y\}) \mapsto z$ on 1-faces. It preserves the poset relation and the face weights, and also respects the relation \sim on the vertices. **Example 7.4.** Let (X, w) be a normalized weighted *d*-poset. Then $\text{NIH}^{0,0,1}(X, w) = \text{Gr}(X, w)$. Moreover precisely, $\text{NIH}^{0,0,1}(X, w)$ is the underlying weighted hypergraph of (X, w) together with the binary relation \sim on X(0) determined by $u \sim v$ if and only if $u \neq v$ and $X(1)_u \cap X(1)_v \neq \emptyset$.

Suppose further that X is a simplicial complex. Then $\varphi_X : \operatorname{NIG}^{0,0,1}(X) \to \operatorname{NIH}^{0,0,1}(X)$ is an isomorphism of weighted related hypergraphs, meaning in particular that $\operatorname{NIH}^{0,0,1}(X)$ is a graph and $\operatorname{NIG}^{0,0,1}(X) = \operatorname{Gr}(X)$.

Example 7.5. Let X be a pure simplicial complex of dimension d. and let $i, j, k \in \{0, ..., d\}$ with i, j < k. The nature of $H := \operatorname{NIH}^{i,j,k}(X)$ and $G := \operatorname{NIG}^{i,j,k}(X)$ depends on whether i + j = k - 1, i + j > k - 1 or i + j < k - 1.

When i + j = k - 1, the map $\varphi_X : \operatorname{NIG}^{i,j,k}(X) \to \operatorname{NIH}^{i,j,k}(X)$ is bijective on vertices and $\frac{1}{2}\binom{k+1}{i+1}$ -to-1 on (hyper)edges. It further maps \sim_G bijectively onto \sim_H . One may therefore think of H as the related hypergraph obtained from the graph G by gluing, for every $z \in X(k)$, the edges of the form $(z, \{x, y\})$ to each other.

When i + j > k - 1, the graph G has no edges, because an *i*-face and a *j*-face of a given k-face must have nonempty intersection. On the other hand, H has hyperedges, but the relation \sim on its vertices is the empty relation.

Finally, if i + j < k - 1, then neighboring vertices in G will usually be connected by many edges, because an (i + j - 1)-face may be contained in many k-faces. The same phenomenon occurs for H, and the relation \sim on H(0) becomes complicated to describe.

7.2 Skeleton Expansion

Skeleton expansion was considered for properly weighted graphs in [EK16], [KM18] and similar sources. We now generalize this concept to related weighted hypergraphs.

Definition 7.6 (Skeleton Expansion). Let (X, w, \sim) be a related weighted hypergraph and let $\alpha, \beta \in [0, \infty)$. Given $A \subseteq X(0)$, define

 $E_2(A) = \{ e \in X(1) : \text{ there are distinct } u, v \in A \text{ s.t. } u \leq e, v \leq e \text{ and } u \sim v \}.$

We say that (X, w, \sim) is an (α, β) -skeleton expander if for every $A \subseteq X(0)$, we have

$$w(E_2(A)) \le \alpha w(A) + \beta w(A)^2.$$

Note that if X is a graph viewed as a related hypergraph, then $E_2(A)$ is just the set of edges having both their vertices in A, commonly denoted as E(A).

In [EK16], [KM18] and related sources, a properly weighted graph (X, w) was called an α -skeleton expander if $w(E(A)) \leq \alpha w(A) + w(A)^2$ for every $A \subseteq X(0)$. This is equivalent to X being an $(\alpha, 1)$ -skeleton expander in our sense. We introduced the additional constant β to better accommodate improperly (even non-normalized) weight functions and non-connected graphs.

Example 7.7. (i) Let (X, w) be a properly weighted graph, and let λ be the second largest eigenvalue of the normalized adjacency operator of (X, w) (§2.1). By Proposition 2.2, (X, w) is a $(\lambda, 1)$ -skeleton expander.

(ii) Let (X, w, \sim) be a weighted related hypergraph. If X has no hyperedges, or \sim is the empty relation, then (X, w, \sim) is a (0, 0)-skeleton expander. Thus, if X is a weighted pure simplicial complex of dimension d and $i, j, k \in \{0, \ldots, d\}$ satisfy $0 \le i, j < k$ and i + j > k - 1, then both $\operatorname{NIH}^{i,j,k}(X)$ and $\operatorname{NIG}^{i,j,k}(X)$ are (0, 0)-skeleton expanders (cf. Example 7.5).

(iii) Let G be a pure graph and let G_n (n > 1) be a graph consisting of n disjoint copies of G. Give G and G_n their natural weight functions. Since G_n is not connected, the second largest

eigenvalue of its weighted adjacency operator is 1, and so Proposition 2.2 only tells us that G_n is a (1, 1)-skeleton expander. However, if G is an (α, β) -skeleton expander, then G_n is an $(\alpha, n\beta)$ -skeleton expander. Indeed, given $A \subseteq G_n(0)$, let A_i be the intersection of A with the *i*-th copy of G in G_n . Then

$$w_{G_n}(E(A)) = \sum_{i=1}^n w_{G_n}(E(A_i)) = \frac{1}{n} \sum_{i=1}^n w_G(E(A_i)) \le \frac{1}{n} \sum_{i=1}^n [\alpha w_G(A_i) + \beta w_G(A_i)^2]$$
$$\le \alpha w_{G_n}(A) + \frac{\beta}{n} (\sum_{i=1}^n w_G(A_i))^2 = \alpha w_{G_n}(A) + n\beta w_{G_n}(A)^2.$$

Let (X, w) be a normalized weighted *d*-poset. We will ultimately be interested in the skeleton expansion of related weighted hypergraphs of the form $\operatorname{NIH}^{i,j,k}(X, w)$, but very little is known about it. The following lemma tells us that it is enough to bound the skeleton expansion of the graph $\operatorname{NIG}^{i,j,k}(X)$, which is much more manageable. In fact, this is the reason why we need to consider $\operatorname{NIG}^{i,j,k}(X)$ in this work. We expect that in general the skeleton expansion of $\operatorname{NIH}^{i,j,k}(X)$ will not be much smaller (in both parameters) than that of $\operatorname{NIG}^{i,j,k}(X)$.

Lemma 7.8. Let (X, w) be a normalized weighted d-poset and let $i, j, k \in \{0, ..., d\}$ with $i, j \leq k$. If $\operatorname{NIG}^{i,j,k}(X, w)$ is an (α, β) -skeleton expander, then $\operatorname{NIH}^{i,j,k}(X, w)$ is an (α, β) -skeleton expander.

Proof. Let A be a set of vertices in $H := \operatorname{NIH}^{i,j,k}(X)$. We may also view A as a set of vertices in $G := \operatorname{NIG}^{i,j,k}(X)$. One readily checks that $\varphi_X(E_G(A)) = E_2(A)$, and hence $w_H(E_2(A)) \leq w_G(E_G(A)) \leq \alpha w(A) + \beta w(A)^2$.

7.3 Intersection Profiles

Notation 7.9. Let (X, w) be a properly weighted *d*-poset, let $z \in X$ and let $i, j, k \in \{0, \ldots, d\}$ with $b := \dim z < i, j < k$. We let

$$\operatorname{NIH}_{z}^{i,j,k}(X,w) = \operatorname{NIH}^{i-b+1,j-b+1,k-b+1}(X_{z},w_{z}).$$

Here, w_z is the proper weight function induced by w on X_z (§4.3).

In our main result, we would need to consider the skeleton expansion of related weighted hypergraphs of the form $\operatorname{NIH}_{z}^{i,j,k}(X,w)$ for various values of i, j, k and $b = \dim z$. For a particular X, it will often be enough to consider a subset of the legal quadruples (k, i, j, b). In order to keep track of the quadruples which are needed, we introduce the notion of an *intersection profile*.

Let $k \ge 0$ be an integer. Broadly speaking, a k-intersection profile for a d-poset X is encodes the dimensions of some quadruples of faces (z, x, y, u) such that x, y lie between z and u and $u \in \text{Inf}\{x, y\}$. The axioms guarantee that all the quadruples (z, x, y, u) with dim z = k + 1 and dim $x = \dim y = k$ must be included, and also that quadruples obtained by taking infima of faces that were previously encountered are also included. The formal definition is as follows.

Definition 7.10 (Abstract Intersection Profile). Let $k \in \mathbb{N} \cup \{0\}$. An abstract k-intersection profile \mathcal{P} consists of a set of integer quadruples $(t, \ell, r, b)^{17}$ with $k + 1 \ge t > \ell \ge r > b \ge -1$, also denoted \mathcal{P} , and a set of pairs of integers (i, j) with $k + 1 \ge i > j \ge -1$, denoted $\operatorname{Ad}(\mathcal{P})$, such that the following conditions are met:

(1) for every $(t, \ell, r, b) \in \mathcal{P}$, we have $(t, \ell), (t, r), (\ell, b), (r, b) \in Ad(\mathcal{P})$;

¹⁷The letters t, ℓ, r, b allude to the words top, left, right and bottom.

- (2) if $t > \ell \ge r > -1$ are integers such that $(t, r), (t, \ell) \in Ad(\mathcal{P})$, then \mathcal{P} contains a quadruple of the form $(t, \ell, r, *)$;
- (3) if $t > \ell > r \ge -1$ are integers such that $(t, \ell), (t, r) \in \mathrm{Ad}(\mathcal{P})$, then $(\ell, r) \in \mathrm{Ad}(\mathcal{P})$;¹⁸

$$(4) (k+1,k) \in \mathrm{Ad}(\mathcal{P}).$$

We call $Ad(\mathcal{P})$ the set of \mathcal{P} -admissible pairs.

The set of admissible pairs of an intersection profile \mathcal{P} is often uniquely determined from its underlying set of integer quadruples. In such cases, we shall sometimes not specify $\operatorname{Ad}(\mathcal{P})$ explicitly.

Definition 7.11 (Intersection Profile). Let X be a d-poset, let $k \in \mathbb{N} \cup \{0\}$ and let \mathcal{P} be an abstract k-intersection profile. Two faces $x, y \in X$ are said to be \mathcal{P} -admissible if x > y and $(\dim x, \dim y) \in Ad(\mathcal{P})$. We say that \mathcal{P} is a k-intersection profile if for every $x, y, z \in X$ such that (x, y) and (x, z) are \mathcal{P} -admissible with $\dim y \geq \dim z$, either $x > y \geq z$, or $(\dim x, \dim y, \dim z, \dim u) \in \mathcal{P}$ for every $u \in Inf\{y, z\}$.

Example 7.12. Let $d, k \ge 0$ be integers. Every *d*-poset X has a unique minimal *k*-intersection profile \mathcal{P} which may be constructed iteratively as follows:

- Start with $\mathcal{P} = \emptyset$ and $\operatorname{Ad}(\mathcal{P}) = \{(k+1,k)\}.$
- For every integer i running down from k + 1 to 0, perform:
 - For every $x \in X(i)$ and every y, z < x with $(\dim x, \dim y), (\dim x, \dim z) \in Ad(\mathcal{P})$ and $\dim y \geq \dim z$, perform:
 - * If x > y > z, add $(\dim y, \dim z)$ to $\operatorname{Ad}(\mathcal{P})$.
 - * Otherwise, for every $u \in \text{Inf}\{y, z\}$, add $(\dim y, \dim u)$, $(\dim z, \dim u)$ to $\text{Ad}(\mathcal{P})$ and $(\dim x, \dim y, \dim z, \dim u)$ to \mathcal{P} .

At the end of this process, \mathcal{P} and $\operatorname{Ad}(\mathcal{P})$ will determine a k-intersection profile for X. (Then name intersection profile comes from the repeated use of intersections (Inf) in this construction.)

Example 7.13 (Intersection Profile for Simplicial Complexes). The abstract k-intersection profile

$$\mathcal{P}_{\triangle}^{(k)} := \{ (i+1, i, i, i-1) \, | \, i \in \{0, \dots, k\} \}$$

is a k-intersection profile for any pure simplicial complex X; the $\mathcal{P}_{\Delta}^{(k)}$ -admissible pairs are $\operatorname{Ad}(\mathcal{P}_{\Delta}^{(k)}) = \{(k, k-1), (k-1, k-2), \ldots, (0, -1)\}$. Indeed, that $\mathcal{P}_{\Delta}^{(k)}$ is an abstract k-intersection follows directly from the definition. Furthermore, if $x, y, z \in X$ are distinct faces such that (x, y) and (x, z) are $\mathcal{P}_{\Delta}^{(k)}$ -admissible, then $-1 < \dim y = \dim z = \dim x - 1$. Since X is simplicial, $\inf\{y, z\} = y \cap z$ and indeed $(\dim x, \dim y, \dim z, \dim(y \cap z)) = (\dim x, \dim x - 1, \dim x - 1, \dim x - 2) \in \mathcal{P}_{\Delta}^{(k)}$.

Given $(t, \ell, r, b) \in \mathcal{P}^{(k)}_{\Delta}$ and $z \in X(b)$, we must have t = b+2 and $\ell = r = b+1$, so $\operatorname{NIH}_{z}^{t,\ell,r}(X) = \operatorname{NIH}^{0,0,1}(X_{z})$ is just the underlying graph of X_{z} .

¹⁸The rationale behind this requirement is that we would have liked the illegal quadruple (t, ℓ, r, r) to be \mathcal{P} . Since we are not allowed to include it, we compensate for that by requiring that (ℓ, r) is in Ad (\mathcal{P}) .

Example 7.14 (Intersection Profile for Cube Complexes). The abstract k-intersection profile

$$\mathcal{P}_{\Box}^{(k)} := \mathcal{P}_{\bigtriangleup}^{(k)} \cup \{(i+1, i, i, -1) \mid i \in \{1, \dots, k\}\}$$

is a k-intersection profile for any pure cube complex X; its set of admissible pairs is $\{(i, i-1) | i \in \{k+1, \ldots, 0\}\} \cup \{(i, -1) | i \in \{k+1, \ldots, 1\}\}$. To see this, let $x, y, z \in X$ be faces as in Definition 7.11 and write $i = \dim x$. Then we must have $-1 < \dim y = \dim z = \dim x - 1$. Since X is a cube complex, the faces y and z must be (i - 1)-faces of the *i*-dimensional cube x, so if they are distinct, their intersection is either an (i - 2)-dimensional cube, or the empty face. In both cases, $(\dim x, \dim y, \dim z, \dim(y \cap z)) \in \mathcal{P}_{\Box}^{(k)}$.

 $(\dim x, \dim y, \dim z, \dim(y \cap z)) \in \mathcal{P}_{\square}^{(k)}$. As in the last example, given $z \in X$ of the dimension i, the hypergraph $\operatorname{NIH}_{z}^{i+1,i+1,i+2}(X)$ (corresponding to $(i+2, i+1, i+1, i) \in \mathcal{P}_{\square}^{(k)}$) is just $\operatorname{Gr}(X_z)$. On the other hand, $\operatorname{NIG}_{\emptyset}^{i,i+1}(X)$ (corresponding to $(i+1, i, i, -1) \in \mathcal{P}_{\square}^{(k)}$) is the graph obtained by taking the *i*-cubes in X as vertices and adding one edge between a pair of *i*-cubes x, y for every (i+1)-cube having x and y as opposite sides.

Example 7.15. (i) $\mathcal{P}^{(0)} := \{(1,0,0,-1)\}$ is a 0-intersection profile for every poset X. In fact, $\mathcal{P}^{(0)}$ is the only abstract 0-intersection profile.

(ii) $\mathcal{P}^{(1)} := \{(2,1,1,0), (2,1,1,-1), (1,0,0,-1)\}$ is a 1-intersection profile for every poset X. It coincides with $\mathcal{P}_{\Box}^{(1)}$.

(iii) Let $\mathcal{P}_{\max}^{(k)}$ denote the set of all quadruples of integers (t, ℓ, r, b) with $k+1 \ge t > \ell \ge r > b \ge -1$. Then $\mathcal{P}_{\max}^{(k)}$ is a k-intersection profile for any d-poset. It is also the largest possible abstract k-intersection profile. In fact, the slightly smaller abstract k-intersection profile

$$\mathcal{P}_{\text{univ}}^{(k)} := \mathcal{P}_{\text{max}}^{(k-1)} \cup \{(k+1,k,k,i) \, | \, i \in \{-1,\dots,k-1\}\}$$

(with the convention $\mathcal{P}_{\max}^{(-1)} = \emptyset$) is also a k-intersection profile for every d-poset X, and is the smallest one having this property.

8 Main Result: Simple Versions

Our main result is a criterion for establishing cosystolic expansion of a sheaf on a poset. The most general form of this criterion — Theorem 11.2 — is technical, and so we find it instructive to first give in this section several special cases which are simpler and easier to apply; they will be proved in Section 11. In fact, these special cases suffice for the applications considered in this paper, and likely for other potential applications.

Theorem 8.1. Let $B \in \mathbb{R}_+$, $F \in \mathbb{N}$, $L \in [1, \infty)$ and $k \in \{0\} \cup \mathbb{N}$. Then there exist (small) constants $K, K' \in (0, 1]$ such that the following hold: Let R be a commutative ring, let $d \ge k+2$, let (X, w) be a properly weighted R-oriented d-poset of lower irregularity at most L and such that $F_{i,j}^{\max}(X) \le F$ for all $-1 \le i \le j \le k+2$, let \mathcal{P} be a k-intersection profile for X and let \mathcal{P}' be a (k+1)-intersection profile for X. Let \mathcal{F} be an R-sheaf on X, let $\varepsilon \in (0, 1]$ and suppose that:

- (1a) $\operatorname{cbe}_{k-\dim u-1}(X_u, w_u, \mathcal{F}_u) \ge \varepsilon$ for every $u \in X(0) \cup \cdots \cup X(k)$;
- (1b) $\operatorname{cbe}_{k-\dim u}(X_u, w_u, \mathcal{F}_u) \ge \varepsilon$ for every $u \in X(0) \cup \cdots \cup X(k+1)$;
- (2) for every $(t, \ell, r, b) \in \mathcal{P} \cup \mathcal{P}'$ and $u \in X(b)$, the related weighted hypergraph $\operatorname{NIH}_{u}^{\ell, r, t}(X)$ (Notation 7.9) is a $((K\varepsilon)^{2^{k+1-\min\{\ell, r\}}}, B(K\varepsilon)^{2^{k+2-\min\{\ell, r\}}-2^{k+1-b}})$ -skeleton expander.

Then

$$\operatorname{cse}_k(X, w, \mathcal{F}) \ge K'(K\varepsilon)^{2^{k+2}-1}$$
 and $\operatorname{ccd}_k(X, w, \mathcal{F}) \ge K'(K\varepsilon)^{2^{k+1}-1}$.

Moreover, if $f \in C^k = C^k(X, \mathcal{F})$ satisfies $\operatorname{dist}(f, Z^k) < K'(K\varepsilon)^{2^{k+2}-1}$, then applying Algorithm 8.2 below to f with $q = (K\varepsilon)^{2^{k+1}}$ returns $f' \in Z^k$ such that $\operatorname{dist}(f, f') \leq K'^{-1}(K\varepsilon)^{-2^{k+1}} \operatorname{dist}(f, Z^k)$.

Algorithm 8.2 (Correction Algorithm). Let $(X, w), \mathcal{F}, d, k$ be as in Theorem 8.1. The input to the algorithm is some $f \in C^k(X, \mathcal{F})$ and a real parameter $q \geq 0$. The algorithm outputs another k-cocycle $f' \in C^k(X, \mathcal{F})$, computed as follows:

- (1) Set $f_0 = f$ and i = 0.
- (2) Do:
 - (a) Look for $u \in X(0) \cup \cdots \cup X(k)$ and $g \in C^{k-\dim u-1}(X_u, \mathcal{F}_u)$ such that $||df_i + d(g^u)|| < ||df_i|| q \cdot w(u)$.¹⁹
 - (b) If no such u and g exist, return f_i . Otherwise, set $f_{i+1} = f_i + g^u$ and increase i by 1.

A few remarks are in order.

Remark 8.3. Concerning Theorem 8.1:

(i) The smaller the reduced parts of intersection profiles \mathcal{P} and \mathcal{P}' are, the more lax condition (2) is. Thus, one should choose the intersection profiles to be as small as possible. Using smaller intersection profiles also improves the constants K and K' by making them larger.

(ii) In applications of Theorem 8.1, F, L, k and ε typically remain constant as X and \mathcal{F} vary, whereas the skeleton expansion of $\operatorname{NIH}_{u}^{\ell,r,t}(X)$ tends to (0,c) with some $c \in \mathbb{R}_{+}$ as the degree of X grows. Thus, once the degree of X is large enough (but constant) and B is chosen to be large enough in advance, conditions (1a), (1b) and (2) will hold and the theorem could be applied.

(iii) Assumptions (1a) and (1b) of Theorem 8.1 are *local* in the sense that they care only about the structure of X_z and \mathcal{F}_z for $\emptyset \neq z \in X$ and not about the global structure of X and \mathcal{F} . The reason why condition (2) is not local in this sense is that it includes the requirement that $\operatorname{NIH}^{\ell,r,t}(X)$ is an $((K\varepsilon)^{2^{k+1}}, B(K\varepsilon)^{2^{k+2-\min\{\ell,r\}}-2^{k+2}})$ -skeleton expander whenever $(-1, \ell, r, t) \in \mathcal{P} \cup \mathcal{P}'$. As we shall see shortly, this non-local requirement can be replaced with a local condition when X is a simplicial complex.

(iv) In the special case where X is a simplicial complex and \mathcal{F} is a constant sheaf, criteria for cosystolic expansion appeared in [KKL16] (dim $X \leq 3$, $\mathcal{F} = (\mathbb{F}_2)_X$), [EK17] ($\mathcal{F} = (\mathbb{F}_2)_X$), [KM21], [DD23]. In addition, the result [KO21, Thm. 1.17] can be interpreted in our language as a criterion for 0-cosystolic expansion of 2-posets equipped with a special kind of a *constraint system* (see §5.5). We surveyed these works in detail and compared them to Theorem 8.1 in §1.6.

Remark 8.4. Concerning Algorithm 8.2, if one fixes some $M \ge 1$ and requires that $w(x) \le Mw(y)$ and $|\mathcal{F}(x)| \le M$ for all $x, y \in X(0) \cup \cdots \cup X(k)$, and $D_{i,k}(X) \le M$ for all $i \in \{0, \ldots, k\}$, then the time complexity of Algorithm 8.2 is *linear* in |X(k)| (the constant depends on F, M); see Proposition 12.5 below. To do the search in (a) in constant time on average, one has to keep a set of the possible faces u which may satisfy $||df_i + d(g^u)|| < ||df_i|| - q \cdot w(u)$ and update it during the search and every time (b) is performed. For details, see Algorithm A.1 in the appendix.

As an immediate corollary of Theorem 8.1 and Lemma 6.2, we get the following criterion for showing that a cocycle code is locally testable and has linear distance.

¹⁹For the definition of g^u , see §5.3.

Corollary 8.5. With notation as in Theorem 8.1, suppose further that:

- $B^k = B^k(X, \mathcal{F}) = 0,$
- there is an R-module Σ such that $\mathcal{F}(x) = \Sigma$ for every $x \in X(k)$ and
- there are $M, M' \in \mathbb{R}_+$ such that $w(x) \leq Mw(y)$ for all $x, y \in X(k)$ and $w(x') \leq M'w(y')$ for all $x', y' \in X(k+1)$.

Then the cycle code

$$Z^k = Z^k(X, \mathcal{F}) \subseteq \Sigma^{X(k)}$$

satisfies $\delta(Z^k) \geq \frac{1}{M}K'(K\varepsilon)^{2^{k+1}-1}$ and its natural tester has soundness $\frac{1}{MM'}K'(K\varepsilon)^{2^{k+2}-1}$. Moreover, Algorithm 8.2 with $q = (K\varepsilon)^{2^{k+1}}$ is a decoding algorithm for Z^k that works for words that are $\frac{K'^2(K\varepsilon)^{2^{k+2}-1}}{M+MK'(K\varepsilon)^{2^{k+1}}}$ -close to Z^k .

By Proposition 4.24, when w is the natural weight function of X, we can take $M = L_{k,d}(X)U_{k,d}(X)$ and $M' = L_{k+1,d}(X)U_{k+1,d}(X)$.

Proof. The only thing that needs proof is the claim about the decoding, and this follows from the next lemma (after recalling that $K \leq 1$).

Lemma 8.6. Let X be an R-oriented d-poset, let \mathcal{F} be an R-sheaf on X, let $k \in \mathbb{Z}$ and let $M \in [1, \infty)$. Suppose that $X(k) \neq \emptyset$, $B^k = B^k(X, \mathcal{F}) = 0$, and for every $x, y \in X(k)$, we have $\mathcal{F}(x) = \Sigma$ and $w(x) \leq Mw(y)$. Assume further that there are A, B > 0 and an algorithm which takes $f \in C^k = C^k(X, \mathcal{F})$ with dist $(f, Z^k) < A$ and returns $f' \in Z^k$ such that dist $(f, f') \leq B$ dist (f, Z^k) . Then this algorithm is also a decoding algorithm for the k-cocycle code $Z^k \subseteq \Sigma^{X(0)}$ which can decode words that are $\frac{1}{M} \min\{\frac{\operatorname{ccd}_k(X, \mathcal{F})}{B+1}, A\}$ -close to Z^k .

Proof. Write $\eta = \frac{1}{M} \min\{\frac{\operatorname{ccd}_k(X,\mathcal{F})}{B+1}, A\}$, and let $f \in C^k$ be η -close to Z_k . Then there is f_0 such that $\operatorname{dist}_{\operatorname{uni}}(f, f_0) < \eta$. As in the proof of Lemma 6.2, we have $\operatorname{dist}(f, f_0) \leq M \operatorname{dist}_{\operatorname{uni}}(f, f_0) < M\eta \leq A$. Thus, applying the algorithm to f yields $f' \in Z^k$ with $\operatorname{dist}(f, f') \leq B \operatorname{dist}(f, Z^k) \leq B \operatorname{dist}(f, f_0) = BM\eta$. Thus, $\operatorname{dist}(f_0, f') \leq \operatorname{dist}(f_0, f) + \operatorname{dist}(f, f') < M\eta + BM\eta \leq \operatorname{ccd}_k(X, \mathcal{F})$. Since $B^k = 0$, this means that $f' = f_0$, so we have shown that the algorithm decoded f.

We proceed with specializing Theorem 8.1 to simplicial complexes and cube complexes. In the former case, it simplifies into the following theorem.

Theorem 8.7 (Cosystolic Expansion for Simplicial Complexes). Let $k \in \mathbb{N}$. Then there exist (small) constants $K, K' \in (0, 1]$ such that the following hold: Let R be a commutative ring, let (X, w) be a properly weighted pure d-dimensional simplicial complex with $d \ge k + 2$, let \mathcal{F} be an R-sheaf on X and let $\varepsilon \in (0, 1]$. Suppose that:

- (1a) $\operatorname{cbe}_{k-\dim u-1}(\mathcal{F}_u) \ge \varepsilon$ for every $u \in X(0) \cup \cdots \cup X(k)$;
- (1b) $\operatorname{cbe}_{k-\dim u}(\mathcal{F}_u) \ge \varepsilon$ for every $u \in X(0) \cup \cdots \cup X(k+1)$;
- (2) for every $u \in X(-1) \cup \cdots \cup X(k)$, the underlying graph of (X_u, w_u) is an $((K\varepsilon)^{2^{k-\dim u}}, 1)$ -skeleton expander.

Then

$$\operatorname{cse}_k(\mathcal{F}) \ge K'(K\varepsilon)^{2^{k+2}-1}$$
 and $\operatorname{ccd}_k(\mathcal{F}) \ge K'(K\varepsilon)^{2^{k+1}-1}$

Moreover, if $f \in C^k = C^k(X, \mathcal{F})$ satisfies $\operatorname{dist}(f, Z^k) < K'(K\varepsilon)^{2^{k+2}-1}$, then applying Algorithm 8.2 to f with the parameter $q = (K\varepsilon)^{2^{k+1}}$ returns $f' \in Z'$ such that $\operatorname{dist}(f, f') \leq K'^{-1}(K\varepsilon)^{-2^{k+1}} \operatorname{dist}(f, Z^k)$.

Proof. This follows from Theorem 8.1 by taking B = 1, L = 1, $F = \binom{k+2}{\lceil (k+2)/2 \rceil}$, $\mathcal{P} = \mathcal{P}^{(k)}_{\bigtriangleup}$ and $\mathcal{P}' = \mathcal{P}^{(k+1)}_{\bigtriangleup}$ (notation as in Example 7.13). Indeed, X is always R-oriented and lower regular (i.e. L(X) = 1), and $F_{i,j}(X) \leq \binom{k+2}{\lceil (k+2)/2 \rceil} =: F$ for all $-1 \leq i \leq j \leq k+2$.

Remark 8.8. The only non-local assumption in Theorem 8.7 is that $X(\leq 1)$ is a $((K\varepsilon)^{2^{k+1}}, 1)$ skeleton expander (take $u = \emptyset_X$ in (2)). This assumption can be replaced by a stronger local
condition thanks to Oppenheim's Trickling Down Theorem [Opp15, Theorem 1.4]. Specifically, we
can fix some $0 \leq \ell \leq d-2$ and (2) with requiring that the underlying graph of (X_z, w_z) is a λ expander (see §2.1) every $z \in X(\ell)$, where λ is small enough so that $\frac{\lambda}{1-(\ell+1)\lambda} \leq (K\varepsilon)^{2^{k+1}}$. (Indeed,
by Oppenheim's Theorem, this would guarantee that $X(\leq 1)$ is a $(K\varepsilon)^{2^{k+1}}$ -expander and hence a $((K\varepsilon)^{2^{k+1}}, 1)$ -skeleton expander.)

For cube complexes, Theorem 8.1 specializes to the following theorem.

Theorem 8.9 (Cosystolic Expansion for Cube Complexes). Let $k \in \mathbb{N}$ and $B \in \mathbb{R}_+$. Then there exist (small) constants $K, K' \in (0, 1]$ such that the following hold: Let R be a commutative ring, let (X, w) be a properly weighted pure d-dimensional cube complex with $d \ge k+2$, let \mathcal{F} be an R-sheaf on X and let $\varepsilon \in (0, 1]$. Suppose that conditions (1a)-(2) of Theorem 8.7 hold and in addition,

(2') for every $i \in \{1, ..., k+1\}$, the weighted related hypergraph $\text{NIH}^{i,i,i+1}(X)$ is an $((K\varepsilon)^{2^k}, B(K\varepsilon)^{-2^{k+1}})$ -skeleton expander.

Then the conclusions of Theorem 8.7 hold verbatim.

By Lemma 7.8, the theorem also holds if we replace the hypergraph $\text{NIH}^{i,i,i+1}(X)$ with the graph $\text{NIG}^{i,i,i+1}(X)$.

Proof. Let $F_{i,j}$ be the number of *i*-faces a *j*-dimensional cube has. The theorem follows by applying Theorem 8.1 with $F = \max\{F_{0,k+2}, \ldots, F_{k,k+2}\}, L = 1, \mathcal{P} = \mathcal{P}_{\Box}^{(k)}$ and $\mathcal{P}' = \mathcal{P}_{\Box}^{(k)}$ (notation as in Example 7.14).

We now restrict our attention to 0-cocycle codes. In this special case, the general form of our criterion for cosystolic expansion (Theorem 11.2) simplifies into the following result.

Theorem 8.10 (Criterion for 0-Cosystolic Expansion). For every $F \in \mathbb{N}$ and $L \in [1, \infty)$, there are (small) real constants E, E', E'', D, D', D'' > 0 such that the following hold: Let R be a ring, let (X, w) be a properly weighted R-oriented d-poset $(d \geq 2)$ such that $L(X) \leq L$ and $F_{0,2}^{\max}(X), F_{1,2}^{\max}(X) \leq F$, and let \mathcal{F} be an R-sheaf on X. Let $\varepsilon, \varepsilon', \alpha_0, \beta_0, \alpha_{-1}, \beta_{-1}, \alpha_{||}, \beta_{||} > 0$ and suppose that:

(1a) $\operatorname{cbe}_{-1}(\mathcal{F}_v) \geq \varepsilon$ for every $v \in X(0)$;

- (1b) $\operatorname{cbe}_{-1}(\mathcal{F}_e) \geq \varepsilon'$ for every $e \in X(1)$;
- (1c) $\operatorname{cbe}_0(\mathcal{F}_v) \geq \varepsilon'$ for every $v \in X(0)$;

- (2a) $X_v(\leq 1)$ is an (α_0, β_0) -skeleton expander for all $v \in X(0)$;
- (2b) $X(\leq 1)$ is an $(\alpha_{-1}, \beta_{-1})$ -skeleton expander.
- (2c) NIH^{1,1,2}(X) (see §7) is an $(\alpha_{\parallel}, \beta_{\parallel})$ -skeleton expander.

Suppose further that

$$\alpha_{-1} < E\varepsilon$$

and one can find $h_{-1}, h_0, h_{||} \in (0, 1]$ satisfying the following inequality:

$$(\alpha_{0} + \beta_{0}h_{0}) + (\alpha_{||} + \beta_{||}h_{||}) + \frac{\alpha_{-1} + \beta_{-1}h_{-1}}{h_{0}} \le E'\varepsilon'.$$

Then

$$\csc_0(X, w, \mathcal{F}) \ge \frac{E''}{h_0^{-1}h_1^{-1} + h_{||}^{-1}} \quad and \quad \operatorname{ccd}_0(X, w, \mathcal{F}) \ge \frac{E'''(E\varepsilon - \alpha_{-1})}{\beta_{-1}}.$$

Moreover, if $f \in C^0 = C^0(X, \mathcal{F})$ satisfies $\operatorname{dist}(f, Z^0) < \frac{D}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}$, then applying Algorithm 8.2 to f with the parameter $a = D'h_0$ returns $f' \in Z'$ such that $\operatorname{dist}(f, f') < \frac{1}{d_0} \operatorname{dist}(f, Z^0)$

f with the parameter $q = D'h_0$ returns $f' \in Z'$ such that $\operatorname{dist}(f, f') \leq \frac{1}{D''h_0} \operatorname{dist}(f, Z^0)$. Explicit values of E, E', E'', E''', D, D', D'' to which this applies are listed in Table 1, both in general, and under some assumptions X.

Assumption on X	E	E'	E''	$E^{\prime\prime\prime}$	D	D'	D''
None	L^{-8}	$\frac{1}{4}L^{-18}F^{-5}$	$\frac{1}{2}L^{-13}F^{-2}$	L^{-3}	$L^{-15}F^{-1}$	$\frac{1}{4}L^{-12}F^{-2}$	$\frac{1}{4}L^{-16}F^{-3}$
Lower regular	1	$\frac{1}{4}F^{-2}$	$\frac{1}{2}F^{-1}$	1	F^{-1}	$\frac{1}{4}F^{-1}$	$\frac{1}{4}F^{-2}$
Cube complex	1	$\frac{1}{12}$	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$
Square complex	1	$\frac{1}{16}$	$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$
$X(\leq 2)$ is an <i>m</i> -gon	L^{-8}	$\frac{1}{4m}L^{-18}$	$\frac{1}{2}L^{-13}$	L^{-3}	$\frac{1}{2}L^{-15}$	$\frac{1}{4}L^{-12}$	$\frac{1}{8}L^{-16}$
complex							

Table 1: Values for the constants of Theorem 8.10.

Remark 8.11. Applying Theorem 8.1 with k = 0 and the intersection profiles $\mathcal{P}^{(0)}$ and $\mathcal{P}^{(1)}$ of Example 7.15 gives a similar, but less flexible result. Indeed, Theorem 8.1 can be applied even when the parameters β_{-1} and $\beta_{||}$ are arbitrarily large. This extra generality will be needed in some of the applications.

Corollary 8.12. With notation as in Theorem 8.10, suppose further that $\mathcal{F}(\emptyset_X) = 0$, that there is an *R*-module Σ such that $\mathcal{F}(v) = \Sigma$ for every $v \in X(0)$ and that there are $M, M' \in [1, \infty)$ such that $w(v) \leq Mw(v')$ for all $v, v' \in X(0)$ and $w(e) \leq M'w(e')$ for all $e, e' \in X(1)$. Then the 0-cocycle code $Z^0 = Z^0(X, \mathcal{F}) \subseteq \Sigma^{X(0)}$ satisfies $\delta(Z^0) \geq \frac{E\varepsilon - \alpha_{-1}}{M\beta_{-1}}$ and its natural tester has soundness $\frac{E''}{MM'(h_0^{-1}h_1^{-1} + h_{||}^{-1})}$. Moreover, Algorithm 8.2 with $q = D'h_0$ is a decoding algorithm for words that are $\min\{\frac{D}{M(h_0^{-1}h_{-1}^{-1} + h_{||}^{-1})}, \frac{D''(E\varepsilon - \alpha_{-1})}{M\beta_{-1}(h_0^{-1} + D'')})\}$ -close to Z^0 .

Again, when w is the natural weight function of X, we can take $M = L_{0,d}U_{0,d}$ and $M' = L_{1,d}U_{1,d}$, by Proposition 4.24. *Proof.* This follows from Theorem 8.10, Lemma 6.2 and Lemma 8.6.

In addition to the above results, our method also gives a criterion for guaranteeing that the k-cocycle distance is large. In its most general form, this is Corollary 12.9 appearing later in the text. For 0-cocycles it simplifies into the following theorem, which we prove in §12.2.

Theorem 8.13. For every $F \in \mathbb{N}$ and $L \in [1, \infty)$, there are real constants E, E''' > 0 (the same constants as in Table 1) such that the following hold: Let R be a ring, let (X, w) be a properly weighted d-poset $(d \ge 1)$ such that $L(X) \le L$ and every 1-face of X contains at most F 0-faces. Let $\varepsilon, \alpha, \beta > 0$ such that

$$\alpha < E\varepsilon$$

and suppose that

(1) $\operatorname{cbe}_{-1}(\mathcal{F}_v) \geq \varepsilon$ for every $v \in X(0)$ and

(2) the underlying graph of (X, w) is an (α, β) -skeleton expander.

Then $\operatorname{ccd}_0(X, w, \mathcal{F}) \geq \frac{E^{\prime\prime\prime}(E\varepsilon - \alpha)}{\beta}.$

9 2-Query LTCs from Sheaves on Square Complexes

In this section we apply Theorem 8.10 to certain sheaves on double Cayley complexes in order to construct good 2-query LTCs. Our LTCs turn out to be the line codes of the good LTCs constructed in $[DEL^+22]$, so we can recover the properties of the latter (with slightly different constants) using the relations between a lifted code and its line code established in Section 3. This offers a new perspective on the good LTCs of $[DEL^+22]$, showing how they can be neatly derived from our criterion for 0-cocystolic expansion.

9.1 The Poset

The poset which we will use is a double Cayley complex — a special kind of square complex constructed as follows: Let G be a finite group and let A and B be two symmetric sets of generates to G such that

$$gag^{-1} \neq b \qquad \forall a \in A, b \in B, g \in G.$$

$$(9.1)$$

The double Cayley complex of G, A, B is the square complex X = Cay(A, G, B) constructed as follows:

- $X(0) = \{\{g\} \mid g \in G\}$ (so G is the set of vertices of X),
- $X(1) = \{\{g, ag\} \mid g \in G, a \in A\} \cup \{\{g, bg\} \mid g \in G, b \in B\},\$
- $X(2) = \{\{g, ag, gb, agb\} \mid g \in G, a \in A, b \in B\}.$

We also set $X(-1) = \{\emptyset\}$ and endow X with the inclusion relation. Condition (9.1) guarantees that this is indeed a square complex. Moreover, it implies that if $e \in X(1)$ and $\{g\}$ $(g \in G)$ is a vertex of e, then either $e = \{g, ag\}$ with unique $a \in A$, or $e = \{g, gb\}$ with unique $b \in B$. Furthermore, if $e = \{g, ag\}$ $(g \in G, a \in A)$ is an edge contained in a square s, then there is a unique $b \in B$ such that $s = \{g, ag, gb, agb\}$. Likewise, if $e = \{g, gb\}$, then there is a unique $a \in A$ such that $s = \{g, ag, gb, agb\}$.

Since X is a square complex, and in particular a regular cell complex, there is a \mathbb{Z} -orientation [:] on X such that $[v: \emptyset] = 1$ for every $v \in X(0)$; see Example 4.26. We fix such a \mathbb{Z} -orientation on X; it induces an R-orientation on X for every commutative ring R.

9.2 The Sheaf

Let $G, A, B, X = \operatorname{Cay}(A, G, B)$ be as above. Fix a finite field \mathbb{F} , and let $C_A \subseteq \mathbb{F}^A$ and let $C_B \subseteq \mathbb{F}^B$ be *linear* codes with alphabet \mathbb{F} .

It will be convenient to view $\mathbb{F}^A \otimes \mathbb{F}^B$ (all tensor products are taken over \mathbb{F}) as the space of matrices with rows indexed by A and columns index by B, denoted $\mathcal{M}_{A \times B}(\mathbb{F})$. (Explicitly, for $u \in \mathbb{F}^A$ and $v \in \mathbb{F}^B$, the tensor $u \otimes v$ corresponds to $(u_a v_b)_{a \in A, b \in B} \in \mathcal{M}_{A \times B}(\mathbb{F})$.) With this interpretation, the subspace $C_A \otimes C_B$ is

$$\{M \in \mathcal{M}_{A \times B}(\mathbb{F}) : r_a(M) \in C_B \text{ and } c_b \in C_A \text{ for all } a \in A, b \in B\},\$$

where, as before, $r_a(-)$ and $c_b(-)$ mean taking the *a*-th row and *b*-th column, respectively.

We define an \mathbb{F} -sheaf \mathcal{F} on X as follows. For every $g \in G$, $a \in A$, $b \in B$, set

- $\mathcal{F}(\emptyset) = 0$,
- $\mathcal{F}(\{g\}) = C_A \otimes_{\mathbb{F}} C_B,$
- $\mathcal{F}(\{g,ag\}) = C_B$,
- $\mathcal{F}(\{g, gb\}) = C_A,$
- $\mathcal{F}(\{g, ag, gb, agb\}) = \mathbb{F},$
- $\operatorname{res}_{\{g,ag\}\leftarrow\{g\}} = r_a \otimes \operatorname{id} : C_A \otimes_{\mathbb{F}} C_B \to C_B,$
- $\operatorname{res}_{\{g,gb\}\leftarrow\{g\}} = c_b : C_A \otimes_{\mathbb{F}} C_B \to C_A,$
- $\operatorname{res}_{\{g,ag,gb,agb\}\leftarrow\{g,ag\}}: C_B \to \mathbb{F}$ sends v to v_b ,
- $\operatorname{res}_{\{g,ag,gb,agb\}\leftarrow\{g,gb\}}: C_A \to \mathbb{F}$ sends u to u_a ,
- $\operatorname{res}_{\{q,aq,qb,aqb\}\leftarrow\{q\}}: C_A \otimes C_B \to \mathbb{F}$ sends m (an $A \times B$ -matrix) to $m_{a,b}$.

Lemma 9.1. With notation as above, \mathcal{F} is a well-defined sheaf on X.

Proof. Recall that condition 9.1 guarantees that that for every $v = \{g\} \in X(0)$ and $e \in X(1)$, either $e = \{g, ag\}$ for a unique a, or $e = \{gb, g\}$ for a unique b. This shows that the restriction map $\operatorname{res}_{e \leftarrow v}^{\mathcal{F}}$ is well-defined. Similarly, all the restriction maps are well-defined. Since $\mathcal{F}(\emptyset) = 0$, it remains to check that if $s \in X(2)$, $e \in s(1)$ and $v \in e(0)$, then $\operatorname{res}_{s \leftarrow v} = \operatorname{res}_{s \leftarrow e} \circ \operatorname{res}_{e \leftarrow v}$. Writing $g = \{g\}$ with $g \in G$, suppose that $e = \{g, ag\}$ for $a \in A$. Then there is $b \in B$ such that $s = \{g, ag, gb, agb\}$ and for every $m \in C_A \otimes C_B$, we have $\operatorname{res}_{s \leftarrow e} \operatorname{res}_{e \leftarrow v}(m) = (r_a(m))_b = m_{a,b} = \operatorname{res}_{s \leftarrow v}$, as required. The case $e = \{g, bg\}$ with $b \in B$ is handled similarly.

9.3 The Code and Its Tester

Keeping the previous notation, put $\Sigma = C_A \otimes C_B$ and observe that $\mathcal{F}(v) = \Sigma$ for every $v \in X(0)$. We may therefore form the 0-cocycle code

$$Z^0 = Z^0(X, \mathcal{F}) \subseteq \Sigma^{X(0)} = \Sigma^G.$$

Recall that we chose the \mathbb{F} -orientation on X in such a manner that $[v:\emptyset] = 1$ for every $v \in X(0)$. This implies that for every edge $e \in X(1)$ with vertices u and v, we have [e:v] = -[e:u]. As a result, Z^0 consists of the words $f = (f(g))_{g \in G} \in \Sigma^G = (C_A \otimes C_B)^G$ which satisfy

$$\operatorname{res}_{\{g,ag\}\leftarrow\{g\}} f(g) = \operatorname{res}_{\{g,ag\}\leftarrow\{g\}} f(ag) \quad \text{and} \quad \operatorname{res}_{\{g,gb\}\leftarrow\{g\}} f(g) = \operatorname{res}_{\{g,gb\}\leftarrow\{gb\}} f(gb)$$

for all $a \in A$, $b \in B$. Since $\{g, ag\} = \{ag, a^{-1}(ag)\}$ and $\{g, gb\} = \{gb, gb(b^{-1})\}$, this is equivalent to

$$r_a(f(g)) = r_{a^{-1}}(f(ag))$$
 and $r_b(f(g)) = r_{b^{-1}}(f(gb)),$

respectively. This can be further restated as saying that $f(g)_{a,b} = f(ag)_{a^{-1},b}$ and $f(g)_{a,b} = f(gb)_{a,b^{-1}}$ for all $a \in A, b \in B, g \in G$. We conclude that Z^0 may be viewed as the space of ensembles of matrices $(m_g)_{g\in G} \in \Sigma^G$ such that

$$(m_g)_{a,b} = (m_{ag})_{a^{-1},b} = (m_{gb})_{a,b^{-1}} \qquad \forall a \in A, b \in B, g \in G.$$
(9.2)

From the above description, we see that the natural tester of Z^0 operates as follows: Given $f = \{f(g)\}_{g \in G} \in \Sigma^G \subseteq M_{A \times B}(\mathbb{F})^G$, choose $g \in G$ and $x \in A \sqcup B$ uniformly at random. When $x \in A$, accept f if and only if $r_x(f(g)) = r_{x^{-1}}(f(xg))$ and when $x \in B$, accept f if and only if $r_x(f(g)) = r_{x^{-1}}(f(xg))$.

Now that we have described the code $Z^0 \subseteq \Sigma^G$, we bound its rate from below.

Proposition 9.2. Let $G, A, B, X, C_A, C_B, \mathcal{F}$ be as above, and write $r_A = r(C_A)$ and $r_B = r(C_B)$. Then $r(Z^0(X, \mathcal{F})) \geq \frac{4r_A r_B - 3}{4r_A r_B}$.

Proof. By replacing g, a, b with agb, a^{-1}, b^{-1} in (9.2), we see that Z^0 is the subspace of Σ^G defined by the constraints

$$(m_g)_{a,b} = (m_{ag})_{a^{-1},b} = (m_{gb})_{a,b^{-1}} = (m_{agb})_{a^{-1},b^{-1}}.$$
(9.3)

for all $g \in G$, $a \in A$, $b \in B$. One readily checks that (9.3) depends only on the square $\{g, ag, gb, agb\}$ and not on g, a, b. Thus, Z^0 is defined by $3|X(2)| = \frac{3|G||A||B|}{4}$ linear constraints inside Σ^G . It follows that

$$\dim Z^0 \ge \dim \Sigma^G - \frac{3|G||A||B|}{4} = |G||A||B|r_A r_B - \frac{3|G||A||B|}{4} = \dim \Sigma^G (1 - \frac{3}{4r_A r_B}),$$

and the proposition follows.

9.4 Interpretation of Local Expansion Conditions

Our goal is not to apply Theorem 8.10 to $X = \operatorname{Cay}(A, G, B)$ and the sheaf \mathcal{F} we constructed. To that end, we now unfold the conditions (1a)–(2c) of that theorem and interpret them (for (X, \mathcal{F}) in terms of the codes C_A and C_B and the one-sided Cayley graphs $\operatorname{Cay}(A, G)$ and $\operatorname{Cay}(G, B)$.

We begin with restating (1a). Recall that $\delta(C_A)$ denotes the distance of the code C_A .

Lemma 9.3. With notation as §9.2, let $g \in G$, $a \in A$ and $b \in B$. Then

(i)
$$\operatorname{cbe}_{-1}(\mathcal{F}_{\{q,aq\}}) = \delta(C_B)$$

(*ii*)
$$\operatorname{cbe}_{-1}(\mathcal{F}_{\{g,gb\}}) = \delta(C_A),$$

(*iii*)
$$\operatorname{cbe}_{-1}(\mathcal{F}_{\{q\}}) = \frac{1}{2}(\delta(C_A) + \delta(C_B)).$$

Proof. (i) Write $e = \{g, ag\}$. Every 2-face in X containing e is of the form $s_b := \{g, ag, gb, agb\}$ for a unique $b \in B$. Recall that $\mathcal{F}(\{g, ag\}) = C_B$, and for every $b \in B$, $\mathcal{F}(s_b) = \mathbb{F}$ and $\operatorname{res}_{s_b \to e} C_B \to \mathbb{F}$ is projection onto the b-component. This means that the map $f \mapsto (\operatorname{res}_{s_b \to e}(f))_{b \in B} : \mathcal{F}(e) \to \prod_{b \in B} \mathcal{F}(s_b)$ is just the inclusion map $C_B \to \mathbb{F}^B$. The natural weight function on the 0-poset X_e is uniform on $X_e(0)$ (Example 4.8), so, as noted in Example 4.8, $\operatorname{cbe}_{-1}(\mathcal{F}_{\{q,aq\}}) = \delta(C_B)$. (ii) This is similar to (i).

(iii) Write $v = \{g\}$. There are |A| + |B| edges containing v, namely, $\{e_a := \{g, ag\}\}_{a \in A}$ and $\{e_b := \{g, gb\}\}_{b \in B}$, and one readily checks that $w_v(e_a) = \frac{1}{2|A|}$ and $w_v(e_b) = \frac{1}{2|B|}$ for all $a \in A$ and $b \in B$. Let $f \in \mathcal{F}(v) - \{0\} = C_A \otimes C_B - \{0\}$, and let $A_0 = \{a \in A : \operatorname{res}_{e_a \leftarrow v} f \neq 0\}$ and $B_0 = \{b \in B : \operatorname{res}_{e_b \leftarrow v} f \neq 0\}$. By Example 6.4, we need to show that $w_v(A_0 \cup B_0) \geq \frac{1}{2}(\delta(C_A) + \delta(C_B))$, and that equality is attained for some choice of f.

Recall that we view f as a matrix in $M_{A \times B}(\mathbb{F})$ with $r_a(f) \in C_B$ and $c_b(f) \in C_A$ for all $a \in A$, $b \in B$. Then $A_0 = \{a \in A : r_a(f) \neq 0\}$ and $B_0 = \{b \in B : c_b(f) \neq 0\}$. Since $f \neq 0$, there are $a_0 \in A$ and $b_0 \in B$ such that $f_{a_0,b_0} \neq 0$. This means that $r_{a_0}(f) \in C_B - \{0\}$, so at least $\delta(C_B)|B|$ entries in the a_0 -th row of f are nonzero. As a result, $|B_0| \geq \delta(C_B)|B|$. Similarly, $|A_0| \geq \delta(C_A)|A|$ and it follows that $w_v(A_0 \cup B_0) = w_v(A_0) + w_v(B_0) \geq \frac{1}{2}(\delta(C_B) + \delta(C_A))$. To see that equality can be attained, choose $f_A \in C_A$, $f_B \in C_B$ with $||f_A|| = \delta(C_A)$ and $||f_B|| = \delta(C_B)$, and take $f = f_A \otimes f_B$ (i.e., the matrix $(f_A(a)f_B(b))_{a \in A, b \in B}$).

We proceed with (1b).

Lemma 9.4. With notation as above, let $g \in G$ and $\kappa \in [0, \infty)$. Then $cbe_0(\mathcal{F}_{\{g\}}) \geq \kappa$ if and only if $C_A \otimes C_B$ is κ -agreement testable (Example 2.6).

Proof. By Proposition 6.6, changing the \mathbb{F} -orientation of $X_{\{g\}}$ is harmless. We therefore choose a \mathbb{Z} -orientation on $X_{\{g\}}$ such that $[v: \{g\}] = 1$ for every $v \in X_{\{g\}}(0)$; this is possible by Example 4.26.

Observe that $X_{\{g\}}$ may be identified with the complete bipartite graph on A and B — simply map the vertices $\{g, ga\}$ and $\{g, bg\}$ to a and b, respectively, and the edge $\{g, ag, gb, agb\}$ to the dege $\{a, b\}$ $(a \in A, b \in B)$. It is now routine to check using Examples 6.5 and 2.6 that $cbe_0(\mathcal{F}_{\{g\}}) \geq \kappa$ if and only if $C_A \otimes C_B$ is κ -agreement testable.

Condition (2a) of Theorem 8.10 holds automatically for X with B = 1.

Lemma 9.5. With notation as above, for every $g \in X$, the graph $X_{\{q\}}$ is a (0,1)-skeleton expander.

Proof. The graph $X_{\{g\}}$ is the complete bipartite graph on the sets $\{\{g, ag\} | a \in A\}$ and $\{\{g, gb\} | b \in B\}$. It is well-known that such a graph is a 0-expander, so the lemma follows from Proposition 2.2.

In order to secure (2b) and (2c), we need to require that the Cayley graphs Cay(A, G) and Cay(G, B) are λ -expanders (§2.1).

Lemma 9.6. With notation as above, suppose that both Cay(A, G) and Cay(G, B) are λ -expanders. Then:

- 1. $X(\leq 1)$ is a λ -expander and a $(\lambda, 1)$ -skeleton expander.
- 2. NIG^{1,1,2}(X) is a $(2\lambda, 4 \max\{|A|, |B|\})$ -skeleton expander.

Proof. (i) Let \mathcal{A} , \mathcal{A}_A , and \mathcal{A}_B denote the weighted adjacency operators of $X(\leq 1)$, $\operatorname{Cay}(A, G)$ and $\operatorname{Cay}(G, B)$, respectively. One readily checks that $\mathcal{A}_A \mathcal{A}_B = \mathcal{A}_B \mathcal{A}_A$ and $\mathcal{A} = \frac{1}{2}(\mathcal{A}_A + \mathcal{A}_B)$. The former means that \mathcal{A}_A and \mathcal{A}_B can be simultaneously diagonalized. Thus, every eigenvalue μ of \mathcal{A} on $C^0_\circ(X(\leq 1), \mathbb{R})$ is of the form $\frac{1}{2}(\mu_A + \mu_B)$ where μ_A , μ_B are eigenvalues of \mathcal{A}_A , \mathcal{A}_B . As both $\mu_A, \mu_B \subseteq [-1, \lambda]$, we conclude that $\mu \leq \lambda$.

(ii) Observe that $G := \text{NIG}^{1,1,2}(G)$ is the disjoint union of two subgraphs: G_A and G_B . The vertex set of G_A is $\{\{g, ag\} | a \in A, g \in G\}$ and the vertex set of G_B is $\{\{g, ag\} | a \in A, g \in G\}$. We will prove (ii) in two steps.

Step 1. We claim that G_A , endowed with its natural weight function, is a $(\lambda, |A|)$ -skeleton expander. Likewise, G_B is a $(\lambda, |B|)$ -skeleton expander.

To see this, let G'_A be the graph with vertex set $\{(g, ag) \mid g \in G, a \in A\}$ and edges

$$\{\{(g, ag), (gb, agb)\} \mid g \in G, a \in A, b \in B\}.$$

The map $p: G'_A \to G_A$ sending (g, ag) to $\{g, ag\}$ and $\{(g, ag), (gb, agb)\}$ to $\{g, ag, gb, agb\}$ is a 2-covering of G_A . On the other hand, the graph G'_A is the disjoint union of |A| copies of $\operatorname{Cay}(G, B)$. By Example 7.7(i) and (iii), $\operatorname{Cay}(G, B)$ is a $(\lambda, 1)$ -skeleton expander and hence G'_A is a $(\lambda, |A|)$ -skeleton expander. Since for every $U \subseteq G_A(0)$, we have $p^{-1}(E_{G_A}(U)) = E_{G'_A}(p^{-1}(U))$ and $w_{G'_A}(p^{-1}(U)) = w_{G_A}(U)$, it follows that G_A is also a $(\lambda, |A|)$ -skeleton expander, as claimed.

Step 2. Write H for the hypergraph $NIH^{1,1,2}(X)$ and let w_H be its weight function. Let w_A and

 w_B denote the natural weight functions of G_A and G_B , respectively. The graph G_A is |B|-regular, so $w_A(e) = \frac{1}{G_A(0)} = \frac{2}{|A||G|}$ and $w_A(s) = \frac{1}{G_A(1)} = \frac{4}{|A||B||G|}$ for every $e \in G_A(0)$ and $s \in G_A(1)$. On the other hand, by unfolding the definition of w_H , one finds that $w_H(e) = \frac{1}{|A||G|} = \frac{1}{2} w_A(e)$ and $w_H(s) = \frac{4}{|A||B||G|} = w_A(s)$. Now, by Step 1, for every $U \subseteq G_A(0)$ with $w_H(U) = \alpha$, we have

$$w_H(E_2(U)) = w_A(E_{G_A}(U)) \le |A|w_A(U)^2 + \lambda w_A(U) = 4|A|w_H(U)^2 + 2\lambda w_H(U) = 4|A|\alpha^2 + 2\lambda\alpha.$$

Similarly, for every $V \subseteq G_B(0)$ with $w_G(V) = \beta$, we have

$$w_G(E_G(V)) \le 4|B|\beta^2 + 2\lambda\beta.$$

Finally, let $Z \subseteq H(0)$ with $w_H(Z) = \gamma$, and put $U = Z \cap G_A(0)$ and $V = Z \cap G_B(0)$. Then, with α and β as before, we have

$$w_H(E_2(Z)) = w_H(E_2(U) \cup E_2(V)) \le w_H(E_2(U)) + w_H(E_2(V)) \le 4(|A|\alpha^2 + |B|\beta^2) + 2\lambda(\alpha + \beta) \le 4\max\{|A|, |B|\}\gamma^2 + 2\lambda\gamma,$$

which is what we want.

9.5 Constructing 2-Query LTCs

We finally plug all our previous observations to Corollary 8.12 to get the following theorem, which implies Theorem 1.4 from the introduction.

Theorem 9.7. Let $G, A, B, X, \mathbb{F}, C_A, C_B, \mathcal{F}$ be as in §9.1 and let $\varepsilon \in (0, 1]$. Suppose that the following conditions are met:

- (1a') $\delta(C_A) > \varepsilon$,
- (1b') $\delta(C_B) \geq \varepsilon$,
- (1c') $C_A \otimes C_B$ is ε -agreement testable,
- (2) the Cayley graphs $\operatorname{Cay}(A,G)$ and $\operatorname{Cay}(G,B)$ are $\frac{\varepsilon^2}{6400}$ -expanders.

	_

View $Z^0 = Z^0(X, \mathcal{F})$ as a code inside $C^0(X, \mathcal{F}) = \Sigma^G$ (where $\Sigma = C_A \otimes C_B$). Then

$$\delta(Z^0) \ge \frac{79}{80}\varepsilon, \qquad r(Z^0) \ge \frac{4r(C_A)r(C_B) - 3}{4r(C_A)r(C_B)}$$

and the natural 2-query tester of Z^0 has soundness

$$\frac{1}{128 \max\{\frac{|A|}{|B|}, \frac{|B|}{|A|}\}[100\varepsilon^{-3} + \max\{|A|, |B|\}\varepsilon^{-1}]}.$$

Moreover, setting $\eta = \frac{1}{128(100\varepsilon^{-3} + \max\{|A|, |B|\}\varepsilon^{-1})}$, Algorithm 8.2 with $q = \frac{\varepsilon}{4(\varepsilon+80)}$ is a correction algorithm for words that are η -close to Z_0 .

When viewed as functions of ε , |A| and |B|, the order of magnitude of the distance and soundness of Z^0 , as well as the required expansion of $\operatorname{Cay}(A, G)$ and $\operatorname{Cay}(B, G)$, is the best we can get by using Corollary 8.12. However, we did not attempt to optimize the constants. We also remark that as |X(0)| = |G| grows, |A| and |B| must be $\Omega(\varepsilon^{-4})$ for (2') to hold, because by the Alon–Boppana Theorem $\lambda(\operatorname{Cay}(A, G)) \geq \frac{2\sqrt{|A|-1}}{|A|} - o(1)$, and likewise for $\operatorname{Cay}(G, B)$.

Proof. The claim about the rate is Lemma 9.2. Let $\lambda = \frac{\varepsilon^2}{6400}$. Assumptions (1a')–(2'), the lemmas in §9.4 and Lemma 7.8 imply that assumptions (1a) and (1b) of Theorem 8.10 hold, and in addition,

- $X_v(\leq 1)$ is an (0, 1)-skeleton expander for all $v \in X(0)$,
- $X(\leq 1)$ is a $(\lambda, 1)$ -skeleton expander,
- NIH^{1,1,2}(X) is a $(2\lambda, 4 \max\{|A|, |B|\})$ -skeleton expander.

Observe also that w(u) = w(v) for every $u, v \in X(0)$ and $w(e) \leq \max\{\frac{|A|}{|B|}, \frac{|B|}{|A|}\}w(e')$ for every $e, e' \in X(1)$. We may therefore apply Corollary 8.12 (with the constants E, E', \ldots taken from the last row of Table 1, M = 1 and $M' = \max\{\frac{|A|}{|B|}, \frac{|B|}{|A|}\}$) for any $h_0, h_{-1}, h_{||} \in \mathbb{R}_+$ such that

$$h_0 + 2\lambda + 4 \max\{|A|, |B|\}h_{||} + \frac{\lambda + h_{-1}}{h_0} \le \frac{\varepsilon}{8}.$$

(Note that the requirement $\lambda < E\varepsilon = \varepsilon$ holds automatically.) Our theorem is obtained by choosing $h_0 = \sqrt{\lambda}, h_1 = \lambda$ and $h_{||} = \frac{\varepsilon}{64 \max\{|A|,|B|\}}$.

In order to get a good 2-query LTC from Theorem 9.7, it remains to show that it can be applied to an infinite family of G, A, B, C_A, C_B satisfying assumptions (1a')-(2'). The existence of a suitable family has been shown in [DEL⁺22, §5–6], but we recall some details for the sake of completeness. Specifically, we will show the following.

Theorem 9.8. For every r > 0 and finite field \mathbb{F} , there are $m \in \mathbb{N}$ and $\varepsilon > 0$ for which there exist:

- (i) a sequence of groups $\{G_i\}_{i\in\mathbb{N}}$ with $|G_i| \to \infty$,
- (ii) symmetric generating subsets $A_i, B_i \subseteq G_i$ (for every $i \in \mathbb{N}$) satisfying (9.1), $|A_i| = |B_i| = m$, and such that $\operatorname{Cay}(A_i, G_i)$ and $\operatorname{Cay}(G_i, B_i)$ are $\frac{\varepsilon^2}{6400}$ -expanders,
- (iii) a linear code $C_0 \subseteq \mathbb{F}^m$ such that $r(C_0) \geq r$, $\delta(C_0) \geq \varepsilon$ and the tensor code $C_0 \otimes C_0$ is ε -agreement testable.

Thus, if $r > \frac{3}{4}$ and \mathcal{F}_i is the sheaf on $X_i := \operatorname{Cay}(A_i, G_i, B_i)$ constructed in §9.2 (with A_i, B_i, G_i in place of A, B, G) by choosing $C_{A_i} = C_{B_i} = C_0$, then Theorem 9.7 tells us that the family

$$\{Z_i(X_i,\mathcal{F}_i)\subseteq (C_0\otimes C_0)^{G_i}\}_{i\geq 0}$$

is a 2-query LTC with alphabet $\Sigma = C_0 \otimes C_0$.

Theorem 9.8 is obtained by combining two results.

Lemma 9.9 ([DEL⁺22, Lem .5.1]). For every 0 < r < 1 and finite field \mathbb{F} , there are $\delta_0, \kappa_0 > 0$ and $d_0 \in \mathbb{N}$ such that for any $D \in \mathbb{N}$ divisible by d_0 , there exists a linear code $C_0 \subseteq \mathbb{F}^D$ with $r(C_0) \ge r$, $\delta(C_0) \ge \delta_0$ and such that $C_0 \otimes C_0$ is κ_0 -agreement testable.

The proof in [DEL⁺22] consists of showing that a random LDPC code will satisfy all the requirements (for suitable δ_0 , κ_0) with positive probability as the length of the code grows. It is written under the assumption that $\mathbb{F} = \mathbb{F}_2$, but works for every finite field \mathbb{F} .

Lemma 9.10 ([DEL⁺22, Lem .5.2]). Let $d_0 \in \mathbb{N}$, let q be an odd prime number with $q \geq d_0^2$, and let $D = d_0 \lfloor \frac{q+1}{d_0} \rfloor$. Then for every $i \in \mathbb{N}$, there is an explicit group G_i of size $\Theta(q^{3i})$ admitting two symmetric generating sets A_i , B_i of size D which satisfy (9.1) and such that $\operatorname{Cay}(A_i, G_i)$ and $\operatorname{Cay}(G_i, B_i)$ are $4D^{-1/2}$ -expanders.

This is shown using known constructions of Ramanujan graphs.

Proof of Theorem 9.8. Recall that we are given 0 < r < 1 and a finite field \mathbb{F} . Let δ_0, κ_0 and d_0 be as in Lemma 9.9, and put $\varepsilon = \min\{\delta_0, \kappa_0\}$. Choose a prime number q sufficiently large so that $q_0 \geq d_0^2$ and $4D^{-1/2} \leq \frac{\varepsilon^2}{6400}$, where D is as in Lemma 9.10. Having fixed such a q, let G_i, A_i, B_i be the family promised by that lemma. Take m = D. Since m is divisible by d_0 , Lemma 9.9 supplies us with a code $C_0 \subseteq \mathbb{F}^m$ such that $r(C_0) \geq r$, $\delta(C_0) \geq \varepsilon$ and $C_0 \otimes C_0$ is ε -agreement testable. This is exactly what we want.

9.6 Realization as a Line Code

Let $G, A, B, X, \mathbb{F}, C_A, C_B, \mathcal{F}, \Sigma = C_A \otimes C_B$ be as in §9.1, §9.2. We finish this section by showing that $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$ is in fact the line code of a linear lifted code $C(A, G, B) \subseteq \mathbb{F}^G$ that was constructed in $[\text{DEL}^+22]$ (in the case $\mathbb{F} = \mathbb{F}_2$). The main result of [op. cit.] states that under conditions resembling those of Theorem 9.7, $C(A, G, B) \subseteq \mathbb{F}^G$ is a good LTC. We shall recover this result (with slightly different parameters) by applying the results of Section 3 to the lifted code of $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ and its line code $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$.

The code $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ is constructed as follows. For every $\{g\} \in X(0)$, there is a bijection $\varphi_g : A \times B \to X(2)_{\{g\}}$ given by $\varphi_g(a, b) = \{g, ag, gb, agb\}$. We use this bijection to identify $\mathbb{F}^{X(2)_{\{g\}}}$ with $\mathbb{F}^{A \times B} = \mathcal{M}_{A \times B}(\mathbb{F})$ and let C_g be the subspace of $\mathbb{F}^{X(2)_{\{g\}}}$ corresponding to $C_A \otimes C_B$ under this identification; formally, once viewing every $f \in C_A \otimes C_B$ as a function $f : A \times B \to \mathbb{F}$, we have

$$C_g = \{ f \circ \varphi_g^{-1} \mid f \in C_A \otimes C_B \}.$$

The code $C = C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ is the lifted code determined by the small codes $\{C_g \subseteq \mathbb{F}^{X(2)}_{\{G\}}\}_{g \in G}$.²⁰ That is,

$$C(A, G, B) = \{ f : X(2) \to \mathbb{F} : f|_{X(2)_{\{g\}}} \in C_g \text{ for all } g \in G \}.$$

Since every small code C_g is canonically identified with $\Sigma = C_A \otimes C_B$, we may form the line code $L = L(\{C_g\}_{g \in G}) \subseteq \Sigma^G$ of C(A, G, B) (§3). As we now show, this code is precisely $Z^0(X, \mathcal{F})$.

²⁰According to the conventions of §2.3, we should have denoted C_g as $C_{X(2)_{fol}}$.

Lemma 9.11. With notation as above, the line code of $C(A, G, B) = C(\{C_g\}_{g \in G}) \subseteq \mathbb{F}^{X(2)}$ is $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$.

Proof. The elements of the line code L of C(A, G, B) are the ensembles $f = (f_g)_{g \in G} \in (C_A \otimes C_B)^G$ which satisfy the following condition for all $g, h \in G$:

(*)
$$f_g \circ \varphi_g^{-1}$$
 agrees with $f_h \circ \varphi_h^{-1}$ on $X(2)_{\{g\}} \cap X(2)_{\{h\}}$.

Let $g, h \in G$. If $X(2)_{\{g\}} \cap X(2)_{\{h\}} = \emptyset$ or g = h, then condition (\star) holds. Otherwise, there are unique $a_0 \in A$ and $b_0 \in B$ such that $h \in \{a_0g, gb_0, a_0gb_0\}$. Suppose that $h = a_0g$. Then $X(2)_{\{g\}} \cap X(2)_{\{h\}} = \{\{g, a_0g, gb, a_0gb\} \mid b \in B\}$. Since $\varphi_g(a_0, b) = \{g, a_0g, gb, a_0gb\} = \{h, a_0^{-1}h, hb, a_0^{-1}hb\} = \varphi_h(a_0^{-1}, b)$, condition (\star) is equivalent to having

$$r_{a_0}(f_g) = r_{a_0^{-1}}(f_{a_0g}). (9.4)$$

Likewise, when $h = gb_0$, condition (\star) is equivalent to

$$c_{b_0}(f_g) = c_{b_0^{-1}}(f_{gb_0}). (9.5)$$

Finally, if $g = a_0 g b_0$, then $X(2)_{\{g\}} \cap X(2)_{\{h\}} = \{\{a_0 g, g b_0, a_0 g b_0\}\}$ and condition (*) becomes

$$(f_g)_{a_0,b_0} = (f_{a_0gb_0})_{a_0^{-1},b_0^{-1}}$$

However, this already follows from (9.4) and (9.5) (for all $g \in G$), because they imply that $(f_g)_{a_0,b_0} = (f_{a_0g})_{a_0^{-1},b_0} = (f_{a_0gb_0})_{a_0^{-1}gb_0^{-1}}$.

By comparing (9.4) and (9.5) with the description of $Z^0(X, \mathcal{F})$ in §9.3, we see that $L = Z^0(X, \mathcal{F})$.

Corollary 9.12. Let $G, A, B, X, \mathbb{F}, C_A, C_B$ be as in §9.1 and §9.2, and let $C = C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ be the lifted code constructed above. Let $\varepsilon \in (0, 1]$, and suppose that conditions (1a')-(2') of Theorem 9.7 hold. Then

$$\delta(C) \ge \frac{79\varepsilon}{80|A||B|}, \qquad r(C) \ge 4r(C_A)r(C_B) - 3\varepsilon$$

and the natural tester of C has soundness

$$\frac{1}{256 \max\{\frac{|A|}{|B|}, \frac{|B|}{|A|}\}[100\varepsilon^{-3} + \max\{|A|, |B|\}\varepsilon^{-1}] + 1}$$

Moreover, provided that $|A|, |B|, |\mathbb{F}|$ are fixed, C has a linear-time decoding algorithm able to correct words that are $\frac{1}{128|A||B|(100\varepsilon^{-3}+\max\{|A|,|B|\}\varepsilon^{-1})}$ -close to C.

Proof. Write $L = Z^0(X, \mathcal{F})$; this is the line code of C by Lemma 9.11. Theorem 9.7 provides us with lower bounds on $\delta(L)$, r(L) and the soundness μ of the natural 2-query tester of L, as well as a decoding algorithm. The lower bounds on $\delta(C)$ and r(C) are now obtained by applying Proposition 3.6; in our case $D_{\min} = D_{\max} = 4$ and $k_{\min} = k_{\max} = |A||B|$. Next, we apply Theorem 3.11 using the graph $X(\leq 1)$ and the labelling ℓ mapping a face x to $X(2)_x$, namely, the set of squares containing x. It implies that the natural tester of C has soundness $\frac{\mu}{2+\mu} = \frac{1}{2\mu^{-1}+1}$ (in our case $d_{\min} = d_{\max} = |A| + |B|$). Finally, the claim about the decoding algorithm follows from Proposition 3.7. By applying Corollary 9.12 to the family $\{G_i, A_i, B_i\}_{i \in \mathbb{N}}$ from Theorem 9.8 (with $\mathbb{F} = \mathbb{F}_2$), we recover the good locally testable lifted codes of [DEL⁺22]. Our bounds on the rate, distance, soundness and decoding are slightly different, though.

Remark 9.13. The proof that $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ is locally testable under assumptions (1a')-(2') of Theorem 9.7 in [DEL⁺22] implicitly establishes the local testability of the line code $Z^0(X, \mathcal{F}) \subseteq (C_A \otimes C_B)^G$ and then deduces from it the local testability of C(A, G, B). Specifically, observe that [DEL⁺22, Algorithm 1] is essentially our Algorithm 8.2 restricted to the X and \mathcal{F} constructed from G, A, B. Our proof of Corollary 9.12 illuminates that aspect of the proof as well as the hidden role of cosystolic expansion of sheaves.

Remark 9.14. In [DEL⁺22, §4.1], the authors give lower bounds on the rate of $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ that are slightly better than those of Corollary 9.12. Using them together with Proposition 3.6(i) gives lower bounds on the rate of $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$ that are better than those given in Lemma 9.2.

Remark 9.15. By Theorem 3.11, we can also reverse the argument in the proof of Corollary 9.12 and use the fact that $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ is locally testable and admits a linear-time decoding algorithm to deduce that $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$ admits a 2-query tester making it into an a good LTC with a linear-time decoding algorithm. Note, however, that the alluded 2-query tester is *not* the natural tester of $Z^0(X, \mathcal{F})$. Rather, it is the tester defined in Section 3 corresponding to the intersection graph of the sets $\{X(2)_{\{g\}} \mid g \in G\}$ (Example 3.8). Explicitly, given $f \in Z^0(X, \mathcal{F})$, this tester chooses uniformly at random a pair of vertices $\{g\}, \{h\} \in X(0)$ that are contained in a common square and checks whether $f(g) \in M_{A \times B}(\mathbb{F}) \cong \mathbb{F}^{X(2)_{\{g\}}}$ agrees with $f(h) \in M_{A \times B}(\mathbb{F}) \cong \mathbb{F}^{X(2)_{\{h\}}}$ on $X_i(2)_{\{g\}} \cap X(2)_{\{h\}}$. This is different from the natural tester of $Z^0(X, \mathcal{F})$ because $\{g\}$ and $\{h\}$ may be the *opposite* vertices of the a square in X_i .

The parameters of the code $Z^0(X, \mathcal{F}) \subseteq \Sigma^G$ that we would get by applying Theorem 3.11 to the LTC $C(A, G, B) \subseteq \mathbb{F}^{X(2)}$ (say, using [DEL⁺22]) would be much worse than those promised by Theorem 9.7.

10 Local Testability of Two-Layer Lifted Codes

In this section, we apply Theorem 8.10 to give a local criterion for a lifted code to be locally testable w.r.t. to its natural tester. This requires the lifted code to have some auxiliary extra structure. In particular, the local codes forming our lifted codes should be lifted codes themselves.

10.1 Two-Layer Lifted Codes

Recall that a lifted code $C \subseteq \Sigma^n$ is determined by small codes $\{C_s \subseteq \Sigma^s\}_{s \in S}$, where $S \subseteq P([n])$. We would like to consider lifted codes in which each small code $C_s \subseteq \Sigma^s$ is itself a lifted code. This structure can be neatly encoded using a 1-poset labelled by subsets of [n].

Definition 10.1 (Two-Layer Lifted Code). Let Σ be a finite alphabet and let $n \in \mathbb{N}$. A two-layer lifted code inside Σ^n consists of a triple $(X, \ell, \{C_e\}_{e \in X(1)})$, where

- X is a 1-poset,
- $\ell: X \to P([n])$ is a function assigning every face of X a subset of [n],
- C_e is a code inside $\Sigma^{\ell(e)}$ for every $e \in X(1)$,

and such that the following conditions hold:

(1) $\ell(x) = \bigcup_{y \in X: y > x} \ell(y)$ for all $x \in X(0) \cup X(-1)$; (2) $\ell(\emptyset) = [n]$.

In this case, for every $x \in X(0) \cup X(-1)$, we assign a lifted code $C_x \subseteq \Sigma^{\ell(x)}$ defined by

$$C_x = C(\{C_e\}_{e \in X(1)_x}) = \{f \in \Sigma^{\ell(x)} : f|_{\ell(e)} \in C_e \text{ for all } e \in X(1)\}.$$

The code $C_{\emptyset} \subseteq \Sigma^n$ will also be denoted as $C := C(X, \ell, \{C_e\}_{e \in X(1)})$. It can can also be realized as a lifted code w.r.t. the "bigger" small codes $\{C_v\}_{v \in X(0)}$.

Definition 10.2 (Natural Tester of a Two-Layer Lifted Code). Let $C = C(X, \ell, \{C_e\}_{e \in X(1)}) \subseteq \Sigma^n$ be a two-layer lifted code. The natural tester of C is its natural tester when realized as a lifted code w.r.t. the small codes $\{C_v\}_{v \in X(0)}$. Explicitly, given $f \in \Sigma^n$, the natural tester picks $v \in X(0)$ uniformly at random, probes f_i for every $i \in \ell(v)$, and accepts f if and only if $f|_{\ell(v)} \in C_v$.

10.2 Subset-Labelled *d*-Posets

The notion of a 1-poset labelled by subsets of [n] extends naturally to *d*-posets.

Definition 10.3 (S-Subset Labelled d-Poset). Let S be a finite set and let X be a d-poset. An S-subset labelling on X is a function $\ell: X \to P(S)$ such that

(1) $\ell(x) = \bigcup_{y \in X: y > x} \ell(y)$ for all $x \in X$ with dim x < d and (2) $\ell(\emptyset) = S$.

In this case, we call (X, ℓ) an S-subset labelled d-poset.

Example 10.4. (i) Let $n \in \mathbb{N}$, let V be a collection of subsets of [n] covering [n] and let X be a d-poset with X(0) = V. Define a labelling $\ell : X \to P([n])$ by sending \emptyset to [n] and every other x to $\bigcap_{v \in x(0)} v$ is an [n]-subset labelling on X. If X is big enough such that for every $x \in X$ with $\dim x < d$, the sets $\{\bigcap_{v \in y(0)} v \mid y \in X(\dim x + 1)_x\}$ cover $\bigcap_{v \in x(0)} v$, then $\ell : X \to P([n])$ is an [n]-subset labelling of X. This generalizes the setting of §1.8, which is essentially the case where X is a pure 2-dimensional simplicial complex.

(ii) Let $d \leq d'$ be natural numbers, let X' be a d'-poset and set S = X'(d'). Put $X = X'(\leq d)$ and define $\ell : X \to P([n])$ by $\ell(x) = Y(d)_x$. Then ℓ is an S-subset labelled d-poset.

An S-subset labelling on a d-poset X induces a normalized weight function $w_{\ell}: X \to \mathbb{R}_+$ given by $w_{\ell}(x) = \frac{1}{n} \sum_{j \in \ell(x)} \frac{1}{\#\{y \in X(\dim x) : j \in \ell(y)\}}$. The number $w_{\ell}(x)$ is also the probability of getting x by choosing $j \in S$ uniformly at random and then choosing a face $y \in X(\dim x)$ with $j \in \ell(y)$ uniformly at random.

Given integers $-1 \leq i \leq j \leq d$, we define the (i, j)-lower regularity and *i*-upper irregularity of the S-subset labelling ℓ to be

$$L_{i,j}(\ell) := \frac{\max_{s \in S, x \in X(i)} \#\{y \in X(j)_x : s \in \ell(y)\}}{\min_{s \in S, x \in X(i)} \#\{y \in X(j)_x : s \in \ell(y)\}} \quad \text{and} \quad U_i(\ell) := \frac{\max\{\#\ell(x) \mid x \in X(i)\}}{\min\{\#\ell(x) \mid x \in X(i)\}},$$

respectively. The *i*-th degree of ℓ is

$$D_i(\ell) = \max_{x \in X(i)} \#\ell(x).$$

Example 10.5. In the setting of Example 4.2(ii), $w_{\ell} : X \to \mathbb{R}_+$ is just the restriction of the natural weight function of X' to X. Moreover, $L_{i,j}(\ell) = L_{i,j,d'}(X')$, $U_i(\ell) = U_{i,d'}(X')$ and $D_i(\ell) = D_{i,d'}(X')$.

Remark 10.6. If X is lower regular and ℓ is lower regular in the sense that $L_{i,j}(\ell) = 1$ for all $-1 \leq i \leq j \leq d$, then w_{ℓ} is a proper weight function on X. This follows from Corollary 4.19 and the following observation.

An S-subset labelling ℓ on a d-poset X may be used to extend X into a (d+1)-poset $Y := X \sqcup S$, where the elements of S are viewed as (d+1)-faces and for $x \in X$ and $s \in S$, we have x < s if and only if $s \in \ell(x)$. The weight function $w_{\ell} : X \to \mathbb{R}_+$ is then just the restriction of the natural weight function of Y to X. Moreover, $L_{i,j}(\ell)$, $U_i(\ell)$ and $D_i(\ell)$ are just $L_{i,j,d+1}(Y)$, $U_{i,d+1}(Y)$ and $D_{i,d+1}(Y)$, respectively.

Example 10.7. Keep the setting of Example 10.4(ii) and suppose further that d' = d + 1. Then then poset $X \sqcup S$ associated to the S-subset labelling ℓ coincides with X'.

We finally note that if (X, ℓ) is an S-subset labelled d-poset and $z \in X$ is of dimension *i*, then the pair $(X_z, \ell|_{X_z})$ is an $\ell(z)$ -subset labelled (d - i - 1)-poset. We shall abbreviate $\ell|_{X_z}$ to ℓ_z . If we let $Y = X \sqcup S$ as above, then w_{ℓ_z} coincides with the natural weight function of Y_z , because this poset is just $X_z \sqcup \ell(z)$.

10.3 A Criterion for a 2-Layer Lifted Code to be Locally Testable

Let $C = C(X, \ell, \{C_e\}_{e \in X(1)}) \subseteq \Sigma^n$ be a two layer lifted code. The following theorem gives a criterion for C to be locally testable when (X, ℓ) is the [n]-subset labeled 1-poset underlying a [n]-subset labeled pure 2-dimensional regular cell complex.

Theorem 10.8. Let $F \in \mathbb{N}$ and $L \in [1, \infty)$. Then there exist constants $S, S', S'', T_1, \ldots, T_5 > 0$ (all are inverse-polynomial in F and L) such that the following hold:²¹ Let $n \in \mathbb{N}$ and let (X, ℓ) be an [n]-subset labelled pure 2-dimensional regular cell complex (see §10.2) such that

- (0a) $F_{i,2}^{\max}(X) \leq F$ for all $i \in \{0,1\}$;
- (0b) $L_{i,j}(\ell) \leq L$ for all integers $-1 \leq i < j \leq 2$;
- (0c) for all $u, v \in X(0)$ and $j \in \ell(u) \cap \ell(v)$, there is a path of edges from u to v such that $j \in \ell(e)$ for every edge e along the path.

Let R be a commutative ring, let Σ be an R-module, and for every $e \in X(1)$, let $C_e \subseteq \Sigma^{\ell(e)}$ be a code which also an R-submodule. Then $(X(\leq 1), \ell, \{C_e\}_{e \in X(1)})$ is a 2-layer lifted code. Let $\alpha_0, \beta_0, \alpha_{-1}, \beta_{-1}, \alpha_{||}, \beta_{||} > 0$ and suppose that

- (1a) $\delta(C_e) \geq \varepsilon$ for all $e \in X(1)$,
- (1b) for every $v \in X(0)$, the quartet $(\{C_e\}_{e \in X(1)_v}, X_v(\leq 1), \ell|_{X_v(<1)}, w_{\ell_v})$ is an ε -agreement tester,²²
- (2a) $(X_v(\leq 1), w_{\ell_v})$ is an (α_0, β_0) -skeleton expander for all $v \in X(0)$;
- (2b) $(X(\leq 1), w_{\ell})$ is an $(\alpha_{-1}, \beta_{-1})$ -skeleton expander;

²¹We encourage the reader to think of F and L (and thus of S, S', T_1, \ldots, T_5) as being constant or $\Theta(1)$ as this is usually what happens in practice.

²²Note that X_v is a graph by our assumption that X is a 2-dimensional regular cell complex. The weight function $w_{\ell_v}: X_v \to \mathbb{R}_+$ is given explicitly by $w_{\ell_v}(x) = \frac{1}{\#\ell(v)} \sum_{i \in \ell(x)} \frac{1}{\#\{y \in X(\dim x)_v : i \in \ell(y)\}}$.

(2c) NIH^{1,1,2}(X, w_ℓ) (see §7) is an $(\alpha_{||}, \beta_{||})$ -skeleton expander.

Suppose further that

$$\alpha_{-1}\beta_0 + S'\alpha_0 < S\varepsilon \tag{10.1}$$

and one can find $h_{-1}, h_0, h_{||} \in (0, 1]$ satisfying the following inequality

$$(\alpha_0 + \beta_0 h_0) + (\alpha_{||} + \beta_{||} h_{||}) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} \le S'' \varepsilon.$$
(10.2)

Then the two-layer lifted code $C = C(X(\leq 1), \ell, \{C_e\}_{e \in X(1)}) \subseteq \Sigma^n$ satisfies

$$\delta(C) > \frac{1}{U_0(\ell)D_0(\ell)} \cdot \frac{T_1(S\varepsilon - S'\alpha_0 - \alpha_{-1}\beta_0)}{\beta_{-1}}$$

and its natural tester has soundness

$$\frac{1}{U_0(\ell)^2 U_1(\ell)} \cdot \frac{T_2}{T_3^{-1} + h_0^{-1} h_1^{-1} + h_{||}^{-1}}$$

Moreover, $C \subseteq \Sigma^n$ has a linear-time decoding algorithm (the constant depends on $D_0(\ell)$, $|\Sigma|$, F, L, h_0) able to correct words that are η -close to C, where

$$\eta = \frac{1}{U_0(\ell)D_0(\ell)} \min\left\{\frac{T_4(S\varepsilon - S'\alpha_0 - \alpha_{-1}\beta_0)}{\beta_{-1}h_0^{-1}}, \frac{T_5}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}\right\}.$$

We prove Theorem 10.8 in the next subsection.

Example 10.9. Similarly to Example 10.4(ii), let X' be a d'-poset $(d' \ge 2)$ with X' = [n], and define $\ell : X \to P([n])$ by $\ell(x) = X'(d)_x$. Then $X := X'(\le 2)$ and $\ell|_X$ satisfy conditions (0a)–(0c) of Theorem 10.8 with L = L(X') and $F = \max\{F_{0,2}^{\max}(X'), F_{0,1}^{\max}(X')\}$. Moreover, the weight functions w_ℓ and w_{ℓ_v} ($v \in X(0)$) coincide with the natural weight functions of X' and X'_v , respectively. We can now choose an R-module Σ and R-submodules $C_e \subseteq \Sigma^{\ell(e)}$ for every $e \in X(1)$ and attempt to apply Theorem 10.8.

When the regular cell complex X from Theorem 10.8 is a simplicial complex and the labelling ℓ is lower regular in the sense of Remark 10.6(ii), we can use Theorem 10.8 together with Oppenheim's Trickling Down Theorem [Opp18] to get the following *local* criterion for showing that a two layered lifted code is locally testable.

Theorem 10.10. There are constants K, K' > 0 such that the following hold. Let X be a pure 2-dimensional simplicial complex, let $n \in \mathbb{N}$ and let $\ell : X \to P([n])$ be an [n]-subset labeling satisfying condition (0c) of Theorem 10.8 and such that $L_{i,j}(\ell) = 1$ for all $-1 \leq i < j \leq 2$. Let $(X(\leq 1), \ell, \{C_e\}_{e \in X(1)})$ is a 2-layer lifted code with alphabet Σ as in Theorem 10.8. Let $\varepsilon \in (0, 1]$ and suppose that:

(1a) $\delta(C_e) \geq \varepsilon$ for every $e \in X(1)$;

(1b) for every $v \in X(0)$, the quartet $(\{C_e\}_{e \in X(1)_v}, X_v(\leq 1), \ell|_{X_v(\leq 1)}, w_{\ell_v})$ is an ε -agreement tester;

(2a) $(X_v (\leq 1), w_{\ell_v})$ is a $K \varepsilon^2$ -spectral expander.

Then the natural tester of the two-layer lifted code $C = C(X(\leq 1), \ell, \{C_e\}_{e \in X(1)}) \subseteq \Sigma^n$ has soundness $\frac{K'\varepsilon^3}{U_0(\ell)^2 U_1(\ell)}$. Moreover, $\delta(C) \geq \frac{K'\varepsilon}{U_0(\ell)D_0(\ell)}$ and there is a linear-time decoding algorithm for words in Σ^n that are $\frac{1}{U_0(\ell)D_0(\ell)}K'\varepsilon^3$ -far from C.

Proof. Suppose for time being that K and K' were specified and $K \leq \frac{1}{2}$. Their values will be given later on.

Since X is simplicial, $F_{i,2}^{\max}(X) = 3$ for all $i \in \{0,1\}$. Moreover, $\text{NIH}^{1,1,2}(X)$ has now edges, so it is a (0,0)-skeleton expander. Since both X and ℓ are lower regular, w_{ℓ} is a proper (Remark 10.6(ii)) and w_{ℓ_v} coincides with $(w_{\ell})_v$ and is therefore also proper (Corollary 4.19). Now, by (2a), $(X_v \leq 1), w_{\ell_v}$ is a $(K\varepsilon^2, 1)$ -skeleton expander, and by Oppenheim's Trickling Down Theorem [Opp18], (X, w_ℓ) is a $\frac{K\varepsilon^2}{1-K\varepsilon^2}$ -spectral expander, and hence a $(2K\varepsilon^2, 1)$ -skeleton expander provided $K \leq \frac{1}{2}.$

Let $S, S', T_1, \ldots, T_5, E'$ be the constants guaranteed by Theorem 10.8 when F = 3 and L = 1. We choose K to be small enough to satisfy $(2 + S')K < \frac{S'}{2}$ and $4\sqrt{K} \le E'$.

We claim that we may apply Theorem 10.8 to our (X, ℓ) and 2-layered lifted code with $(\alpha_0, \beta_0) =$ $(K\varepsilon^2, 1), (\alpha_{-1}, \beta_{-1}) = (2K\varepsilon^2, 1), (\alpha_{||}, \beta_{||}) = (0, 0), h_0 = \sqrt{K\varepsilon}, h_{-1} = K\varepsilon^2 \text{ and } h_{||} = 1.$ Note that conditions (0a)-(2c) of Theorem 10.8 hold by the last paragraph, or by our assumptions, so it remains to check the inequalities (10.1) and (10.2). Indeed, by our choice of K,

$$\alpha_{-1}\beta_0 + S'\alpha_0 = (2+S')K\varepsilon^2 \le (2+S')K\varepsilon < S'\varepsilon$$

and

$$(\alpha_0 + \beta_0 h_0) + (\alpha_{||} + \beta_{||} h_{||}) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} = K\varepsilon^2 + \sqrt{K}\varepsilon + \frac{K\varepsilon^2 + K\varepsilon^2}{\sqrt{K}\varepsilon} \le 4\sqrt{K}\varepsilon < E'\varepsilon.$$

Note also that $S\varepsilon - S'\alpha_0 - \alpha_{-1}\beta_0 \geq \frac{S'}{2}\varepsilon$. Now, Theorem 10.8 tells that natural tester of the code $C \subseteq \Sigma^n$ has soundness $\frac{1}{U_0(\ell)^2 U_1(\ell)} \frac{T_2}{T_3^{-1} + K^{-1.5}\varepsilon^{-3} + 1}$. Moreover, $\delta(C) > \frac{1}{U_0(\ell)D_0(\ell)} \cdot T_1 \cdot \frac{1}{2}S_1\varepsilon$ and C has a linear-time decoding algorithm able to correct works that are η -far from C with $\eta = \frac{1}{D_0(\ell)U_0(\ell)} \min\{\frac{T_4S_1\varepsilon}{K^{-0.5}\varepsilon}, \frac{T_5}{K^{-1.5}\varepsilon^{-3}+1}\}$. From this, one sees that there is a constant K' > 0 for which the assertions of the theorem hold.

Proof of Theorem 10.8 10.4

We prove Theorem 10.8 by realizing the line code of $C = C(\{C_v\}_{v \in X(0)})$ as the 0-cocycle code of a sheaf on X, applying Theorem 8.10 to that sheaf, and then deducing the good properties of Cusing the results of Section 3. This will be done a series of lemmas. A byproduct of this approach is that the line code of $C = C(\{C_v\}_{v \in X(0)})$ is also locally testable and has linear distance; this is Lemma 10.14.

Throughout, we will use the following general notation: Let $n \in \mathbb{N}$ and let (X, ℓ) be an [n]subset labelled pure 2-dimensional regular cell complex. Let R be a ring and let Σ be an R-module. For every $e \in X(1)$, let $C_e \subseteq \Sigma^{\ell(e)}$ be a submodule, and let $C_v = C(\{C_e\}_{e \in X(1)_v}) \subseteq \Sigma^{\ell(v)}$ for all $v \in X(0)$. Recall that $C := C(X, \ell, \{C_e\}_{e \in X(1)})$ also equals $C(\{C_v\}_{v \in X(0)})$.

As in §10.2, let $Y = X \sqcup [n]$ be the 3-poset associated to (X, ℓ) . We denote the natural weight function of Y by w. Recall that $w_{\ell} = w|_X$ and $w_{\ell_v} = w_v$ for all $v \in X(0)$. In addition, $F_{i,2}^{\max}(Y) = F_{i,2}^{\max}(X)$ for all $i \in \{0,1\}$ and condition (0b) of Theorem 10.8 is equivalent to saying that $L_{i,j,3}(Y) \leq L$ for all $-1 \leq i \leq j \leq 3$. As $L_{i,j,k}(Y) \leq F_{i,j,k}^{\max}(Y)$ for all $i \leq j \leq k$, it follows that if conditions (0a) and (0b) of Theorem 10.8 hold, then $L(Y) \leq \max\{L, F\}$. Since $X = Y(\leq 2)$ is a regular cell complex, it admits a \mathbb{Z} -orientation with $[v:\emptyset] = 1$ for every $v \in X(0)$ (Example 4.26). We fix such an orientation once and for all. Note, however, that this orientation may not extends to Y. (Admittedly, Y is introduced only for the sake of the weights.)

Define a sheaf \mathcal{F} on Y as follows:

- $\mathcal{F}(v) = C_v \subseteq \Sigma^{\ell(v)}$ for all $v \in Y(0)$;
- $\mathcal{F}(e) = C_e \subseteq \Sigma^{\ell(e)}$ for all $e \in Y(1)$;
- $\mathcal{F}(x) = \Sigma^{\ell(x)}$ for all $x \in Y(2)$;
- $\mathcal{F}(y) = 0$ for every other face, i.e., $y \in Y(-1) \cup Y(3)$;
- $\operatorname{res}_{y \leftarrow x} : \mathcal{F}(x) \to \mathcal{F}(y)$ is given by $\operatorname{res}_{y \leftarrow x}(f) = f|_{\ell(y)}$ whenever $0 \le \dim x < \dim y \le 2$ (recall that $f \in \Sigma^{\ell(x)}$ and thus $f|_{\ell(y)} \in \Sigma^{\ell(y)}$);
- $\operatorname{res}_{y\leftarrow x}$ is the zero map in all other cases.

Lemma 10.11. Assume that condition (0c) of Theorem 10.8 holds. Then $Z^0(Y, \mathcal{F})$ is the line code of the lifted code $C := C(\{C_v\}_{v \in X(0)}) \subseteq \Sigma^n$.

Proof. Observe that $C^0(Y, \mathcal{F}) = \prod_{v \in X(0)} \mathcal{F}(v) = \prod_{v \in X(0)} C_v$. Let $f = (f_v)_{v \in X(0)} \in C^0(X, \mathcal{F})$. Since we chose the \mathbb{Z} -orientation on $X(\leq 2)$ such that $[v:\emptyset] = 1$ for all $v \in X(0)$, every $e \in X(1)$ admits exactly one vertex v with [e:v] = 1 and the other vertex u satisfies [e:u] = -1. Consequently, the condition (df)(e) = 0 is equivalent to having $\operatorname{res}_{e \leftarrow u} f(u) = \operatorname{res}_{e \leftarrow v} f(v)$. As a result, $Z^0(X, \mathcal{F})$ consists of the set of $f = (f_v)_{v \in X(0)} \in \prod_{v \in X(0)} C_v$ such that $f_u|_{\ell(e)} = f_v|_{\ell(e)}$ for every $e = \{u, v\} \in X(1)$. This means that $Z^0(X, \mathcal{F})$ contains the line code L of C. It remains to show that $L \supseteq Z^0(X, \mathcal{F})$.

Suppose that $f \in Z^0(X, \mathcal{F})$. In order to show that $f \in L$, we need to show that for every $u, v \in X(0)$ with $\ell(u) \cap \ell(v) \neq \emptyset$, we have $f_u|_{\ell(u) \cap \ell(v)} = f_v|_{\ell(u) \cap \ell(v)}$. Let $i \in \ell(u) \cap \ell(v)$. By condition (0d) of Theorem 10.8, there is a path of edges $e_1, \ldots, e_r \in X(0)$ from u to v such that $i \in \ell(e_j)$ for every $j \in \{1, \ldots, r\}$. Write $e_j = \{u_{j-1}, u_j\}$ so that $u = u_0$ and $v = u_r$. Our assumption that $f \in Z^0(X, \mathcal{F})$ implies that $f_{u_{j-1}}|_{\ell(e_j)} = f_{u_j}|_{\ell(e_j)}$, and in particular $f_{u_{j-1},i} = f_{u_{j,i}}$. Thus, $f_{u,i} = f_{u_0,i} = \cdots = f_{u_r,i} = f_{v,i}$. As this holds for all $i \in \ell(u) \cap \ell(v)$, we conclude that $f_u|_{\ell(u)\cap\ell(v)} = f_v|_{\ell(u)\cap\ell(v)}$.

Lemma 10.12. With notation as above, let $e \in X(1)$. Then $cbe_{-1}(Y_e, \mathcal{F}_e) \geq \delta(C_e)$.

Proof. Let $f \in \mathcal{F}(e) - 1$ and let $A = \{x \in Y(2)_e : \operatorname{res}_{x \leftarrow e} f \neq 0\}$. The $||df|| = w_e(A)$ and ||f|| = 1. We therefore need to show that $w_e(A) \geq \delta(C_e)$. In what follows, x ranges over $Y(2)_e$ and j ranges over $Y(3)_e = \ell(e)$. We have

$$w_e(A) = \sum_{x:f|_{\ell(x)}\neq 0} w_e(x) = \sum_{x:f|_{\ell(x)}\neq 0} \sum_{j:j>x} \frac{1}{|\ell(e)||\{y \in X(2)_e : e < y < j\}|}$$

$$\geq \sum_{j:f_j\neq 0} \sum_{x:e< x < j} \frac{1}{|\ell(e)||\{y \in X(2)_e : e < y < j\}|} = \frac{|\{j \in \ell(e) : f_j \neq 0\}|}{|\ell(e)|} \geq \delta(C_e). \quad \Box$$

Lemma 10.13. With notation as above, let $v \in Y(0)$, let $\alpha, \beta, \varepsilon > 0$ and let $F \in \mathbb{N}$. Suppose that $Y_v(\leq 1)$ is an (α, β) -skeleton expander, $\delta(C_e) \geq \varepsilon$ for every $e \in Y(1)_v$, and $F_{i,2}^{\max}(Y_v) \leq F$ for all $i \in \{0, 1\}$. Then there are constants Q, Q' > 0, depending only on F, such that

$$\operatorname{cbe}_{-1}(Y_v, \mathcal{F}_v) \ge \frac{Q'(Q\varepsilon - \alpha_0)}{\beta_0}.$$

The constants Q and Q' are E and E''' from Table 1 once taking $X = Y_v$. In particular, we can take Q = Q' = 1 when Y is lower regular.

Proof. Let \mathcal{F}'_v be the sheaf on Y_v obtained from \mathcal{F}_v by changing $\mathcal{F}_v(v) = C_v$ to 0. Then for every $e \in Y(1)_v$, we have $\operatorname{cbe}_{-1}(\mathcal{F}'_e) = \operatorname{cbe}_{-1}(\mathcal{F}_e) \geq \varepsilon$, where the inequality is By Lemma 10.12. Applying Theorem 8.13 to (Y_v, \mathcal{F}'_v) now tells us that $\operatorname{cse}_0(\mathcal{F}'_v) \geq \frac{E'''(E\varepsilon - \alpha)}{\beta}$. It is therefore enough to show that $\operatorname{cbe}_{-1}(\mathcal{F}_v) \geq \operatorname{cse}_0(\mathcal{F}'_v)$. Indeed, let $f \in C^{-1}(Y_v, \mathcal{F}_v) - \{0\}$. Since $\ell(v) = \bigcup_{e \in Y(1)_v} \ell(e)$, there is some $e \in X(1)_v$ such that $\operatorname{res}_{e \leftarrow v} f = f|_{\ell(e)} \neq 0$. As $df(e) = \pm \operatorname{res}_{e \leftarrow v} f$, it follows that $df \in Z^0(Y_v, \mathcal{F}_v) - \{0\} = Z^0(Y_v, \mathcal{F}'_v) - B^0(Y_v, \mathcal{F}'_v)$, and so $\|df\| \geq \operatorname{cse}_0(\mathcal{F}'_v) = \operatorname{cse}_0(\mathcal{F}'_v)\|f\|$, as claimed.

Lemma 10.14. Keep the notation above, let $F \in \mathbb{N}$ and $L \in [1, \infty)$, and let E, E', E'', E''', D, D', D'', Q, Q' be as in Table 1 and Lemma 10.13 (recall that these constants depend only on F and L). Suppose that $F_{i,2}^{\max}(X) \leq F$ for every $i \in \{0,1\}$, that $L(Y) \leq L$ and that assumptions (1a)-(2c) of Theorem 10.8 hold for some $\alpha_0, \beta_0, \alpha_{-1}, \beta_{-1}, \alpha_{||}, \beta_{||} > 0$. Suppose further that

$$\alpha_{-1}\beta_0 + EQ'\alpha_0 < EQQ'\varepsilon$$

and there are $h_0, h_{-1}, h_{||} \in (0, 1]$ such that

$$(\alpha_0 + \beta_0 h_0) + (\alpha_{||} + \beta_{||} h_{||}) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} \le E' \varepsilon.$$

Then

$$\csc_0(Y, \mathcal{F}) \ge \frac{E''}{h_0^{-1}h_1^{-1} + h_{||}^{-1}} \qquad and \qquad \operatorname{ccd}_0(Y, \mathcal{F}) \ge \frac{E'''(EQQ'\varepsilon - EQ'\alpha_0 - \alpha_{-1}\beta_0)}{\beta_{-1}}.$$

Moreover, if $f \in C^0 = C^0(Y, \mathcal{F})$ satisfies $\operatorname{dist}(f, Z^0) < \frac{D}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}$, then applying Algorithm 8.2 to f with the parameter $q = D'h_0$ returns $f' \in Z^0$ such that $\operatorname{dist}(f, f') \leq \frac{1}{D''h_0} \operatorname{dist}(f, Z^0)$.

Proof. We show this by applying Theorem 8.10 to (Y, \mathcal{F}) . (Note that Y may not admit an *R*-orientation, but $X = Y(\leq 2)$ has one and that is enough.) Condition (1b) of Theorem 8.10 holds with $\varepsilon' = \varepsilon$ by Lemma 10.12, condition (1c) of Theorem 8.10 holds by condition (1b) of Theorem 10.8 and condition (1a) of Theorem 8.10 holds with $\varepsilon = \frac{Q'(Q\varepsilon - \alpha_0)}{\beta_0}$ by Lemma 10.13. Our assumption that $\alpha_{-1}\beta_0 + EQ'\alpha_0 < EQQ'\varepsilon$ implies readily that $\alpha_{-1} < E \cdot \frac{Q'(Q\varepsilon - \alpha_0)}{\beta_0}$. As all other assumptions of Theorem 8.10 clearly hold, we may apply it (with $\frac{Q'(Q\varepsilon - \alpha_0)}{\beta_0}$ and ε in place of ε and ε') and derive all the assertions of the lemma.

In what follows, we shall view $Z^0(Y, \mathcal{F})$ as a code with coordinate-dependent alphabet inside $\prod_{v \in Y(0)} C_v = C^0(Y, \mathcal{F})$. It can be regarded as an honest code when all the C_v have the same cardinality, in which case they can all be identified with some alphabet Σ' . The following lemma says that $Z^0(Y, \mathcal{F}) \subseteq \prod_{v \in X(0)} C_v$ is locally testable and has linear distance under the assumptions of Theorem 10.8.

Lemma 10.15. Keep the notation and assumptions as in Lemma 10.14. Then the code $Z^0(Y, \mathcal{F}) \subseteq \prod_{v \in X(0)} C_v$ satisfies

$$\delta(Z^{0}(Y,\mathcal{F})) \geq \frac{1}{U_{0}(\ell)L_{0}(\ell)} \cdot \frac{E'''(EQQ' - EQ'\alpha_{0} - \alpha_{-1}\beta_{0})}{\beta_{-1}} := \delta_{0}$$

and its natural tester has soundness

$$\mu_0 := \frac{1}{U_0(\ell)U_1(\ell)L_0(\ell)L_1(\ell)} \cdot \frac{E''}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}.$$

Moreover, Algorithm 8.2 with the parameter $q = D'h_0$ is a decoding algorithm for $Z^0(Y, \mathcal{F}) \subseteq \Sigma^n$ which can fix words that are η_0 -close to $Z^0(Y, \mathcal{F})$, where

$$\eta_0 = \frac{1}{U_0(\ell)L_0(\ell)} \min\left\{\frac{E'''(EQQ'\varepsilon - EQ'\alpha_0 - \alpha_{-1}\beta_0)}{\beta_{-1}(1 + D''^{-1}h_0^{-1})}, \frac{D}{h_0^{-1}h_1^{-1} + h_{||}^{-1}}\right\}$$

Proof. By Proposition 4.24, for every $i \in \{0,1\}$ and $x, y \in X(i)$, we have $w(x) \leq U_{i,0}(Y)L_{i,0}(Y) \cdot w(y) = U_i(\ell)L_i(\ell) \cdot w(y)$. The assertion about the relative distance and soundness is therefore a consequence of Lemma 10.14 and Lemma 6.2 (with $M = U_0(\ell)L_0(\ell)$ and $M' = U_1(\ell)L_1(\ell)$). Moreover, Lemma 8.6 (with $M = U_0(\ell)L_0(\ell)$) tells us that Algorithm 8.2 with the parameter $q = D'h_0$ can decode words which are η_0 -close to $Z^0(X, \mathcal{F})$.

We finally prove Theorem 10.8.

Proof of Theorem 10.8. In short, this follows from Lemma 10.11, Lemma 10.15 and the results in Section 3. All that remains is checking the claims about the constants.

We use the notation of the Theorem 10.8 and construct Y and \mathcal{F} above using $X, R, \Sigma, \{C_e\}_{e \in X(1)}$ given in the theorem. We write $U_i := U_i(\ell) = U_{i,3}(Y)$ and $L_i := L_{-1,i}(\ell) = L_{i,3}(Y)$. Recall that assumptions (0a) and (0b) imply that $L(Y) \leq \max\{L, F\}$. We will specify the constants $S, S', S'', T_1, \ldots, T_5$ at the end.

Let $E, E', \ldots, \delta_0, \mu_0, \eta_0$ be as in Lemma 10.15 when applied with $\max\{L, F\}$ in place of L. We take S = EQQ', S' = EQ' and S'' = E'. The assumptions of Theorem 10.8 now imply that we may apply Lemma 10.15 with with the same (X, ℓ) , and $\{C_e\}_{e \in X(1)}$. Thus, $\delta(Z^0(Y, \mathcal{F})) \geq \delta_0$, the natural tester of $Z^0(Y, \mathcal{F})$ has soundness μ_0 and there is a decoding algorithm for words that are η_0 -close to $Z^0(Y, \mathcal{F})$ with time complexity O(|Y(0)|) = O(n). Here, the constant in the O(n) depends on $D_0, F, L, \max\{\#C_v \mid v \in Y(0)\}$ and h_0 (see Remark 8.4 and Proposition 4.24). However, $\max\{\#C_v \mid v \in Y(0)\} \leq |\Sigma|^{D_0(\ell)}$.

Let us realize $C := C(Y(\leq 1), \ell, \{C_e\}_{e \in Y(1)})$ as the lifted code $C(\{C_v\}_{v \in Y(0)}) \subseteq \Sigma^n$. By Lemma 10.11, the line code of C is $Z^0(Y, \mathcal{F}) \subseteq \prod_{v \in X(0)} C_v$. Now, by Proposition 3.6 (with $D_{\max} = F_{0,3}^{\max}(Y), D_{\min} = F_{0,3}^{\min}(Y), k_{\max} = D_0(\ell)$),

$$\delta(C) \geq \frac{F_{0,3}^{\min}(Y)}{F_{0,3}^{\max}(Y)D_0(\ell)} \delta_0 \geq \frac{1}{U_0 D_0} \cdot \frac{E'''(EQQ' - EQ'\alpha_0 - \alpha_{-1}\beta_0)}{L_0^2\beta_{-1}}$$

Furthermore, by Proposition 3.7, C has a linear-time decoding algorithms that can decode words that are η -close to C for

$$\eta = \frac{F_{0,3}^{\min}}{F_{0,3}^{\max} D_0(\ell)} \eta_0 \ge \frac{1}{U_0 D_0 L_0^2} \min\left\{ \frac{E'''(EQQ'\varepsilon - EQ'\alpha_0 - \alpha_{-1}\beta_0)}{\beta_{-1}(1 + D''^{-1}h_0^{-1})}, \frac{D}{h_0^{-1}h_1^{-1} + h_{||}^{-1}} \right\}$$
$$\ge \frac{1}{U_0 D_0} \min\left\{ \frac{E'''(EQQ'\varepsilon - EQ'\alpha_0 - \alpha_{-1}\beta_0)}{L^2(1 + D''^{-1})\beta_{-1}h_0^{-1}}, \frac{D}{L^2(h_0^{-1}h_1^{-1} + h_{||}^{-1})} \right\}.$$

Finally, by Theorem 3.10 (in our setting $\frac{d_{\max}}{d_{\min}} = \frac{F_{0,1}^{\max}}{F_{0,1}^{\min}} \leq F$ and $\frac{k_{\min}}{k_{\max}} = U_0(\ell)^{-1}$) the natural tester of C has soundness

$$\begin{aligned} \frac{1}{U_0(\ell)} \cdot \frac{F_{0,3}^{\min}(Y)}{F_{0,3}^{\max}(Y)} \cdot \frac{\mu_0}{\mu_0 + 2F} &\geq \frac{1}{U_0 L_0} \cdot \frac{1}{1 + 2F\mu_0^{-1}} = \frac{1}{U_0 L_0} \cdot \frac{1}{1 + 2FE''^{-1} U_0 U_1 L_0 L_1 (h_0^{-1} h_1^{-1} + h_{||}^{-1})} \\ &\geq \frac{1}{U_0 L_0} \cdot \frac{1}{U_0 U_1 + 2U_0 U_1 FE''^{-1} L_0 L_1 (h_0^{-1} h_1^{-1} + h_{||}^{-1})} \\ &\geq \frac{1}{U_0^2 U_1} \frac{\frac{1}{2} 2F^{-1} E'' L^{-3}}{\frac{1}{2} 2F^{-1} E'' L^{-2} + h_0^{-1} h_1^{-1} + h_{||}^{-1}}. \end{aligned}$$

One can now read from these assertions the required values for $S, S', S'', T_1, \ldots, T_5$.

11 Main Result: Technical Version

The remainder of this paper is dedicated to proving Theorems 8.1 and 8.10. We shall derive both theorems from a single, more general theorem, which we state in this section and prove in the following sections.

11.1 Notation

Let X be d-poset, let $k \in \{0, \ldots, d-2\}$ and let \mathcal{P} be an k-intersection profile for X (Section 7). We shall associate with X, k, \mathcal{P} a constant $U_{\mathcal{P}}$ and Laurent polynomials $T_k, \ldots, T_{-1}, S_{\alpha,\beta}$ in the variables $\{x_{\rho}\}_{\rho \in \mathcal{P}}$, making repeated use of the subface counting constants $F_{i,j,d}^{\min}$ and $F_{i,j,d}^{\max}$ defined in §4.4.

The constant $U_{\mathcal{P}} \in \mathbb{N}$ is defined as follows: Given $x \in X$, let $\mathcal{V}(x)$ denote the set of subsets A of $\{y \in X : y \leq x\}$ with $x \in A$ and the property that for every $a \in A$, there are faces $x = a_0 \geq \cdots \geq a_t = a$ in A such that $(\dim a_{i-1}, \dim a_i)$ is \mathcal{P} -admissible for all $i \in \{1, \ldots, t\}$. For such an A, let M(A) denote the set of minimal elements in A, and set

$$U(x) = \max_{A \in \mathcal{V}(x)} |M(A)|.$$

Finally, set

$$U_{\mathcal{P}} = U_{\mathcal{P}}(X) = \max_{x \in X(k+1)} U(x).$$

Next, let $\mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ be the ring of Laurent polynomials with real coefficients in the variables $\{x_{\rho}\}_{\rho \in \mathcal{P}}$. We define Laurent polynomials $T_k, T_{k-1}, \ldots, T_{-1} \in \mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ (depending on \mathcal{P} and X) inductively by the formula

$$T_{i} = \begin{cases} 1 & i = k, \\ \sum_{\rho \in \mathcal{P}: hgt \, \rho = i} x_{\rho}^{-1} \left[c_{\rho} T_{\ell(\rho)} + c_{\rho}' T_{r(\rho)} \right] & i = k - 1, k - 2, \dots, -1, \end{cases}$$

where

$$c_{(t,\ell,r,i)} = \frac{\frac{F_{\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{i,\ell,d}^{\min}}}{\frac{F_{\ell,d}^{\min}}{F_{i,\ell,d}^{\min}} + \frac{F_{r,d}^{\min}}{F_{i,r,d}^{\max}} \frac{F_{i,r,d}^{\min}F_{i,d}^{\min}}{F_{r,d}^{\max}}} \quad \text{and} \quad c_{(t,\ell,r,i)}' = c_{(t,r,\ell,i)}.$$

Suppose further that we are given vectors $\alpha = \{\alpha_{\rho}\}_{\rho \in \mathcal{P}}$ and $\beta = \{\beta_{\rho}\}_{\rho \in \mathcal{P}}$ in $\mathbb{R}^{\mathcal{P}}$. We then define $S_{\alpha,\beta} \in \mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ by

$$S_{\alpha,\beta} = \sum_{\rho=(t,\ell,r,i)\in\mathcal{P}} \frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}} \left[\frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}}T_{\ell} + \frac{F_{i,r,d}^{\max}F_{i,r}^{\max}}{F_{r,d}^{\min}}T_{r}\right] (\alpha_{\rho} + \beta_{\rho}x_{\rho}).$$

Example 11.1. We will see later in the proof of Lemma 11.4(i) below that if X is lower regular, then $c_{\rho} = c'_{\rho} = \frac{1}{2}$ for every $\rho \in \mathcal{P}$. This can also be seen directly using Lemma 4.11.

Suppose moreover that X is a pure d-dimensional simplicial complex and \mathcal{P} is the k-intersection profile from Example 7.15(i). Then $\mathcal{P} = \{(i+1, i, i, i-1) \mid i \in \{0, \ldots, k\}\}$ and $F_{i,j,\ell}^{\min} = F_{i,j,\ell}^{\max} = \binom{\ell-i}{\ell-j}$. Abbreviating $x_{(i+1,i,i,i-1)}$ to x_{i-1} , it is straightforward to check that

$$T_i(x_{k-1},\ldots,x_{-1}) = \frac{1}{x_i x_{i+1} \cdots x_{k-1}}$$

More bounds on the coefficients of the T_i and $S_{\alpha,\beta}$, as well as explicit computations for small k, will be given in §11.3 below.

11.2 Main Theorem

Let R be a commutative ring, let (X, w) be a properly weighed R-oriented d-poset, let $k \in \{0, \ldots, d-2\}$, let \mathcal{P} be a k-intersection profile for X and let \mathcal{P}' be a (k+1)-profile for X.

Let $\{\varepsilon_i\}_{i=0}^k$, $\{\varepsilon'_i\}_{i=0}^k$, $\alpha = \{\alpha_\rho\}_{\rho\in\mathcal{P}}$, $\beta = \{\beta_\rho\}_{\rho\in\mathcal{P}}$, $\alpha' = \{\alpha'_\rho\}_{\rho\in\mathcal{P}'}$, $\beta' = \{\beta'_\rho\}_{\rho\in\mathcal{P}'}$ be lists of non-negative real numbers. Let $U_{\mathcal{P}}$ and $T_k, \ldots, T_{-1}, S_{\alpha,\beta} \in \mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ as in §11.1. We define define $U_{\mathcal{P}'}$ and $T'_{k+1}, \ldots, T'_{-1}, S'_{\alpha',\beta'} \in \mathbb{R}[x_{\rho'}^{\pm 1} | \rho' \in \mathcal{P}']$ similarly by replacing \mathcal{P} and k with \mathcal{P}' and k+1, respectively. Finally, let

$$\tilde{\varepsilon} = \min\left\{\frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}\varepsilon_i \mid i \in \{0,\dots,k\}\right\},\\ \tilde{\varepsilon}' = \min\left\{\frac{F_{i,k+2,d}^{\min}F_{k+1,d}^{\min}}{F_{i,k+1,d}^{\max}F_{k+2,d}^{\max}}\varepsilon_i' \mid i \in \{0,\dots,k+1\}\right\}$$

and let

$$C = \max\left\{\frac{F_{i,k,d}^{\max}F_{i,d}^{\max}}{F_{k,d}^{\min}} \mid i \in \{0,\dots,k\}\right\}.$$

Theorem 11.2. With notation as above, let \mathcal{F} be an R-sheaf on X such that:

- (1a) $\operatorname{cbe}_{k-\dim u-1}(X_u, w_u, \mathcal{F}_u) \ge \varepsilon_{\dim u}$ for every $u \in X(0) \cup \cdots \cup X(k)$;
- (1b) $\operatorname{cbe}_{k-\dim u}(X_u, w_u, \mathcal{F}_u) \ge \varepsilon'_{\dim u}$ for every $u \in X(0) \cup \cdots \cup X(k+1)$;
- (2a) $\operatorname{NIH}_{u}^{\ell,r,t}(X)$ is an $(\alpha_{\rho},\beta_{\rho})$ -skeleton expander for every $\rho = (t,\ell,r,b) \in \mathcal{P}$ and $u \in X(b)$;
- (2b) $\operatorname{NIH}_{u}^{\ell,r,t}(X)$ is an $(\alpha'_{\rho},\beta'_{\rho})$ -skeleton expander for every $\rho = (t,\ell,r,b) \in \mathcal{P}'$ and $u \in X(b)$.²³

²³This condition may overlap with condition (2a) when $\mathcal{P} \cap \mathcal{P}' \neq \emptyset$.

Suppose that there exist $h = \{h_{\rho}\}_{\rho \in \mathcal{P}} \in (0,1]^{\mathcal{P}}, h' = \{h'_{\rho}\}_{\rho \in \mathcal{P}'} \in (0,1]^{\mathcal{P}} and q \in [0,1]$ such that

$$p := \tilde{\varepsilon} - U_{\mathcal{P}} S_{\alpha,\beta}(h) > 0 \qquad and \tag{11.1}$$

$$p' := \tilde{\varepsilon}' - U_{\mathcal{P}'} S'_{\alpha',\beta'}(h') - q \sum_{i=0}^{k} \frac{F_{i,k+2,d}^{\min} F_{k+1,d}^{\min}}{F_{i,k+1,d}^{\max} F_{k+2,d}^{\max}} \varepsilon'_{i} T'_{i}(h') > 0.$$
(11.2)

Then

$$\operatorname{cse}_{k}(X,\mathcal{F}) \geq \min\left\{T_{-1}'(h')^{-1}, \frac{q}{C}\right\} \quad and \quad \operatorname{ccd}_{k}(X,\mathcal{F}) \geq T_{-1}(h)^{-1}.$$

 $\begin{aligned} & \text{Moreover, if } f \in C^k(X, \mathcal{F}) \text{ satisfies } \operatorname{dist}(f, Z^k(X, \mathcal{F})) < \frac{F_{k+1,d}^{\min}}{F_{k,k+1,d}^{\max}} T'_{-1}(h')^{-1}, \text{ then applying Algorithm 8.2 to } f \text{ and } q \text{ returns } f' \in Z^k(X, \mathcal{F}) \text{ such that } \operatorname{dist}(f, f') < \frac{C}{q} \frac{F_{k,k+1,d}^{\max}F_{k,d}^{\max}}{F_{k+1,d}^{\max}} \operatorname{dist}(f, Z^k(X, \mathcal{F})). \end{aligned}$

Remark 11.3. The assumption that X is R-oriented in Theorem 11.2 can be relaxed to assuming that there is an R-orientation on the subposet $X(k) \cup X(k+1) \cup X(k+2)$.

The remainder of this section is dedicated to deriving Theorems 8.1 and 8.10 from Theorem 11.2. Theorem 11.2 itself will be proved in the next two sections.

11.3 Bounds on Constants

We first prove some lemmas which bound the constants and the coefficients of the Laurent polynomials defined in §11.1 and §11.2 in terms of the lower irregularity of X and the constants $F_{i,j}^{\max}$ (see §4.4).

Lemma 11.4. With notation as in §11.1, write $F = \max\{F_{i,j}^{\max} | -1 \le i \le j \le k+1\}$ and L = L(X). Then for every $\rho = (t, \ell, r, i) \in \mathcal{P}$, we have:

$$(i) \ c_{(t,\ell,r,i)} \leq \frac{L_{\ell,d}L_{i,\ell}L_{i,\ell,d}L_{i,d}}{L_{\ell,d}^{-1}L_{i,\ell,d}^{-1} + L_{r,d}^{-1}L_{i,r,d}^{-1}} \leq \frac{1}{2}L^{6},$$

$$(ii) \ \frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}} \leq L_{t,k+1,d}L_{k+1,d}L_{i,t,d}L_{t,d}^{2}\frac{F_{t,k+1}^{\max}F_{i,t}^{\max}}{2F_{i,d}^{\min}} \leq \frac{1}{2}L^{5}F^{2}$$

$$(iii) \ \frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}} \leq L_{i,\ell,d}L_{\ell,d}\frac{(F_{i,\ell}^{\max})^{2}}{F_{i,d}^{\min}} \leq L^{2}F^{2}.$$

In addition, for every $-1 \leq i \leq k$, we have:

$$(iv) \ \frac{1}{L^5 F} \le \frac{1}{L_{i,d} L_{i,k+1,d} L_{k,d} L_{i,k,d} L_{k+1,d}} \frac{F_{i,k+1}^{\min}}{F_{i,k}^{\max}} \le \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} \le L_{i,d} \frac{F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \le LF_{i,k}$$

$$(v) \ 1 \le F_{i,k}^{\min} \le \frac{F_{i,k,d}^{\max} F_{i,d}^{\max}}{F_{k,d}^{\min}} \le L_{i,k,d} L_{i,d} L_{k,d} F_{i,k}^{\max} \le L^3 F.$$

Proof. By Lemma 4.11, whenever $0 \le i \le j \le d$, we have

$$\frac{F_{i,j}^{\min}F_{j,d}^{\min}}{F_{i,d}^{\max}F_{i,j,d}^{\max}} \leq 1 \qquad \text{and} \qquad \frac{F_{i,d}^{\min}F_{i,j,d}^{\min}}{F_{i,j}^{\max}F_{j,d}^{\max}} \leq 1.$$

Making repeated use of these inequalities and the definition of $L_{i,j,k}$, we now prove each of (i)–(v) in turn.

$$\begin{split} c_{(t,\ell,r,i)} &= \frac{\frac{F_{\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{i,\ell,d}^{\min}}}{\frac{F_{i,\ell,d}^{\min}F_{i,\ell,d}^{\min}}{F_{i,\ell,d}^{\max}} + \frac{F_{r,d}^{\min}}{F_{r,d}^{\max}}\frac{F_{i,r,d}^{\min}F_{i,d}^{\min}}{F_{r,d}^{\max}}} \leq \frac{L_{\ell,d}L_{i,\ell}L_{i,\ell,d}\frac{F_{\ell,d}^{\min}F_{i,\ell}^{\min}}{F_{i,\ell,d}^{\max}F_{i,d}^{\max}}F_{i,d}^{\max}}{L_{\ell,d}^{-1}L_{i,\ell,d}^{-1}F_{i,d}^{\min}} + L_{r,d}^{-1}L_{i,r,d}^{-1}F_{i,d}^{\min}} \\ &\leq \frac{L_{\ell,d}L_{i,\ell,\ell}L_{i,\ell,d}L_{i,\ell,d}L_{i,d}}{L_{\ell,d}^{-1}L_{i,\ell,d}^{-1} + L_{r,d}^{-1}L_{i,r,d}^{-1}} \end{split}$$

This proves (i).

$$\frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}} = \frac{F_{t,k+1,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}}L_{t,d} \\
= \frac{F_{t,k+1,d}^{\max}F_{t,d}^{\min}}{F_{k+1,d}^{\min}F_{t,k+1}^{\max}}F_{t,k+1}^{\max}\frac{F_{i,t,d}^{\max}F_{i,d}^{\min}}{F_{t,d}^{\min}F_{i,t}^{\max}}\frac{F_{i,t,d}^{\max}}{2F_{i,d}^{\min}}L_{t,d} \\
= \frac{F_{t,k+1,d}^{\min}F_{t,d}^{\min}}{F_{k+1,d}^{\max}F_{t,k+1}^{\max}}L_{t,k+1,d}L_{k+1,d}\frac{F_{i,t,d}^{\min}F_{i,d}^{\min}}{F_{t,d}^{\max}F_{i,t}^{\max}}L_{i,t,d}L_{t,d}\frac{F_{t,k+1}^{\max}F_{i,t}^{\max}}{2F_{i,d}^{\min}}L_{t,d} \\
\leq L_{t,k+1,d}L_{k+1,d}L_{i,t,d}L_{t,d}L_{t,d}\frac{F_{t,k+1}^{\max}F_{i,t}^{\max}}{2F_{i,d}^{\min}}$$

This proves (ii).

$$\frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}} = \frac{F_{i,\ell,d}^{\min}F_{i,d}^{\min}}{F_{\ell,d}^{\max}F_{i,\ell}^{\max}} L_{i,\ell,d}L_{\ell,d}\frac{(F_{i,\ell}^{\max})^2}{F_{i,d}^{\min}} \le L_{i,\ell,d}L_{\ell,d}\frac{(F_{i,\ell}^{\max})^2}{F_{i,d}^{\min}}$$

This proves (iii). Next,

$$\frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}} = \frac{F_{k,d}^{\min}F_{i,k}^{\min}}{F_{i,k,d}^{\max}F_{i,k}^{\max}} \frac{F_{i,k+1,d}^{\min}F_{i,d}^{\min}}{F_{k+1,d}^{\max}F_{i,k+1}^{\max}} \frac{F_{i,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}F_{i,d}^{\min}} \le L_{i,d} \frac{F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \le L_{i,d} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \le L_{i,d} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\min}} \le L_{i,d} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\max}} \le L_{i,d} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1,d}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,k+1}^{\max}F_{i,k+1}^{\max}}{F_{i,k+1}^{\max}}} \frac{F_{i,$$

On the other hand,

$$\frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}} = \frac{F_{i,k+1,d}^{\max}F_{k,d}^{\max}}{F_{i,k,d}^{\min}F_{k+1,d}} \frac{1}{L_{i,k+1,d}L_{k,d}L_{i,k,d}L_{k+1,d}} \\
= \frac{F_{k,d}^{\max}F_{i,k}^{\max}}{F_{i,k,d}^{\min}F_{i,d}^{\min}} \frac{F_{i,k+1,d}^{\max}F_{i,d}^{\max}}{F_{k+1,d}F_{i,k+1}^{\min}} \frac{F_{i,d}^{\min}F_{i,k+1}^{\min}}{F_{i,k}^{\max}F_{i,d}^{\max}} \frac{1}{L_{i,k+1,d}L_{k,d}L_{k+1,d}} \\
\ge \frac{1}{L_{i,d}L_{i,k+1,d}L_{k,d}L_{i,k,d}L_{k+1,d}} \frac{F_{i,k+1,d}^{\min}F_{i,k+1}^{\min}}{F_{i,k}^{\min}}.$$

This proves (iv). Finally, for (v), note that

$$F_{i,k}^{\min} \le \frac{F_{i,k,d}^{\max} F_{i,d}^{\max}}{F_{k,d}^{\min}} = \frac{F_{i,k,d}^{\min} F_{i,d}^{\min}}{F_{i,k}^{\max} F_{k,d}^{\max}} L_{i,k,d} L_{i,d} L_{k,d} F_{i,k}^{\max} \le L_{i,k,d} L_{i,d} L_{k,d} F_{i,k}^{\max}.$$

This completes the proof.

Lemma 11.5. With notation as in §11.1, $U_{\mathcal{P}} \leq \sum_{i=0}^{k} F_{i,k+1}^{\max}$.
In practice, \mathcal{U}_P is smaller than $\sum_{i=0}^{k} F_{i,k+1}^{\max}$.

Proof. Let $x \in X(k + 1)$ and let $A \in \mathcal{V}(x)$ (see §11.1). We need to show that M(A), the set of minimal elements in A, contains at most $\sum_{i=0}^{k} F_{i,k+1}^{\max}$ elements. If $x \in M(A)$, then we must have $A = \{x\}$ and claim holds. If $\emptyset \in M(A)$, then we must have $M(A) = \{\emptyset\}$, and again the lemma claim holds. When both x and \emptyset are not in M(A), we have $M(A) \subseteq \bigcup_{i=0}^{k} x(i)$ and thus $|M(A)| \leq \sum_{i=0}^{k} F_{i,k+1}^{\max}$.

Lemma 11.6. With notation as in §11.1, let $F = \max\{F_{i,j}^{\max} \mid -1 \leq i \leq j \leq k+1\}$ and L = L(X). Let $u \in (0,1]$ and $A, B \in \mathbb{R}_+$, and for every $\rho = (t, \ell, r, b) \in \mathcal{P}$, let

$$h_{\rho} = u^{2^{k-b} - 2^{k-\min\{\ell,r\}}}, \qquad \alpha_{\rho} = A u^{2^{k-\min\{\ell,r\}}}, \qquad \beta_{\rho} = B u^{2^{k+1-\min\{\ell,r\}} - 2^{k-b}}.$$

Then:

- (i) For every $i \in \{-1, \ldots, k\}$, there is a real constant $Q_i > 0$ depending only on k such that $T_i(h) \leq Q_i L^{6(k-i)} u^{1-2^{k-i}}$.
- (ii) There is a real constant P > 0 depending only on k such that $S_{\alpha,\beta}(h) \leq PL^{6k+7}F^4(A+B)u$.

(*iii*)
$$\sum_{i=0}^{k-1} \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} T_i(h) \le (\sum_{i=0}^{k-1} Q_i) L^{6k+1} F u^{1-2^k} \text{ with } Q_0, \dots, Q_{k-1} \text{ as in } (i).$$

Proof. If \mathcal{P}' is a k-intersection profile for X containing \mathcal{P} , then replacing \mathcal{P} by \mathcal{P}' cannot decrease the left hand sides of the inequalities in (i), (ii) and (iii). We may therefore replace \mathcal{P} with the maximal k-intersection profile of Example 7.15. This ensures that the constants Q_i and P that we shall define depend only on k and not on \mathcal{P} . (However, a smaller \mathcal{P} allows for smaller constants Q_i and P.)

(i) Define the Q_i inductively for i = k, k - 1, ..., -1 by setting $Q_k = 1$ and

$$Q_i = \frac{1}{2} \sum_{\rho \in \mathcal{P}: \operatorname{hgt}(\rho) = i} (Q_{\ell(\rho)} + Q_{r(\rho)})$$

for i < k. The desired inequality now follows by decreasing induction on i. Indeed, the case i = k is clear since $T_k(h) = 1$. Assuming the inequality was verified for all $i \in \{k, k - 1, \dots, j + 1\}$, Lemma 11.4(i) tells us that

$$\begin{split} T_{j}(h) &= \sum_{\rho \in \mathcal{P}: \mathrm{hgt} \, \rho = j} h_{\rho}^{-1} \frac{L^{6}}{2} (T_{\ell(\rho)}(h) + T_{r(\rho)}(h)) \\ &\leq \frac{L^{6}}{2} \sum_{\rho \in \mathcal{P}: \mathrm{hgt} \, \rho = j} u^{2^{k - \min\{\ell(\rho), r(\rho)\}} - 2^{k - j}} (Q_{\ell(\rho)} L^{6(k - \ell(\rho))} u^{1 - 2^{k - \ell(\rho)}} + Q_{r(\rho)} L^{6(k - r(\rho))} u^{1 - 2^{k - r(\rho)}}) \\ &\leq \frac{L^{6}}{2} \sum_{\rho \in \mathcal{P}: \mathrm{hgt} \, \rho = j} (Q_{\ell(\rho)} + Q_{r(\rho)}) L^{6(k - j - 1)} u^{2^{k - \min\{\ell(\rho), r(\rho)\}} - 2^{k - j} + 1 - 2^{k - \min\{\ell(\rho), r(\rho)\}}} \\ &= Q_{j} L^{6(k - j)} u^{1 - 2^{k - j}}. \end{split}$$

(ii) Take $P = \frac{1}{2} \sum_{\rho \in \mathcal{P}} (Q_{\ell(\rho)} + Q_{r(\rho)})$. By (i) and Lemma 11.4(ii)–(iii),

$$\begin{split} S_{\alpha,\beta}(h) &= \sum_{\substack{\rho = (t,\ell,r,i) \in \mathcal{P} \\ \rho = (t,\ell,r,i) \in \mathcal{P} \\ \rho = (t,\ell,r,i) \in \mathcal{P} \\ } \frac{1}{2} L^{7} F^{2} \left(L^{2} F^{2} T_{\ell}(h) + L^{2} F^{2} T_{r}(h) \right) \left(\alpha_{\rho} + \beta_{\rho} h_{\rho} \right) \\ &\leq \sum_{\substack{\rho = (t,\ell,r,i) \in \mathcal{P} \\ \gamma = (t,\ell,r,i) \in \mathcal{P} \\ } \frac{1}{2} L^{7} F^{4} \left(Q_{\ell} L^{6(k-\ell)} u^{1-2^{k-\ell}} + Q_{r} L^{6(k-r)} u^{1-2^{k-r}} \right) \\ &\quad \cdot \left(A u^{2^{k-\min\{\ell,r\}}} + B u^{2^{k+1-\min\{\ell,r\}} - 2^{k-b} + 2^{k-b} - 2^{k-\min\{\ell,r\}}} \right) \\ &\leq \sum_{\substack{\rho = (t,\ell,r,i) \in \mathcal{P} \\ \gamma = (t,\ell,r,i) \in \mathcal{P} \\ } \frac{1}{2} L^{6k+7} F^{4} \left(Q_{\ell} + Q_{r} \right) \left(A + B \right) u^{1-2^{k-\min\{\ell,r\}} + 2^{k-\min\{\ell,r\}}} \\ &= P L^{6k+7} F^{4} (A + B) u. \end{split}$$

(iii) By (i) and Lemma 11.4(iv),

$$\sum_{i=0}^{k-1} \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} T_i(h) \le \sum_{i=0}^{k-1} LF \cdot Q_i L^{6(k-i)} u^{1-2^{k-i}} \le (\sum_{i=0}^{k-1} Q_i) L^{6k+1} F u^{1-2^k}.$$

Lemma 11.7. With notation as in §11.1, suppose that k = 0 and $\mathcal{P} = \{(1,0,0,-1)\}$ (cf. Example 7.15(ii)). We abbreviate the variable $x_{(1,0,0,-1)}$ to x_{-1} and similarly for other variables. Then, for every $\alpha_{-1}, \beta_{-1} \in \mathbb{R}_+$, we have:

$$T_{-1}(x_{-1}) = \frac{L_{0,d}^3}{x_{-1}} \quad and \quad S_{\alpha,\beta}(x_{-1}) = L_{1,d}^2 L_{0,d}(\alpha_{-1} + \beta_{-1}x_{-1}).$$

Proof. By direct computation.

Lemma 11.8. With notation as in §11.1, suppose that k = 1 and $\mathcal{P} = \{(2, 1, 1, 0), (2, 1, 1, -1), (1, 0, 0, -1)\}$ (cf. Example 7.15(iii)). We abbreviate the variables $x_{(2,1,1,0)}, x_{(1,0,0,-1)}, x_{(2,1,1,-1)}$ to $x_0, x_{-1}, x_{||}$ respectively, and similarly for other variables indexed by $\rho \in \mathcal{P}$. Suppose that $h, \alpha, \beta \in \mathbb{R}^{\mathcal{P}}_+$ and let L = L(X). Then:

(i)
$$T_0(h) \leq \frac{L^6}{h_0}$$
,
(ii) $T_{-1}(h) \leq \frac{L^{12}}{h_0h_{-1}} + \frac{L^6}{h_{||}}$,
(iii) $S_{\alpha,\beta}(h) \leq L^6 \frac{F_{0,2}^{\max}(F_{0,1}^{\max})^2}{(F_{0,d}^{\min})^2} (\alpha_0 + \beta_0 h_0) + L^6 (\alpha_{||} + \beta_{||} h_{-1}) + L^{13} F_{1,2}^{\max} \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0}$

Proof. For (i), note that $T_0(h) = \frac{c_{(2,1,1,0)}}{h_0}$ and $c_{(2,1,1,0)} \le L^6$ by Lemma 11.4(i). To see (ii), we use (i) and Lemma 11.4(i) to get

$$T_{-1}(h) = \frac{c_{(1,0,0,-1)}T_0(h)}{h_{-1}} + \frac{c_{(2,1,1,-1)}T_1(h)}{h_{||}} \le \frac{L^{12}}{h_0h_{-1}} + \frac{L^6}{h_{||}}.$$

Finally, by the definition of $S_{\alpha,\beta}$, parts (i) and (ii), and Lemma 11.4(ii)–(iii), we have

$$\begin{split} S_{\alpha,\beta}(h) &= L^4 L_{2,2,d} \frac{F_{2,2}^{\max} F_{0,2}^{\max}}{F_{0,d}^{\min}} \cdot L^2 \frac{(F_{0,1}^{\max})^2}{F_{0,d}^{\min}} \cdot (\alpha_0 + \beta_0 h_0) T_1(h) \\ &+ L^4 L_{2,2,d} \frac{F_{2,2}^{\max} F_{-1,2}^{\max}}{F_{-1,d}^{\min}} \cdot L^2 \frac{(F_{-1,1}^{\max})^2}{F_{-1,d}^{\min}} \cdot (\alpha_{||} + \beta_{||} h_{||}) T_1(h) \\ &+ L^5 \frac{F_{1,2}^{\max} F_{-1,1}^{\max}}{F_{-1,d}^{\min}} \cdot L^2 \frac{(F_{-1,0}^{\max})^2}{F_{-1,d}^{\min}} \cdot (\alpha_{-1} + \beta_{-1} h_{-1}) T_0(h) \\ &\leq L^6 \frac{F_{0,2}^{\max} (F_{0,1}^{\max})^2}{(F_{0,d}^{\min})^2} (\alpha_0 + \beta_0 h_0) + L^6 (\alpha_{||} + \beta_{||} h_{-1}) + L^{13} F_{1,2}^{\max} \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0}. \end{split}$$

This proves (iii).

11.4 Proofs of Theorems 8.1 and 8.10

We now use Theorem 11.2 to prove Theorems 8.1 and 8.10.

Proof of Theorem 8.1 assuming Theorem 11.2. We use the notation of Theorem 8.1. Recall that we are given $B \in \mathbb{R}_+$, $F \in \mathbb{N}$, $L \in [1, \infty)$ and $k \in \{0\} \cup \mathbb{N}$. Recall also that (X, w) is a properly weighted R-oriented d-poset $(d \ge k+2)$ with $L(X) \le L$ and $F_{i,j}^{\max}(X) \le F$ whenever $-1 \le i \le j \le k+2$. In addition \mathcal{P} is a k-intersection profile for X, \mathcal{P}' is a (k+1)-intersection profile for X and \mathcal{F} is an R-sheaf on X.

We will apply Lemma 11.6 both to k, \mathcal{P} and k + 1, \mathcal{P}' . The constants provided by the lemma in the latter case will be denoted $Q'_{k+1}, \ldots, Q'_{-1}$ and P'.

Let K > 0 be a constant depending on B, F, L, k to be specified later. Let $\varepsilon > 0$ and set

$$\varepsilon_i = \varepsilon'_i = \varepsilon$$

for every $i \in \{0, \ldots, k\}$ and $j \in \{0, \ldots, k+1\}$. For every $\rho = (t, \ell, r, b) \in \mathcal{P}$ and $\rho' = (t', \ell', r', b') \in \mathcal{P}'$, define

$$\begin{aligned} \alpha_{\rho} &= (K\varepsilon)^{2^{k-\min\{\ell,r\}}}, & \alpha'_{\rho'} &= (K\varepsilon)^{2^{k+1-\min\{\ell',r'\}}}, \\ \beta_{\rho} &= B(K\varepsilon)^{2^{k+1-\min\{\ell,r\}}-2^{k-b}}, & \beta'_{\rho'} &= B(K\varepsilon)^{2^{k+2-\min\{\ell',r'\}}-2^{k+1-b'}}, \\ h_{\rho} &= (K\varepsilon)^{2^{k-b}-2^{k-\min\{\ell,r\}}}, & h'_{\rho'} &= (K\varepsilon)^{2^{k+1-b'}-2^{k+1-\min\{\ell',r'\}}}. \end{aligned}$$

In addition, let

$$q = (K\varepsilon)^{2^{k+1}}$$

In order to prove Theorem 8.1, it is enough to show that K can chosen in such a way that that $K\varepsilon \leq 1$ and the inequalities (11.1) and (11.2) of Theorem 11.2 hold. Indeed, suppose that assumptions (0)–(2) of Theorem 8.1 hold. Then conditions (1a)–(2b) of Theorem 11.2 with the ε_i , ε'_i , α_ρ and β_ρ defined above hold for (X, w, \mathcal{F}) . Provided that (11.1) and (11.2) also hold, all the conclusions of Theorem 11.2 are true. Now, by applying Lemma 11.6(i) with $u = K\varepsilon \in (0, 1]$ (both for k, \mathcal{P} and k+1, \mathcal{P}'), we see that $T_{-1}(h) \leq Q_{-1}L^{6k+6}(K\varepsilon)^{1-2^{k+1}}$ and $T'_{-1}(h') \leq Q'_{-1}L^{6k+12}(K\varepsilon)^{1-2^{k+2}}$. In addition, by Lemma 11.4(v), we have $C \leq L^3F$ and $L^{-3}F^{-1} \leq \frac{F_{k+1,d}^{\min}}{F_{k,k+1,d}^{\max}} \leq 1$. Combining this

with the conclusions of Theorem 11.2, we get that

$$\operatorname{cse}_{k}(X, w, \mathcal{F}) \geq \min\left\{T_{-1}'(h')^{-1}, \frac{q}{C}\right\} \geq \min\left\{\frac{(K\varepsilon)^{2^{k+2}-1}}{Q_{-1}'L^{6k+12}}, \frac{Q(K\varepsilon)^{2^{k+1}}}{L^{3}F}\right\}$$
$$\geq \min\{(Q_{-1}')^{-1}L^{-6k-12}, QL^{-3}F^{-1}\}(K\varepsilon)^{2^{k+2}-1}$$

and

$$\operatorname{ccd}_k(X, w, \mathcal{F}) \ge T_{-1}(h)^{-1} \ge \frac{(K\varepsilon)^{2^{k+1}-1}}{Q_{-1}L^{6k+6}}.$$

Furthermore, applying Algorithm 8.2 with the parameter $q = (K\varepsilon)^{2^{k+1}}$ to any $f \in C^k$ with

$$\operatorname{dist}(f, Z^k) < L^{-3} F^{-1} (Q'_{-1} L^{6k+12})^{-1} (K\varepsilon)^{2^{k+2}-1} \le \frac{F_{k+1,d}^{\min}}{F_{k,k+1,d}^{\max} F_{k,d}^{\max}} T'_{-1} (h')^{-1}$$

results in $f' \in Z^k$ with

$$\operatorname{dist}(f, f') < \frac{C}{q} \frac{F_{k,k+1,d}^{\max} F_{k,d}^{\max}}{F_{k+1,d}^{\min}} \operatorname{dist}(f, Z^k(X, \mathcal{F})) \le (L^3 F)^2 (K\varepsilon)^{-2^{k+1}} \operatorname{dist}(f, Z^k(X, \mathcal{F})).$$

From this one readily sees that there is K' > 0, depending only on F, L, B, k, for which the assertions of Theorem 8.1 about (X, w, \mathcal{F}) hold. Explicitly,

$$K' = \min\left\{ (Q'_{-1})^{-1} L^{-6k-12}, L^{-3} F^{-1}, (Q_{-1})^{-1} L^{-6k-6}, (Q'_{-1})^{-1} L^{-6k-15} F^{-1}, L^{-6} F^{-2} \right\}$$

so provided that k is fixed, $K' = \Omega(L^{-6k-15}F^{-2})$.

We now show the existence of the constant K > 0. Note first that we can secure $K\varepsilon \leq 1$ by choosing $K \leq 1$, because $\varepsilon \leq 1$. Next, observe that Lemma 11.4(iv) implies that $\tilde{\varepsilon} \geq \frac{1}{L^5 F} \varepsilon$ and likewise for $\tilde{\varepsilon}'$. Consider the inequality (11.1). By Lemma 11.6 (with $u = K\varepsilon \leq 1$ and A = 1) and Lemma 11.5, we have

$$p = \tilde{\varepsilon} - U_{\mathcal{P}} S_{\alpha,\beta}(h) \ge \frac{\varepsilon}{L^5 F} - (k+1)F \cdot P L^{6k+7} F^4(1+B) K \varepsilon$$

so we can guarantee that p > 0 by taking $K < P^{-1}(k+1)^{-1}L^{-6k-12}F^{-6}(1+B)^{-1}$. Next, consider (11.2). By the same lemmas applied with k + 1 and \mathcal{P}' , we get

$$p' = \tilde{\varepsilon}' - U_{\mathcal{P}'} S'_{\alpha',\beta'}(h') - q \sum_{i=0}^{k} \frac{F_{i,k+2,d}^{\min} F_{k+1,d}^{\min}}{F_{i,k+1,d}^{\max} F_{k+2,d}^{\max}} \varepsilon'_{i} T'_{i}(h')$$

$$\geq \frac{\varepsilon}{L^{5}F} - (k+2)F \cdot PL^{6k+13}F^{4}(1+B)K\varepsilon - (K\varepsilon)^{2^{k}} (\sum_{i=0}^{k} Q'_{i})L^{6k+7}F(K\varepsilon)^{1-2^{k}}$$

$$\frac{\varepsilon}{L^{5}F} - \left[(k+2)F \cdot PL^{6k+13}F^{4}(1+B) + L^{6k+7}F\sum_{i=0}^{k} Q'_{i} \right] K\varepsilon,$$

and again, we can guarantee that p' > 0 by taking $K < \frac{1}{L^5 F}[(k+2)PL^{6k+13}F^5(1+B) + L^{6k+7}F\sum_{i=0}^k Q'_i]^{-1}$. In particular, once k is fixed, $K = \Omega(L^{-6k-18}F^{-6}(1+B)^{-1})$ works. This completes the proof. \Box Proof of Theorem 8.10 assuming Theorem 11.2. Recall that we are given an *R*-oriented properly weighted *d*-poset (X, w) of lower irregularity at most *L* and such that $F_{i,j,\ell}^{\max} \leq F$ whenever $-1 \leq i \leq j \leq \ell \leq 2$. We need to show that there are constants E, E', E'', D, D', D'' > 0, depending only *L* and *F*, for which the conclusions of Theorem 8.10 hold.

To that end, consider the 1-intersection profile \mathcal{P}_1 of Example 7.15(iii) and define

$$K = \max\left\{\frac{F_{0,1,d}^{\max}F_{2,d}^{\max}}{F_{0,2,d}^{\min}F_{1,d}^{\min}}, \frac{F_{1,1,d}^{\max}F_{2,d}^{\max}}{F_{1,2,d}^{\min}F_{1,d}^{\min}}\right\} \le L^5 \max\left\{\frac{F_{0,1}^{\max}}{F_{0,2}^{\min}}, \frac{1}{F_{1,2}^{\min}}\right\} \le L^5 F$$

$$N = U_{\mathcal{P}_1} \max\left\{L^6 \frac{F_{0,2}^{\max}(F_{0,1}^{\max})^2}{(F_{0,d}^{\min})^2}, L^{13}F_{1,2}^{\max}\right\} \le 2FL^6 \max\{F^3, L^7\}$$

$$V = L^7 \frac{F_{0,2}^{\max}}{F_{0,1}^{\min}} \le L^7 F.$$

The inequalities hold by Lemma 11.4(iv) and Lemma 11.5. Also let $C = \frac{F_{0,d}^{\max}F_{0,d}^{\max}}{F_{0,d}^{\min}} = L_{0,d} \leq L$ (this is the same as in Theorem 11.2 with k = 0). We will show that Theorem 8.10 holds for our X with any positive E, E', E'', E''', D, D', D'' satisfying

$$D \leq L^{-12} \frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max} F_{0,d}^{\max}} \qquad E \leq L_{0,d}^{-1} L_{1,d}^{-3} L^{-4}$$

$$D' \leq \frac{1}{4VK} \qquad E' \leq \frac{1}{2NK}$$

$$D'' \leq \frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max} F_{0,d}^{\max}} \frac{1}{4VCK} \qquad E'' \leq \min\left\{\frac{1}{L^{12}}, \frac{1}{2VKC}\right\}$$

$$E''' \leq L_{0,d}^{-3}$$

Since $\frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max}F_{0,d}^{\max}} \ge \frac{1}{L^3 F_{0,1}^{\max}}$ (Lemma 11.4(v)), we can choose values for these constants which depend only on L and F.

Let \mathcal{F} be an *R*-sheaf on *X* such that conditions (1a)–(2c) of Theorem 8.10 hold, $\alpha_{-1} < E\varepsilon$, and there are and $h_0, h_{-1}, h_{||} > 0$ satisfying

$$(\alpha_0 + \beta_0 h_0) + (\alpha_{||} + \beta_{||} h_{||}) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} \le E' \varepsilon'.$$
(11.3)

Let k = 0, and let \mathcal{P} and \mathcal{P}' be the intersection profiles \mathcal{P}_0 and \mathcal{P}_1 of Example 7.15. Now define $T_0, T_{-1}, T'_1, T'_0, T'_{-1}, S_{\alpha,\beta}, S'_{\alpha',\beta'}$ as in §11.2. Note that $\mathcal{P} = \{(1,0,0,-1)\}$ and $\mathcal{P}' = \{(2,1,1,0), (2,1,1,-1), (1,0,0,-1)\}$. We abbreviate the variables $x_{(2,1,1,0)}, x_{(1,0,0,-1)}$ and $x_{(2,1,1,-1)}$ used in the definition of T_0, T_1 , etc. to x_0, x_{-1} and $x_{||}$, respectively. Similarly, we view $\alpha' := (\alpha_0, \alpha_{-1}, \alpha_{||}), \beta' := (\beta_0, \beta_{-1}, \beta_{||})$ and $h' := (h_0, h_{-1}, h_{||})$ from Theorem 8.10 as vectors in $\mathbb{R}^{\mathcal{P}'}$. We also view the numbers α_{-1} and β_{-1} as a vectors $\alpha, \beta \in \mathbb{R}^{\mathcal{P}}$. Choose some $\gamma \in [\frac{1}{2}, 1)$, and set

$$h := \gamma \frac{E\varepsilon - \alpha_{-1}}{\beta_{-1}}$$
 and $q = \frac{\gamma h_0}{2VK}$.

We view h as a vector in $\mathbb{R}^{\mathcal{P}}$. Finally, set $\varepsilon_0 = \varepsilon$ and $\varepsilon'_0 = \varepsilon'_1 = \varepsilon'$, where ε and ε' are those given in Theorem 8.10.

We will prove Theorem 8.10 by applying Theorem 11.2 to (X, w, \mathcal{F}) and the parameters we chose. Assumptions (1a)–(2b) of Theorem 11.2 are precisely assumptions (1a)–(2c) of Theorem 8.10, so it remains to verify the inequalities (11.1) and (11.2). By Lemma 11.4(iv), $\tilde{\varepsilon} = \frac{F_{0.1.d}^{\min}F_{0.d}^{\min}}{F_{0.0,d}^{\max}F_{1.d}^{\max}} \varepsilon \ge L^{-4}F_{0,1}^{\min}\varepsilon$ and by Lemma 11.5, $U_{\mathcal{P}} \le F_{0,1}^{\max}$. Now, by Lemma 11.7 and our assumption $\alpha_{-1} < E\varepsilon$,

$$p = \tilde{\varepsilon} - U_{\mathcal{P}} S_{\alpha,\beta}(h) \ge L^{-4} F_{0,1}^{\min} \varepsilon - F_{0,1}^{\max} L_{0,1}^2 L_{0,d}(\alpha_{-1} + \beta_{-1} \gamma \frac{E\varepsilon - \alpha_{-1}}{\beta_{-1}})$$

= $L^{-4} F_{0,1}^{\min} \varepsilon - E^{-1} L^{-4} F_{0,1}^{\min}(\alpha_{-1} + \gamma (E\varepsilon - \alpha_{-1}))$
= $E^{-1} L^{-4} F_{0,1}^{\min}(E\varepsilon - \alpha_{-1} - \gamma (E\varepsilon - \alpha_{-1}))$
= $E^{-1} L^{-4} F_{0,1}^{\min}(1 - \gamma)(E\varepsilon - \alpha_{-1}) > 0.$

Next, in order to prove (11.2), note that K was chosen so that $\tilde{\varepsilon}' \geq \frac{\varepsilon'}{K}$. Now, by the definition of N, Lemma 11.8, Lemma 11.4(iv) and (11.3), we have

$$\begin{split} p' &= \tilde{\varepsilon}' - U_{\mathcal{P}'} S'_{\alpha',\beta'}(h') - q \frac{F_{0,2,d}^{\min} F_{1,d}^{\min}}{F_{0,1,d}^{\max} F_{2,d}^{\max}} \varepsilon'_0 T'_0(h') \\ &\geq \frac{\varepsilon'}{K} - N \left(\left(\alpha_0 + \beta_0 h_0 \right) + \left(\alpha_{||} + \beta_{||} h_{||} \right) + \frac{\alpha_{-1} + \beta_{-1} h_{-1}}{h_0} \right) - q L_{0,d} \frac{F_{0,2}^{\max}}{F_{0,1}^{\min}} \frac{L^6}{h_0} \varepsilon' \\ &\geq \frac{\varepsilon'}{K} - N E' \varepsilon' - \frac{\gamma h_0}{2VK} L_{0,d} \frac{F_{0,2}^{\max}}{F_{0,1}^{\min}} \frac{L^6}{h_0} \varepsilon' \\ &\geq \frac{\varepsilon'}{K} - \frac{\varepsilon'}{2K} - \frac{\gamma}{2K} \varepsilon' = (1 - \gamma) \frac{\varepsilon'}{2K} > 0. \end{split}$$

This completes the verification of the assumptions of Theorem 11.2.

Therefore, the assertions of Theorem 11.2 hold for our (X, w, \mathcal{F}) . Thanks to Lemmas 11.8 and 11.7, this means that

$$\operatorname{cse}_{k}(X,\mathcal{F}) \geq \min\left\{\frac{1}{T_{-1}'(h')}, \frac{q}{C}\right\} \geq \min\left\{\frac{1}{L^{12}(h_{0}^{-1}h_{-1}^{-1} + h_{||}^{-1})}, \frac{\gamma h_{0}}{2VKC}\right\}$$
$$\geq \min\left\{\frac{1}{L^{12}}, \frac{\gamma}{2VKC}\right\} \frac{1}{h_{0}^{-1}h_{-1}^{-1} + h_{||}^{-1}} \geq \frac{E''}{h_{0}^{-1}h_{-1}^{-1} + h_{||}^{-1}},$$
$$\operatorname{ccd}_{k}(X,\mathcal{F}) \geq T_{-1}(h)^{-1} = \frac{h}{L_{0,d}^{3}} \geq E''' \gamma \frac{E\varepsilon - \alpha_{-1}}{\beta_{-1}}.$$

Moreover, for every $f \in C^0$ with $\operatorname{dist}(f, Z^0) < \frac{\gamma D}{(h_0^{-1}h_{-1}^{-1}+h_{||}^{-1})} \leq \frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max}F_{0,d}^{\max}}T'_{-1}(h')^{-1}$, applying Algorithm 8.2 to f with the parameter being $D'h_0 \leq \frac{\gamma h_0}{2VK} = q$ results in $f' \in Z^k$ such that $\operatorname{dist}(f, f') \leq \frac{C}{q} \frac{F_{0,1,d}^{\max}F_{0,d}^{\max}}{F_{1,d}^{\min}} \operatorname{dist}(f, Z^k) \leq D''^{-1}h_0^{-1}\operatorname{dist}(f, Z^k)$. By Letting γ approach 1, we obtain the required bounds.

We finish with explaining the values listed in Table 1. For the values in the first row of the table, we simply substitute $K = L^5 F$, $N = 2F^4 L^{13}$, $V = L^7 F$ and C = L in the upper bounds for D, D'D,'', E, E', E'', E''' and replace $\frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max}F_{0,d}^{\max}}$ and $L_{i,d}^{-1}$ with the smaller quantities $\frac{1}{L^3 F}$. and L^{-1} , respectively.

respectively. Suppose now that X is lower regular. Then L = 1, $F_{0,1}^{\max} \leq F_{0,2}^{\max} = F_{0,2}^{\min}$ and $F_{0,1}^{\max} \leq F_{0,d}^{\max} = F_{0,d}^{\min}$. From this it follows that $K \leq 1$, $N \leq U_{\mathcal{P}_1} \max\{F_{0,2}^{\max}, F_{1,2}^{\max}\} \leq 2F^2$. Substituting K = 1, $N = 2F^2$, V = F, C = 1 and noting that $\frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max}F_{0,d}^{\max}} = \frac{1}{F_{0,1}^{\max}} \geq \frac{1}{F}$ gives the second row of the table. Finally, assume that $X(\leq 2)$ is an *m*-gon complex. Then

$$\begin{split} &K \leq L^5 \max\left\{\frac{F_{0,1}^{\max}}{F_{0,2}^{\min}}, \frac{1}{F_{1,2}^{\min}}\right\} = L^5 \max\left\{\frac{2}{m}, \frac{1}{m}\right\} = \frac{2L^5}{m} \\ &N = U_{\mathcal{P}_1} \max\left\{L^6 \frac{F_{0,2}^{\max}(F_{0,1}^{\max})^2}{(F_{0,d}^{\min})^2}, L^{13} F_{1,2}^{\max}\right\} \leq m \max\left\{L^6 \frac{m2^2}{m^2}, L^{13} m\right\} = L^{13} m^2 \\ &V = L^7 \frac{F_{0,1}^{\max}}{F_{0,1}^{\min}} = \frac{L^7 m}{2}. \end{split}$$

We get the last row of the table by substituting the right hand sides in the upper bounds for D, D'D, ", E, E', E'', E''' as well as replacing C with L and $\frac{F_{1,d}^{\min}}{F_{0,1,d}^{\max}F_{0,d}^{\max}}$ with $\frac{1}{L^3F_{0,1}^{\max}} = \frac{1}{2L^3}$. The third (resp. fourth) row then follow by taking m = 3 (resp. m = 4) and L = 1.

12 Proof of Theorem 11.2

Throughout this section, R is a commutative ring, (X, w) is a properly weighted R-oriented d-poset and \mathcal{F} is an R-sheaf on X. When there is no risk of confusion, we shall write $C^k = Z^k(X, \mathcal{F})$, $Z^k = Z^k(X, \mathcal{F})$ and $B^k = B^k(X, \mathcal{F})$.

12.1 Mock Locally Minimal Cochains

Definition 12.1 (Mock Locally q-Minimal Cochain). Let $q \in [0,1]$ be as before and let $k \in \{0,\ldots,d\}$. Given $u \in X$ with $i := \dim u \leq k$, a k-cochain $f \in C^k(X,\mathcal{F})$ is called mock q-locally minimal at u if for all $b \in B^{k-i-1}(X_u,\mathcal{F}_u)$, we have

$$||f|| \le ||f + b^u|| + q \cdot w(u).$$

We say that f is mock q-locally minimal if it is mock q-locally minimal at every face u with $0 \le \dim u < k$. A mock locally minimal cochain is a mock 0-locally minimal cochain.

Remark 12.2. (i) Following [EK17] and other sources, it is natural to call a k-cochain $f \in C^k(X, \mathcal{F})$ *q-locally minimal* at $u \in X(i)$ if

$$||f_u|| \le ||f_u + b|| + q$$

for every $b \in B^{k-i-1}(X_u, \mathcal{F}_u)$. When X is lower-regular, this is equivalent to f being mock q'-locally minimal at u for $q' = \frac{F_{i,k,d}F_{i,d}}{F_{k,d}}q$ (use Lemma 4.17 and Corollary 4.18). In general, however, there is no relation between being q-locally minimal being mock q-locally minimal; this is why we use the word "mock" in Definition 12.1. It will be important to use *mock* locally q-minimal cochains, rather than q-minimal cochains, in the proof of Proposition 12.5 below.

(ii) Every $f \in C^k(X, \mathcal{F})$ is mock (0-)locally minimal at every $u \in X(k)$, because $B^{-1}(X_u, \mathcal{F}_u) = 0$. Consequently, every 0-cochain is mock locally minimal.

(iii) One can introduce a coefficient depending on dim u before the factor qw(u) in Definition 12.1. This has no effect beyond altering the constants in Theorem 8.1 and its more technical versions. We did not attempt to look for coefficients that would give better constants.

Algorithm 12.3 (Algorithm for Making a Cochain Mock q-Locally Minimal). Let $k \in \{-1, \ldots, d-1\}$. The algorithm takes as input $h \in C^{k+1}(X, \mathcal{F})$ and some $q \in [0, 1]$ and outputs $g \in C^k(X)$ such that h + dg is mock q-locally minimal. The procedure is as follows:

- (1) Set $g_0 = 0 \in C^k(X, \mathcal{F})$ and i = 0.
- (2) While $h_i := h + dg_i$ is not mock *q*-locally minimal:
 - (a) Choose some $u \in \bigcup_{i=0}^{k} X(i)$ such that h_i is not mock q-locally minimal at u.
 - (b) Find some $g'_i \in C^{k-1}(X_u, \mathcal{F}_u)$ such that $h_i + dg'^u_i$ is mock q-locally minimal at u.
 - (c) Set $g_{i+1} = g_i + g'^u_i$ and increase *i* by 1.
- (3) Return g_i .

Remark 12.4. If, in Algorithm 12.3, we would take h = df for $f \in C^k$ and return $f + g_i$ instead of g_i , then we would get Algorithm 8.2.

Proposition 12.5. Let $k \in \{-1, \ldots, d-1\}$, $h \in C^{k+1}(X, \mathcal{F})$ and $q \in [0, 1]$. Suppose that Algorithm 12.3 is applied to h and q and let g be its output (assuming it stops). Then:

- (i) The algorithm stops. If q > 0 and there is $M \in \mathbb{R}$ such that $w(x) \leq Mw(y)$ for all $x, y \in X(1) \cup \cdots \cup X(k)$, then the loop (2) is executed at most $Mq^{-1}|X(k)| \cdot ||h||$ times.
- (ii) h + dg is mock q-locally minimal and $||h + dg|| \le ||h||$.

(*iii*)
$$||g|| \le \max\left\{\frac{F_{i,k,d}^{\max}F_{i,d}^{\max}}{F_{k,d}^{\min}} \mid i \in \{0, \dots, k\}\right\} q^{-1} ||h||$$

Proof. Let u_i denote the face u chosen at the *i*-th iteration of the loop (2) and let $r_i = \dim u_i$.

(i) By the definition of mock q-local minimality, we have $||h_{i+1}|| < ||h_i|| - qw(u_i)$. In particular, $||h_i|| > ||h_{i+1}||$ for all *i*. Since X is a finite, $||\cdot|| : C^{k+1}(X, \mathcal{F}) \to \mathbb{R}$ attains only finitely many values and so the algorithm must stop.

Suppose now that q > 0 and $w(x) \le Mw(y)$ for all $x, y \in X(1) \cup \cdots \cup X(k)$. There is some $x \in X(k)$ with $w(x) \le \frac{1}{|X(k)|}$. Then $||h_{i+1}|| < ||h_i|| - qw(u_i) \le ||h_i|| - qM^{-1}w(x) \le ||h_i|| - qM^{-1}|X(k)|^{-1}$. By iterating this, we see that $||h_i|| < ||h|| - \frac{qi}{|X(k)|M}$. Since $||h_i|| \ge 0$, this means that the loop (2) is executed at most $Mq^{-1}|X(k)| \cdot ||h||$ times.

(ii) The first claim is immediate from the stopping condition of the loop ((2)). The second claim follows from our earlier observation that $||h + dg_i|| = ||h_i|| > ||h_{i+1}|| = ||h + dg_{i+1}||$ whenever both sides are defined.

(iii) Let n be the value of i when the algorithm stops. Recall that $||h_i|| - ||h_{i+1}|| \ge qw(u_i)$. Using this, the definition of g'_i and Lemma 4.15 we see that

$$\|g_i'^{u_i}\| \le w(X(k)_{u_i}) \le \frac{F_{r_i,k,d}^{\max}F_{r_i,d}^{\max}}{F_{k,d}^{\min}}w(u_i) \le \frac{F_{r_i,k,d}^{\max}F_{r_i,d}^{\max}}{F_{k,d}^{\min}}q^{-1}(\|h_i\| - \|h_{i+1}\|).$$

This means that

$$||g|| \leq \sum_{i=1}^{n} ||g_{i}'^{u_{i}}|| \leq \sum_{i=1}^{n} \frac{F_{r_{i},k,d}^{\max} F_{r_{i},d}^{\max}}{F_{k,d}^{\min}} q^{-1}(||h_{i+1}|| - ||h_{i}||)$$
$$\leq \max\left\{\frac{F_{i,k,d}^{\max} F_{i,d}^{\max}}{F_{k,d}^{\min}} \mid i \in \{0,\dots,k\}\right\} q^{-1}(||h_{0}|| - ||h_{n}||)$$

and (iii) follows because $h_0 = h$.

Corollary 12.6. Let $k \ge \{0, \ldots, d-1\}$ and $h \in C^k(X, \mathcal{F})$. Then there exists $g \in C^{k-1}(X, \mathcal{F})$ such that h + dg is mock locally minimal and $||h + dg|| \le ||h||$.

Proof. Apply Algorithm 12.3 to h with q = 0. The algorithm stops by 12.5(i) and its output is the required g.

12.2 Reduction to Expansion of Small Locally Minimal Cochains

Let $\gamma \in [0,1]$. We call a cochain $f \in C^k(X, \mathcal{F})$ γ -small if $||f|| < \gamma$. Given a subset $S \subseteq C^k(X, \mathcal{F})$ and $\beta \in [0, \infty)$, we say that \mathcal{F} β -expands S (or cochains in S) if $||df|| \ge \beta ||f||$ for every $f \in S$. It turns out that if \mathcal{F} expands small mock locally minimal k-cochains and (k + 1)-cochains, then \mathcal{F} has good cosystolic expansion in dimension k:

Proposition 12.7. Let X, \mathcal{F} , d be as in the beginning of the section, let $k \in \{0, \ldots, d-2\}$ and $C = \max\left\{\frac{F_{i,k,d}^{\max}F_{i,d}^{\max}}{F_{k,d}^{\min}} \mid i \in \{0, \ldots, k\}\right\}$. Let $\beta, \beta' \in [0, \infty), \ \gamma, \gamma', q \in [0, 1]$, and suppose that

- (1) $\mathcal{F} \beta$ -expands γ -small mock locally minimal k-cochains, and
- (2) $\mathcal{F} \beta'$ -expands γ' -small q-mock locally minimal (k+1)-cochains.

Then

$$\operatorname{ccd}_k(X,\mathcal{F}) \ge \gamma$$
 and $\operatorname{cse}_k(X,\mathcal{F}) \ge \min\{\gamma',\frac{q}{C}\}$

The assertion about $\operatorname{ccd}_k(X,\mathcal{F})$ holds even without assuming (2). Moreover, if $f \in C^k(X,\mathcal{F})$ satisfies $\operatorname{dist}(f, Z^k(X,\mathcal{F})) < \frac{F_{k+1,d}^{\min}}{F_{k,k+1,d}^{\max}F_{k,d}^{\max}}\gamma'$, then Algorithm 12.3, applied to to h := df and q returns $g \in C^k(X,\mathcal{F})$ such that $f+g \in Z^k(X,\mathcal{F})$ and $\operatorname{dist}(f, f+g) = \|g\| < \frac{C}{q} \frac{F_{k,k+1,d}^{\max}F_{k,d}^{\max}}{F_{k+1,d}^{\min}} \operatorname{dist}(f, Z^k(X,\mathcal{F})).$

Proof. Let $f \in Z^k$ and suppose that $||f|| < \gamma$. By Corollary 12.6, there is $g \in B^{k-1}$ such that f + dg is mock locally minimal and $||f + dg|| \le ||f|| < \gamma$. By assumption (1), $0 = ||d(f + dg)|| \ge \beta ||f + dg||$. Thus f + dg = 0 and in particular, $f \in B^k$.

Next, let $f \in C^k$. We need to show that $||df|| \ge \min\{\gamma', Cq\} \operatorname{dist}(f, Z^k)$. If $||df|| \ge \gamma'$, then this holds automatically, so assume $||df|| < \gamma'$. By applying Proposition 12.5 with h = df and the parameter q, we see that there is $g \in C^k$ such that df + dg is mock q-locally minimal, $||df + dg|| \le$ $||df|| < \gamma'$ and $||g|| \le Cq^{-1}||df||$. By assumption (2), $0 = ||d(df + dg)|| \ge \beta' ||df + dg||$, so df + dg = 0. This means that $f + g \in Z^k$. As a result, $\operatorname{dist}(f, Z^k) \le ||g|| \le Cq^{-1}||df||$. By rearranging, we get $||df|| \ge \frac{q}{C} \operatorname{dist}(f, Z^k) \ge \min\{\frac{q}{C}, \gamma'\} \operatorname{dist}(f, Z^k)$.

To finish, suppose that $f \in C^k$ satisfies $\operatorname{dist}(f, Z^k) < \frac{F_{k+1,d}^{\min}}{F_{k,k+1,d}^{\max}F_{k,d}^{\max}}\gamma'$. Choose some $g' \in Z^k$ which minimizes ||f - g'||. By Lemma 4.15,

$$\begin{aligned} \|df\| &= \|d(f - g')\| \le \sum_{x \in \text{supp}(f - g')} w(X(k+1)_x) \le \\ &\le \frac{F_{k,k+1,d}^{\max} F_{k,d}^{\max}}{F_{k+1,d}^{\min}} w(\text{supp}(f - g')) = \frac{F_{k,k+1,d}^{\max} F_{k,d}^{\max}}{F_{k+1,d}^{\min}} \operatorname{dist}(f, Z^k). \end{aligned}$$

Thus, $\|df\| < \gamma'$. Let g be the output of Algorithm 12.3 applied to df. Then, as in the last paragraph, it follows that $f + g \in Z^k$ and $\operatorname{dist}(f, f + g) = \|g\| \le Cq^{-1} \|df\| < \frac{C}{q} \frac{F_{k,k+1,d}^{\max}F_{k,d}^{\max}}{F_{k+1,d}^{\min}} \operatorname{dist}(f, Z^k)$. \Box

Using Proposition 12.7, we can reduce Theorem 11.2 into proving the following theorem. Recall that, given a k-intersection profile \mathcal{P} for X, and lists $\alpha = \{\alpha_{\rho}\}_{\rho \in \mathcal{P}}, \{\beta_{\rho}\}_{\rho \in \mathcal{P}}$ of non-negative real numbers, we defined in §11.1 Laurent polynomials

$$T_k, T_{k-1}, \ldots, T_{-1}, S_{\alpha,\beta} \in \mathbb{R}[x_{\rho}^{\pm 1} \mid \rho \in \mathcal{P}^{\mathrm{red}}]$$

and a natural number $U_{\mathcal{P}} \in \mathbb{N}$.

Theorem 12.8. Let R be a commutative ring, let (X, w) be a properly weighted R-oriented d-poset, let \mathcal{F} be an R-sheaf on X, let $k \in \{0, \ldots, d-1\}$, and let \mathcal{P} be a k-intersection profile for X. Let $\{\varepsilon_i\}_{i=0}^k, \alpha = \{\alpha_\rho\}_{\rho \in \mathcal{P}}, \beta = \{\beta_\rho\}_{\rho \in \mathcal{P}}$ be lists of non-negative real numbers. Put

$$\tilde{\varepsilon} = \min\left\{\frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}\varepsilon_i \,|\, i \in \{0,\dots,k\}\right\}$$

and suppose that exist $h = \{h_{\rho}\}_{\rho \in \mathcal{P}} \in (0, 1]^{\mathcal{P}}$ and $q \in [0, 1]$ such that

$$p := \tilde{\varepsilon} - U_{\mathcal{P}} S_{\alpha,\beta}(h) - q \sum_{i=0}^{k-1} \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} \varepsilon_i T_i(h) > 0.$$

Suppose further that the following conditions are met:

- (1) $(X_u, w_u, \mathcal{F}_u)$ is an $\varepsilon_{\dim u}$ -coboundary expander in dimension $k \dim u 1$ for every $u \in X(0) \cup \cdots \cup X(k)$;
- (2) $\operatorname{NIH}_{u}^{\ell,r,t}(X)$ is an $(\alpha_{\rho},\beta_{\rho})$ -skeleton expander for every $\rho = (t,\ell,r,b) \in \mathcal{P}^{\operatorname{red}}$ and $u \in X(b)$.

Then \mathcal{F} p-expands $T_{-1}(h)^{-1}$ -small mock q-locally minimal k-cochains. That is, for every mock q-locally minimal $f \in C^k(X, \mathcal{F})$ with $||f|| < T_{-1}(h)^{-1}$, we have $||df|| \ge p||f||$.

Proof of Theorem 11.2 assuming Theorem 12.8. We use the notation of Theorem 11.2. By Proposition 12.7 and Remark 12.4, it is enough to show that \mathcal{F} (a) *p*-expands $T_{-1}(h)$ -small mock locally minimal *k*-cochains and (b) *p'*-expands $T'_{-1}(h)$ -small mock *q*-locally minimal (k + 1)-cochains. To that end, we apply Theorem 12.8 twice: once for $X, \mathcal{F}, k, \mathcal{P}, \alpha, \beta, h$ with *q* being 0, and again for $X, \mathcal{F}, k+1, \mathcal{P}', \alpha', \beta', h'$ with *q* from Theorem 11.2. This gives (a) and (b), respectively.

We will prove Theorem 12.8 in the next section. Before turning to that, we note that Theorem 12.8 also gives a criterion for bounding the cocycle distance from below:

Corollary 12.9. Keep the notation of Theorem 12.8. Suppose that assumptions (1) and (2) of that theorem hold and there is $h = \{h_{\rho}\}_{\rho \in \mathcal{P}}$ in $(0,1]^{\mathcal{P}}$ such that $\tilde{\varepsilon} > \mathcal{U}_P S_{\alpha,\beta}(h)$. Then $\operatorname{ccd}_k(X, w, \mathcal{F}) \geq T_{-1}(h)^{-1}$.

Proof. By Theorem 12.8, applied with q = 0, the sheaf \mathcal{F} expands $T_{-1}(h)^{-1}$ -small mock locally minimal k-cochains. The lower bound on $\operatorname{ccd}_k(X, \mathcal{F})$ now follows from Proposition 12.7.

We use Corollary 12.9 to prove Theorem 8.13 from earlier.

Proof of Theorem 8.13. In short, this is just unfolding Corollary 12.9 in the case where k = 0 and \mathcal{P} is $\mathcal{P}^{(0)}$ from Example 7.15(i).

Recall that we are given an *R*-sheaf on a properly weighted *d*-poset (X, w). It is further given that assumptions (1) and (2) of Theorem 12.8 hold for k = 0, $\varepsilon_1 = \varepsilon$, $\alpha_{(1,0,0,-1)} = \alpha$, $\beta_{(1,0,0,-1)} = \beta$. We choose the constants *E* and *E'''* as in the proof of Theorem 8.10 (given in §11.2), i.e.,

$$E = L_{0,d}^{-1} L_{1,d}^{-3} L^{-4}, \qquad E''' = L_{0,d}^{-3}$$

Let $\gamma \in [\frac{1}{2}, 1)$ and $h = \gamma \frac{E_{\varepsilon} - \alpha}{\beta}$, and let p be as in Theorem 12.8 (with k = 0 and $\mathcal{P} = \mathcal{P}^{(0)}$). Then, as in the proof of Theorem 8.10, we have p > 0. By Lemma 11.7, $T_{-1}(h) = \frac{1}{E'''h}$. We may therefore apply Corollary 12.9 and assert that $\operatorname{ccd}_k(X, \mathcal{F}) \geq E'''h = E'''\gamma \frac{E_{\varepsilon} - \alpha}{\beta}$. As this holds for all $\gamma \in [\frac{1}{2}, 1)$, we are done.

13 Proof of Theorem 12.8

Throughout, R is a commutative ring, (X, w) is a properly weighted R-oriented d-poset and \mathcal{F} is an R-sheaf on X. If not indicated otherwise, $k \in \{0, \ldots, d-1\}$ and \mathcal{P} is a k-intersection profile for X. In fact, the assumption that X is R-oriented may be relaxed to assuming that the subposet $X(k-1) \cup X(k) \cup X(k+1)$ is R-oriented.

Recall that given $A \subseteq X$ and $z \in X$, we write A_z for $\{x \in A : z \ge x\}$.

13.1 Heavy Faces

For this section, we fix $k \in \{0, \ldots, d-1\}$, a k-cochain $f \in C^k(X, \mathcal{F})$, a k-intersection profile \mathcal{P} for X, and choose real numbers $h \in (h_\rho)_{\rho \in \mathcal{P}_{red}}$ in the interval (0, 1].

Following [EK17], for every $i \in \{-1, \ldots, k\}$, we define a set of *i*-faces $A_i = A_i(f, h, \mathcal{P})$ by decreasing induction on *i* as follows:

- $A_k = \operatorname{supp}(f).$
- Assuming A_{i+1}, \ldots, A_k were defined, define A_i to be the set of face $u \in X(i)$ such that for some $\rho = (t, \ell, r, b) \in \mathcal{P}$ with b = i, we have

$$w_z(A_{\ell,z}) + w_z(A_{r,z}) \ge 2h_\rho.$$
(13.1)

(When $r = \ell$, this simplifies to $w_z(A_{\ell,z}) \ge h_{\rho}$.)

Elements of $A_{-1} \cup \cdots \cup A_k$ will be called (f, h, \mathcal{P}) -heavy, or just heavy for short. Informally, a face is heavy if it is in $\operatorname{supp}(f)$, or it is contained in relatively many heavy faces of larger dimension.

Recall from §11.1 that we defined Laurent polynomials $T_k, \ldots, T_{-1} \in \mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ inductively by setting $T_k = 1$ and

$$T_i = \sum_{\rho \in \mathcal{P}: \text{hgt } \rho = i} x_{\rho}^{-1} \left[c_{\rho} T_{\ell(\rho)} + c_{\rho}' T_{r(\rho)} \right]$$

where

$$c_{(t,\ell,r,i)} = \frac{\frac{\frac{F_{\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{i,\ell,d}^{\min}}}{\frac{F_{\ell,d}^{\min}F_{i,\ell,d}^{\min}}{F_{\ell,d}^{\min}} + \frac{F_{r,d}^{\min}F_{i,r,d}^{\min}F_{r,d}^{\min}}{F_{r,d}^{\max}}} \quad \text{and} \quad c_{(t,\ell,r,i)}' = c_{(t,r,\ell,i)}$$

We now show that the weight of the (f, h, \mathcal{P}) -heavy *i*-faces, $w(A_i)$, is at most proportional to ||f||. Lemma 13.1. Let h and f be as above and suppose $-1 \le i \le k$. Then

$$w(A_i) \le T_i(h) \|f\|.$$

Proof. Let z be a heavy *i*-face and let $\rho = (t, \ell, r, i) \in \mathcal{P}$. By Lemmas 4.17 and 4.15, we have

$$2h_{\rho} = h_{\rho}(w_{z}(X(\ell)_{z}) + w_{z}(X(r)_{z})) \ge h_{\rho}\frac{|X(d)|}{|X(d)_{z}|} \left[\frac{F_{\ell,d}^{\min}}{F_{i,\ell,d}^{\max}}w(X(\ell)_{z}) + \frac{F_{r,d}^{\min}}{F_{i,r,d}^{\max}}w(X(r)_{z})\right] \ge h_{\rho}\frac{|X(d)|}{|X(d)_{z}|} \left[\frac{F_{\ell,d}^{\min}}{F_{i,\ell,d}^{\max}} \cdot \frac{F_{i,\ell,d}^{\min}F_{i,d}^{\min}}{F_{\ell,d}^{\max}} + \frac{F_{r,d}^{\min}}{F_{i,r,d}^{\max}} \cdot \frac{F_{i,r,d}^{\min}F_{i,d}^{\min}}{F_{r,d}^{\max}}\right]w(z),$$

whereas on the other hand,

$$2h_{\rho} \le w_{z}(A_{\ell,z}) + w_{z}(A_{r,z}) \le \frac{|X(d)|}{|X(d)_{z}|} \left[\frac{F_{\ell,d}^{\max}}{F_{i,\ell,d}^{\min}} w(A_{\ell,z}) + \frac{F_{r,d}^{\max}}{F_{i,r,d}^{\min}} w(A_{r,z}) \right].$$

Together, we get

$$w(z) \leq h_{\rho}^{-1} \left[\frac{\frac{F_{\ell,d}^{\max}}{F_{r,\ell,d}^{\min}} w(A_{\ell,z})}{\frac{F_{\ell,d}^{\min}}{F_{r,d}^{\max}} \frac{F_{r,d}^{\min}F_{r,d}^{\min}}{F_{r,d}^{\min}} + \frac{F_{r,d}^{\min}}{F_{r,d}^{\min}} \frac{F_{r,d}^{\min}F_{r,d}^{\min}}{F_{r,d}^{\max}} + \frac{F_{r,d}^{\min}F_{r,d}^{\min}}{F_{r,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}}{F_{r,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}}{F_{r,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\min}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\max}F_{\ell,d}^{\max}}{F_{\ell,d}^{\max}} + \frac{F_{\ell,d}^{\max}F_$$

We now prove the lemma by decreasing induction on i. The case i = k is clear because $w(A_k) = w(\operatorname{supp} f) = T_k(h) ||f||$. Suppose now that i < k and the lemma was established for larger values of i. Then by what we have shown and Lemma 4.16,

$$w(A_{i}) = \sum_{z \in A_{i}} w(z) \leq \sum_{z \in A_{i}} h_{\rho}^{-1} \left[\frac{c_{\rho}}{F_{i,\ell(\rho)}^{\max}} w(A_{\ell(\rho),z}) + \frac{c_{\rho}'}{F_{i,r(\rho)}^{\max}} w(A_{r(\rho),z}) \right]$$
$$\leq \sum_{z \in X(i)} h_{\rho}^{-1} \left[\frac{c_{\rho}}{F_{i,\ell(\rho)}^{\max}} w(A_{\ell(\rho),z}) + \frac{c_{\rho}'}{F_{i,r(\rho)}^{\max}} w(A_{r(\rho),z}) \right]$$
$$\leq h_{\rho}^{-1} (c_{\rho} w(A_{\ell}) + c_{\rho}' w(A_{r})) = T_{i}(h).$$

13.2 Exceptional Faces

We continue to assume that $k \in \{0, \ldots, d-1\}$, $f \in C^k(X, \mathcal{F})$, \mathcal{P} is a k-intersection profile \mathcal{P} for X, and $h = (h_\rho)_{\rho \in \mathcal{P}} \in (0, 1]^{\mathcal{P}}$.

Let $\rho = (t, \ell, r, b) \in \mathcal{P}$. By a ρ -square (in X), we mean a quadruple (x, y, z, u) in X of respective dimensions (t, ℓ, r, b) such that $y, z \leq x$ and $u \in \text{Inf}\{y, z\}$. A ρ -square (x, y, z, u) is called (f, h, \mathcal{P}) exceptional, or just exceptional for short, if y and z are heavy, but u is not. A (k+1)-face $s \in X(k+1)$ is called (f, h, \mathcal{P}) -exceptional if there is an exceptional square (x, y, z, u) with $x \leq s$. The set of exceptional (k + 1)-face will be denoted by

$$\Upsilon = \Upsilon(f, h, \mathcal{P}).$$

Non-exceptional faces $s \in X(k+1) - \Upsilon$ have the property that if (x, y, z, u) is a ρ -square $(\rho \in \mathcal{P})$ with $x \leq s$, and if y and z are heavy, then u is also heavy.

We will show that the weight of exceptional (k+1)-faces is at most proportional to ||f||, provided that the non-intersection hypergraphs of the links of X are good skeleton expanders (Section 7). Recall that given a $\beta \in [0, \infty)^{\mathcal{P}}$ we defined $S = C \mathbb{P}[\pi^{\pm 1} \mid \alpha \in \mathcal{P}]$ by

Recall that given $\alpha, \beta \in [0, \infty)^{\mathcal{P}}$, we defined $S_{\alpha, \beta} \in \mathbb{R}[x_{\rho}^{\pm 1} | \rho \in \mathcal{P}]$ by

$$S_{\alpha,\beta} = \sum_{\rho=(t,\ell,r,i)\in\mathcal{P}} \frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}} \left[\frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}}T_{\ell} + \frac{F_{i,r,d}^{\max}F_{i,r}^{\max}}{F_{r,d}^{\min}}T_{r}\right] (\alpha_{\rho} + \beta_{\rho}x_{\rho}).$$

Lemma 13.2. With notation as before, let $\alpha, \beta, h \in [0, \infty)^{\mathcal{P}}$. Suppose that for every $\rho = (t, \ell, r, b) \in \mathcal{P}$ and $u \in X(b)$, the non-intersecting hypergraph $\operatorname{NIH}_{u}^{\ell,r,t}(X)$ (see Notation 7.9) is an $(\alpha_{\rho}, \beta_{\rho})$ -skeleton expander. Then

$$w(\Upsilon) \le S_{\alpha,\beta}(h) \|f\|.$$

Proof. We first note that a (k+1)-face $s \in X(k+1)$ is exceptional if and only if there is $\rho \in \mathcal{P}$ and an exceptional ρ -square (x, y, z, u) with $x \leq s$. The "if" part is clear. For the converse, the fact that s is exceptional means that there is $\rho = (t, \ell, r, b) \in \mathcal{P}$ (but maybe $\rho \notin \mathcal{P}$) and an exceptional ρ -square (x, y, z, u) such that $x \leq s$. Since y and z are heavy and u is not, we have $u \neq y, z$, which means that $b < \ell, r$ and $\ell, r < t$. If $\ell \geq r$, then $\rho \in \mathcal{P}$, and otherwise, $\rho' := (t, r, \ell, b) \in \mathcal{P}$ and (x, z, y, u) is an exceptional ρ' -square with $x \leq s$. This proves our claim.

Fix some $\rho = (t, \ell, r, i) \in \mathcal{P}$ and $u \in X(i) - A_i$. Since u is not heavy, we have

$$w_u(A_{\ell,u}) + w_u(A_{r,u}) < 2h_\rho.$$

We claim that

$$w_u(E_2(A_{\ell,u} \cup A_{r,u})) < \alpha_\rho \frac{w_u(A_{\ell,u}) + w_u(A_{r,u})}{2} + \beta_\rho \left(\frac{w_u(A_{\ell,u}) + w_u(A_{r,u})}{2}\right)^2.$$

Indeed, if $\ell = r$, then this holds because $A_{\ell,u} = A_{r,u}$ and $H := \operatorname{NIH}_{u}^{\ell,r,t}(X)$ is an $(\alpha_{\rho}, \beta_{\rho})$ -skeleton expander, and if $\ell \neq r$, then $w_H(A_{\ell,u} \cup A_{r,u}) = \frac{1}{2}(w_z(A_{\ell,u}) + w_r(A_{r,u}))$ while $w_H(E^2(A_{\ell,u} \cup A_{r,u})) = w_u(E^2(A_{\ell,u} \cup A_{r,u}))$, and again we reach the same conclude using the skeleton expansion of H. Since $w(A_{\ell,u}) + w_u(A_{r,u}) < 2h_{\rho}$, the inequality implies that

$$w_u(E^2(A_{\ell,u} \cup A_{r,u})) < \frac{1}{2}(\alpha_\rho + \beta_\rho h_\rho)(w_u(A_{\ell,u}) + w_u(A_{r,u})),$$

and by applying Lemma 4.17 to both sides, we get that

$$w(E^{2}(A_{\ell,u}, A_{r,u})) < \frac{1}{2}(\alpha_{\rho} + \beta_{\rho}h_{\rho})\frac{F_{i,t,d}^{\max}}{F_{t,d}^{\min}}\left(\frac{F_{\ell,d}^{\max}}{F_{i,\ell,d}^{\min}}w(A_{\ell,u}) + \frac{F_{r,d}^{\max}}{F_{i,r,d}^{\min}}w(A_{r,u})\right)$$

Let us abbreviate $E_2(A_{\ell,u} \cup A_{r,u})$ to $E(u,\rho)$. Then $E(u,\rho)$ is precisely the set of faces $x \in X(t)$ for which there exists an exceptional ρ -square of the form (x, *, *, u). As a result,

$$w(\Upsilon) \le \sum_{\rho=(t,\ell,r,i)\in\mathcal{P}} \sum_{u\in X(i)-A_i} \sum_{x\in E(u,\rho)} w(X(k+1)_x).$$
(13.2)

Let $\rho = (t, \ell, r, i) \in \mathcal{P}^{\text{red}}$. By Lemma 4.15 and our upper bound on $w(E(u, \rho))$, we have

$$\begin{split} \sum_{u \in X(i) - A_i} \sum_{x \in E(u,\rho)} w(X(k+1)_x) &\leq \sum_{u \in X(i) - A_i} \sum_{x \in E(u,\rho)} \frac{F_{t,k+1,d}^{\max} F_{t,d}^{\max}}{F_{k+1,d}^{\min}} w(x) \\ &= \sum_{u \in X(i) - A_i} \frac{F_{t,k+1,d}^{\max} F_{t,d}^{\max}}{F_{k+1,d}^{\min}} w(E(u,\rho)) \\ &\leq \sum_{u \in X(i)} \frac{1}{2} (\alpha_{\rho} + \beta_{\rho} h_{\rho}) \frac{F_{t,k+1,d}^{\max} F_{t,d}^{\max}}{F_{k+1,d}^{\min}} \frac{F_{t,d}^{\max}}{F_{t,d}^{\min}} \left(\frac{F_{\ell,d}^{\max}}{F_{i,\ell,d}^{\min}} w(A_{\ell,u}) + \frac{F_{r,d}^{\max}}{F_{i,r,d}^{\min}} w(A_{r,u}) \right) \end{split}$$

By Lemma 4.16, the end result is at most

$$(\alpha_{\rho} + \beta_{\rho}h_{\rho})\frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}}\left(\frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}}w(A_{\ell}) + \frac{F_{i,r,d}^{\max}F_{i,r}^{\max}}{F_{r,d}^{\min}}w(A_{r})\right)$$

and by Lemma 13.1, this is bounded from above by

$$(\alpha_{\rho}+\beta_{\rho}h_{\rho})\frac{F_{t,k+1,d}^{\max}F_{t,d}^{\max}F_{i,t,d}^{\max}}{2F_{k+1,d}^{\min}F_{t,d}^{\min}}\left(\frac{F_{i,\ell,d}^{\max}F_{i,\ell}^{\max}}{F_{\ell,d}^{\min}}T_{\ell}(h)+\frac{F_{i,r,d}^{\max}F_{i,r}^{\max}}{F_{r,d}^{\min}}T_{r}(h)\right)\|f\|.$$

Plugging this into (13.2) gives the lemma.

13.3 Ladders and Terminal Faces

We continue to use the notation of \$13.2.

Recall from Definition 7.11 that \mathcal{P} is associated with a set of \mathcal{P} -admissible pairs Ad(calP). We will call a pair of faces $(x, y) \in X \times X$ \mathcal{P} -admissible if $x \ge y$ and $(\dim x, \dim y)$ is \mathcal{P} -admissible. Given $x, y \in X$, a $(\mathcal{P}$ -)ladder from x to y is a sequence of faces $x = x_0 \ge x_1 \ge \cdots \ge x_n = y$ such that (x_{i-1}, x_i) is \mathcal{P} -admissible or $x_{i-1} = x_i$ for all $i \in \{1, \ldots, n\}$. We say that the ladder $x = x_0 \ge x_1 \ge \cdots \ge x_n = y$ is $((f, h, \mathcal{P})$ -)heavy if the faces x_0, \ldots, x_n are heavy or have dimension k+1. If there exists a heavy ladder from x to y, we also say that $x ((f, h, \mathcal{P})$ -)descends to y. A face $x \in X$ is called $((f, h, \mathcal{P})$ -)terminal if it is heavy or (k+1)-dimensional and it does not descend to any of its proper subfaces.

It turns out that a non-exceptional (k + 1)-face descends to exactly one terminal face.

Lemma 13.3. With notation as above, let $x \in X(k+1) - \Upsilon$. Then there is exactly one terminal face u descended from x. Furthermore, any face y descended from x descends to u.

Proof. The second statement follows from the first because y descends to some terminal face u' which is descended from x and must therefore coincide with u. We turn to prove the first statement.

Since x descends to itself, it descends to some terminal subface, call it u. Suppose that x descends to another terminal face v. We need to prove that u = v. Let $x = x_0 \ge \cdots \ge x_s = u$ and $x = y_0 \ge \cdots \ge y_t = v$ be a heavy ladders from x to u and v, respectively.

We claim that there exist faces $u_{i,j}$ for all $i \in \{0, 1, \dots, s\}, j \in \{0, 1, \dots, t\}$ such that:

- (i) $u_{i,0} = x_i$ and $u_{0,j} = y_j$ for all i, j;
- (ii) for every $i \ge 1$, we have $u_{i-1,j} \ge u_{i,j}$, and if $u_{i-1,j} > u_{i,j}$, then $(u_{i-1,j}, u_{i,j})$ is \mathcal{P} -admissible;
- (iii) for every $j \ge 1$, we have $u_{i,j-1} \ge u_{i,j}$, and if $u_{i,j-1} > u_{i,j}$, then $(u_{i,j-1}, u_{i,j})$ is \mathcal{P} -admissible;
- (iv) $u_{i,j}$ is heavy for all i, j.

Indeed, if i = 0 or j = 0, the we define $u_{i,j}$ as in (i); conditions (ii)–(iv) hold in this case because $x_0 \ge \cdots \ge x_s$ and $y_0 \ge \cdots \ge y_t$ are heavy ladders. We construct the remaining $u_{i,j}$ by induction on i + j: Assuming $x := u_{i-1,j-1}$, $y := u_{i,j-1}$ and $z := u_{i-1,j}$ were defined in such a manner that (ii)–(iv) hold, choose $u_{i,j}$ to be some member b of $\inf\{y, z\}$. We need to show that (ii)–(iv) continue to hold. To that end, we split into cases.

If $b \notin \{y, z\}$, then we must have $x \notin \{y, z\}$. By the induction hypothesis, (x, y) and (x, z) are both \mathcal{P} -admissible. Since \mathcal{P} is a k-intersection profile for X, there is $\rho \in \mathcal{P}$ such that (x, y, z, b) or (x, z, y, b) is a ρ -square. This, means that $(u_{i,j-1}, u_{i,j} = (y, b)$ and $(u_{i-1,j}, u_{i,j}) = (z, b)$ are both \mathcal{P} -admissible, proving (ii) and (iii). Moreover, since y and z are heavy and x is not exceptional, b is also heavy.

Suppose next that b = y and $b \neq z$. Then (ii) and (iv) hold and $x \geq z \geq y = b$. If one of the last two inequalities is an equality, then (iii) holds by the induction hypothesis. Otherwise, (x, z) and (x, y) are both \mathcal{P} -admissible, hence $(z, y) = (u_{i,j-1}, u_{i,j})$ must also be \mathcal{P} -admissible and again (iii) holds.

The case where b = z and $b \neq y$ is handled similarly. In the remaining case where b = z = y, (ii)–(iv) follow readily from the induction hypothesis. This completes the proof of our claim.

To finish observe that $u = u_{s,0} \ge u_{s,1} \ge \cdots \ge u_{s,t}$ is a heavy ladder. Since u is terminal, we must have $u_{s,t} = u$. Similarly, $u_{s,t} = v$, and we conclude that u = v.

Our next goal is to relate the weight of the (k + 1)-faces descending to a given terminal face u with the weight of the k-faces (i.e., faces in supp f) which descend to u. To that end, we shall need the following general lemma:

Lemma 13.4. Let \mathcal{F} be a sheaf on a d-poset X, let $f, g \in C^k(X, \mathcal{F})$ $(0 \le k \le d)$, let $q \in [0, 1]$ and let $u \in X$. Suppose that $g(x) \in \{f(x), 0\}$ for every $x \in X(k)$. If f is mock q-locally minimal at u, then so is g.

Proof. Write $i = \dim u$ and let $b \in B^{k-i-1}(X_u, \mathcal{F}_u)$. Also let $Q = \frac{F_{k,d}^{\max}F_{i,d}^{\max}}{F_{i,k,d}^{\min}F_{i,d}^{\min}} \frac{|X(d)_u|}{|X(d)|}q$. We need to prove that $||g|| \le ||g+b^u|| + Q$. Our assumption on g means that ||f|| = ||g|| + ||f-g||. Furthermore, since f is mock q-locally minimal at u, we have $||f|| \le ||f+b^u|| + Q$. Together, we get that $||g+b^u|| = ||(f+b^u) - (f-g)|| \ge ||f+b^u|| - ||f-g|| = ||f+b^u|| - ||f|| + ||g|| \ge ||g|| - Q$, which is what we want.

We continue using the general notation introduced so far: X is an R-oriented d-poset, \mathcal{F} is an R-sheaf, $f \in C^k(X, \mathcal{F})$, etc.

Lemma 13.5. With $X, \mathcal{F}, \mathcal{P}, k, f, \Upsilon$ be as before and let $u \in X$ be a terminal face of dimension $i \in \{0, \ldots, k\}$. Define

$$D(u) := \{x \in X(k) : x \text{ descends to } u\},\$$

$$D'(u) := \{x \in X(k+1) : x \text{ descends to } u\}.$$

Suppose also that f is mock q-locally minimal at u and $(X_u, w_u, \mathcal{F}_u)$ is an ε -CBE in dimension k - i - 1 ($\varepsilon > 0$). Then

$$\varepsilon \frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}w(D(u)) \le w(D'(u) \cap [\operatorname{supp}(df) \cup \Upsilon]) + q \frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}\varepsilon w(u)$$

Proof. Define $g \in C^k(X, \mathcal{F})$ by

$$g(x) = \begin{cases} f(x) & x \in D(u) \\ 0 & x \notin D(u). \end{cases}$$

Lemma 13.4 tells us that g is mock q-locally minimal at u, because f is. In addition, ||g|| = w(D(u)). Step 1. We claim that $\operatorname{supp}(dg) \subseteq D'(u) \cap (\Upsilon \cup \operatorname{supp}(df))$.

Let $x \in \operatorname{supp}(dg)$. Then there is $y \in x(k)$ with $y \in \operatorname{supp} g$. By the definition of g, the face y descends to u, so $x \in D'(u)$. It remains to show that $x \in \Upsilon \cup \operatorname{supp}(df)$. Suppose that $x \notin \Upsilon$. We need to show that $x \in \operatorname{supp}(df)$. Let $y \in x(k) \cap \operatorname{supp}(f)$. Then y is heavy and so x descends to y. By Lemma 13.3 and our assumption that x is not exceptional, y descends to u. This means that g(y) = f(y). Since g(y) = 0 whenever f(y) = 0, we conclude that g and f agree on every $y \in x(k)$, so $df(x) = dg(x) \neq 0$, or rather, $x \in \operatorname{supp}(df)$.

Step 2. By Step 1, we have f(x) = 1

$$w(\operatorname{supp}(dg)) \le w(D'(u) \cap [\operatorname{supp}(df) \cup \Upsilon]).$$

On the other hand, by our assumption that (X_u, \mathcal{F}_u) is an ε -CBE in dimension $k - \dim u - 1$, Lemma 4.17 and our earlier observation that g is mock q-locally minimal at u we have:

$$\begin{split} w(\operatorname{supp}(dg)) &\geq w(X(d)_u) \frac{F_{i,k+1,d}^{\min}}{F_{k+1,d}^{\max}} w_u(\operatorname{supp}(dg_u)) \\ &\geq w(X(d)_u) \frac{F_{i,k+1,d}^{\min}}{F_{k+1,d}^{\max}} \varepsilon \operatorname{dist}(g_u, B(X_u, \mathcal{F}_u)) \\ &\geq w(X(d)_u) \frac{F_{i,k+1,d}^{\min}}{F_{k+1,d}^{\max}} \varepsilon w(X(d)_u)^{-1} \frac{F_{k,d}^{\min}}{F_{i,k,d}^{\max}} \operatorname{dist}(g, B(X_u, \mathcal{F}_u)^u) \\ &\geq \frac{F_{i,k+1,d}^{\min}}{F_{k+1,d}^{\max}} \frac{F_{k,d}^{\min}}{F_{i,k,d}^{\max}} \varepsilon (\|g\| - qw(u)) \\ &= \varepsilon \frac{F_{i,k+1,d}^{\min}}{F_{i,k,d}^{\max}} F_{k+1,d}^{\min}} \|g\| - \varepsilon \frac{F_{i,k+1,d}^{\min}}{F_{i,k,d}^{\max}} F_{k+1,d}^{\min}} qw(u) \end{split}$$

Combining this with the last inequality and rearranging gives the lemma.

13.4 Completion of The Proof

We will derive Theorem 12.8 from the following key lemma. We continue to assume that $f \in C^k(X, \mathcal{F})$, \mathcal{P} is a k-intersection profile, $h \in (0, 1]^{\mathcal{P}}$, Υ is as in §13.2, and T_i , $S_{\alpha,\beta}$, $U_{\mathcal{P}}$ are as in §11.1.

Lemma 13.6. With notation as before, let $\varepsilon_0, \ldots, \varepsilon_k \ge 0$ and $q \in [0, 1]$. Suppose that:

- (1) f is mock q-locally minimal;
- (2) for every $i \in \{0, ..., k\}$ and $u \in X(i)$, (X_u, \mathcal{F}_u) is an ε_i -coboundary expander in dimension k i 1;
- (3) \emptyset is not (f, h, \mathcal{P}) -heavy;

and put

$$\tilde{\varepsilon} = \min\left\{\frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}\varepsilon_i \,|\, i \in \{0,\dots,k\}\right\}.$$

Then

$$\tilde{\varepsilon}||f|| \le ||d_0f|| + U_{\mathcal{P}}w(\Upsilon) + q\sum_{i=0}^{k-1} \frac{F_{i,k+1,d}^{\min}F_{k,d}^{\min}}{F_{i,k,d}^{\max}F_{k+1,d}^{\max}}\varepsilon_i w(A_i)$$

Proof. We abbreviate $c_i := \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} \varepsilon_i$.

Let u denote a terminal face with dim $u \leq k$. Then $u \neq \emptyset$, because \emptyset is not heavy, and thus f is q-locally minimal at u. Therefore, by Lemma 13.5,

$$\varepsilon_{\dim u} \frac{F_{\dim u,k+1,d}^{\min} F_{k+1,d}^{\min}}{F_{\dim u,k,d}^{\max} F_{k+1,d}^{\max}} w(D(u)) \le w(D'(u) \cap [\operatorname{supp}(df) \cup \Upsilon]) + qc_{\dim u}w(u)$$
$$= w(D'(u) \cap [\operatorname{supp}(df) - \Upsilon]) + w(D'(u) \cap \Upsilon) + c_{\dim u}w(u).$$

When dim u = k, this moreover holds with q = 0, because f is 0-locally minimal at every $u \in X(k)$. By summing over all terminal faces u of dimension k or less, we get

$$\sum_{u} \varepsilon_{u} \frac{F_{\dim u,k+1,d}^{\min} F_{k}^{\min}}{F_{\dim u,k,d}^{\max} F_{k+1,d}^{\max}} w(D(u)) \leq$$

$$\sum_{u} w(D'(u) \cap [\operatorname{supp}(df) - \Upsilon]) + \sum_{u} w(D'(u) \cap \Upsilon) + \sum_{u:\dim u < k} qc_{\dim u} w(u).$$
(13.3)

We shall prove the lemma by bounding from about or below the four sums appearing in this inequality.

First, we have

$$\sum_{u} \varepsilon_{u} \frac{F_{\dim u,k+1,d}^{\min} F_{k,d}^{\min}}{F_{\dim u,k,d}^{\max} F_{k+1,d}^{\max}} w(D(u)) \ge \tilde{\varepsilon} \sum_{u} w(D(u)) \ge \tilde{\varepsilon} ||f||,$$

because every face in supp f descends to some terminal face u with dim $u \leq k$. Next,

$$\sum_{u} w(D'(u) \cap [\operatorname{supp}(df) - \Upsilon]) \le \|d_0 f\|$$

because every non-exceptional (k + 1)-face descends to a unique terminal subface (Lemma 13.3). On the other hand, if $y \in \Upsilon$, then y can descend to at most $U_{\mathcal{P}}$ terminal faces, see §11.1, so

$$\sum_{u} w(D'(u) \cap \Upsilon) \le U_{\mathcal{P}} \|\Upsilon\|.$$

Finally, since every terminal face of dimension $\leq k$ is heavy and nonempty, we have

$$\sum_{u:\dim u < k} qc_{\dim u} w(u) \le q \sum_{i=0}^{k-1} c_i w(A_i).$$

Plugging these inequalities into (13.3) gives the lemma.

Proof of Theorem 12.8. Recall that we are given $h \in (0,1]^{\mathcal{P}}$ such that

$$p := \tilde{\varepsilon} - U_{\mathcal{P}} S_{\alpha,\beta}(h) - q \sum_{i=0}^{k} \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}^{\max}} \varepsilon_i T_i(h) > 0$$

and $\tilde{\varepsilon}$ is as in Lemma 13.6. Furthermore, $f \in C^k(X, \mathcal{F})$ is mock *q*-locally minimal and $T_{-1}(h)^{-1}$ small. Again, write $c_i := \frac{F_{i,k+1,d}^{\min} F_{k,d}^{\min}}{F_{i,k,d}^{\max} F_{k+1,d}} \varepsilon_i$.
By Lemma 12.1 and $(L_{i,k,d})^{\min} C_{i,k+1,d} = C_i$.

By Lemma 13.1, $w(A_{-1}) \leq T_{-1}(h) ||f|| < 1$. Since X has only one -1-face and its weight is 1, the face \emptyset is not heavy. Now, Lemma 13.6 tells us that

$$\|df\| \ge \tilde{\varepsilon} \|f\| - U_{\mathcal{P}}w(\Upsilon) - q \sum_{i=0}^{k} c_i w(A_i)$$

By Lemmas 13.1 and 13.2, this means that

$$\|df\| \ge \tilde{\varepsilon} \|f\| - U_{\mathcal{P}} S_{\alpha,\beta}(h) \|f\| - q \sum_{i=0}^{k} c_i T_i(h) = p \|f\|.$$

A Correction Algorithm

The following is a time-efficient variant of Algorithm 8.2.

Algorithm A.1. Let $(X, w), \mathcal{F}, d, k$ be as in Theorem 8.1. The input to the algorithm is some $f \in C^k(X, \mathcal{F})$ and a real number $q \ge 0$. The algorithm outputs another k-cocycle $f' \in C^k(X, \mathcal{F})$, computed as follows:

- (1) $f' \leftarrow f$
- (2) $L \leftarrow \text{empty queue}$
- (3) $B \leftarrow$ boolean array indexed by $X(0) \cup \cdots \cup X(k)$
- (4) For each $z \in X(0) \cup \cdots \cup X(k)$:
 - (4a) L.push(z)
 - (4b) $B[z] \leftarrow \text{True} // z \text{ is in } L$
- (5) While L is not empty:
 - (5a) $z \leftarrow L.pop()$
 - (5b) $B[z] \leftarrow$ False // z is not in L
 - (5c) Search for $u \in X(0) \cup \cdots \cup X(k)$ and $g \in C^{k-\dim u-1}(X_z, \mathcal{F}_z)$ with $||df' d(g^u)|| < ||df|| q \cdot w(u)$.
 - (5d) If such u and g were found:
 - i. f' ← f' g^z.
 ii. For every z' ∈ X(0) satisfying inf{z, z'} = Ø and B[z'] = False:
 A. L.push(z')
 B. B[z'] ← True // z' is in L
- (6) Return f'.

Bibliography

- [AB08] Peter Abramenko and Kenneth S. Brown. *Buildings*, volume 248 of *Graduate Texts* in *Mathematics*. Springer, New York, 2008. Theory and applications.
- [Bjö84] A. Björner. Posets, regular CW complexes and Bruhat order. European J. Combin., 5(1):7–16, 1984.
- [Cur14] Justin Michael Curry. Sheaves, cosheaves and applications. ProQuest LLC, Ann Arbor, MI, 2014. Thesis (Ph.D.)–University of Pennsylvania.
- [DD19] Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science, pages 1495–1524. IEEE Comput. Soc. Press, Los Alamitos, CA, [2019] ©2019.
- [DD23] Yotam Dikstein and Irit Dinur. Coboundary and cosystolic expansion without dependence on dimension or degree, 2023.

- [DDHRZ20] Yotam Dikstein, Irit Dinur, Prahladh Harsha, and Noga Ron-Zewi. Locally testable codes via high-dimensional expanders. 2020.
 - [DEL⁺22] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In STOC '22—Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 357–374. ACM, New York, [2022] ©2022.
 - [DK17] Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In 58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017, pages 974–985. IEEE Computer Soc., Los Alamitos, CA, 2017.
 - [DKW18] Dominic Dotterrer, Tali Kaufman, and Uli Wagner. On expansion and topological overlap. *Geom. Dedicata*, 195:307–317, 2018.
 - [DLV24] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of higher-dimensional cubical complexes with application to quantum locally testable codes. 2024.
 - [EK16] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. In STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, pages 36–48. ACM, New York, 2016.
 - [EK17] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. 2017. Summary appeared in STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing.
 - [EK24] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. J. Amer. Math. Soc., 37(1):39–68, 2024. Summary appeared in STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing.
 - [FK23a] Uriya A. First and Tali Kaufman. The cheeger inequality and coboundary expansion: Beyond constant coefficients. 2023.
 - [FK23b] Uriya A. First and Tali Kaufman. On good 2-query locally testable codes from sheaves on high dimensional expanders. 2023.
 - [Fri15] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture: with an appendix by Warren Dicks. Mem. Amer. Math. Soc., 233(1100):xii+106, 2015. With an appendix by Warren Dicks.
 - [Gol11] Oded Goldreich. Short locally testable codes and proofs. In Studies in complexity and cryptography, volume 6650 of Lecture Notes in Comput. Sci., pages 333–372. Springer, Heidelberg, 2011.
 - [Gro10] Mikhail Gromov. Singularities, expanders and topology of maps. Part 2: From combinatorics to topology via algebraic isoperimetry. *Geom. Funct. Anal.*, 20(2):416–526, 2010.
 - [HG19] Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. J. Appl. Comput. Topol., 3(4):315–358, 2019.
 - [Ive86] Birger Iversen. Cohomology of sheaves. Universitext. Springer-Verlag, Berlin, 1986.

- [JL97] Bruce W. Jordan and Ron Livné. Ramanujan local systems on graphs. Topology, 36(5):1007–1024, 1997.
- [KKL16] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for Ramanujan complexes and topological expanders. *Geom. Funct. Anal.*, 26(1):250–287, 2016.
 - [KL14] Tali Kaufman and Alexander Lubotzky. High dimensional expanders and property testing. In ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science, pages 501–506. ACM, New York, 2014.
- [KM18] Tali Kaufman and David Mass. Cosystolic expanders over any abelian group. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:134, 2018.
- [KM21] Tali Kaufman and David Mass. Unique-neighbor-like expansion and groupindependent cosystolic expansion. In 32nd International Symposium on Algorithms and Computation, ISAAC 2021, December 6-8, 2021, Fukuoka, Japan, volume 212 of LIPIcs, pages 56:1–56:17, 2021.
- [KM22] Tali Kaufman and David Mass. Double Balanced Sets in High Dimensional Expanders. In Amit Chakrabarti and Chaitanya Swamy, editors, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022), volume 245 of Leibniz International Proceedings in Informatics (LIPIcs), pages 3:1–3:17, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [KO21] Tali Kaufman and Izhar Oppenheim. High dimensional expansion implies amplified local testability. 2021.
- [KO22] Tali Kaufman and Izhar Oppenheim. High Dimensional Expansion Implies Amplified Local Testability. In Amit Chakrabarti and Chaitanya Swamy, editors, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022), volume 245 of Leibniz International Proceedings in Informatics (LIPIcs), pages 5:1–5:10, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [KR06] Tali Kaufman and Dana Ron. Testing polynomials over general fields. SIAM J. Comput., 36(3):779–802, 2006.
- [KT23] Tali Kaufman and Ran J. Tessler. Garland's technique for posets and high dimensional Grassmannian expanders. In 14th Innovations in Theoretical Computer Science Conference, volume 251 of LIPIcs. Leibniz Int. Proc. Inform., pages Art. No. 78, 22. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023.
- [Li04] W.-C. W. Li. Ramanujan hypergraphs. Geom. Funct. Anal., 14(2):380–399, 2004.
- [LM06] Nathan Linial and Roy Meshulam. Homological connectivity of random 2-complexes. Combinatorica, 26(4):475–487, 2006.
- [LSV05] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Explicit constructions of Ramanujan complexes of type \tilde{A}_d . European J. Combin., 26(6):965–993, 2005.
 - [LZ22] Anthony Leverrier and Gilles Zémor. Quantum tanner codes. 2022.

- [Mes18] Roy Meshulam. Graph codes and local systems, 2018.
- [MLM94] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in geometry and logic*. Universitext. Springer-Verlag, New York, 1994. A first introduction to topos theory, Corrected reprint of the 1992 edition.
- [MW09] R. Meshulam and N. Wallach. Homological connectivity of random k-dimensional complexes. *Random Structures Algorithms*, 34(3):408–417, 2009.
- [Opp15] Izhar Oppenheim. Vanishing of cohomology and property (T) for groups acting on weighted simplicial complexes. *Groups Geom. Dyn.*, 9(1):67–101, 2015.
- [Opp18] Izhar Oppenheim. Local spectral expansion approach to high dimensional expanders Part I: Descent of spectral gaps. *Discrete Comput. Geom.*, 59(2):293–330, 2018.
- [PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In STOC '22—Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 375–388. ACM, New York, [2022] ©2022.
- [She85] Allen Dudley Shepard. A Cellular Description of The Derived Ccategory of a Stratified Space. ProQuest LLC, Ann Arbor, MI, 1985. Thesis (Ph.D.)–Brown University.
- [SS96] Michael Sipser and Daniel A. Spielman. Expander codes. volume 42, pages 1710–1722. 1996. Codes and complexity.