

Good Quantum Locally Testable Codes from Lossless Cubical Complexes

Itay Cohen^{*1} Itai Leigh^{†1} Assaf Reiner^{‡2} Amnon Ta-Shma^{§1}
Elad Tzalik^{¶3}

¹*Department of Computer Science, Tel Aviv University, Israel.*

²*Department of Mathematics, Hebrew University of Jerusalem, Israel.*

³*Department of Mathematics and Computer Science, Weizmann Institute of Science, Israel.*

Abstract

Sipser and Spielman constructed LDPC codes from either bipartite *spectral* expanders or one-sided *lossless* expanders. In higher dimensions, *spectral* expansion similarly played a central role in the constructions of asymptotically good classical LTCs and qLDPC codes by Dinur, Evra, Livne, Lubotzky, and Mozes and by Panteleev and Kalachev. Alternatively, Lin and Hsieh constructed classical LTCs and qLDPC codes from two-dimensional *lossless* cubical complexes.

In this work we develop the higher-dimensional *lossless* approach. We do not construct the required high-dimensional lossless cubical complexes; rather, we investigate what their existence would imply. We associate with a high-dimensional cubical complex a *level chain complex*, whose chain groups are supported on the level sets of the Boolean cube rather than on its cells. Our main technical contribution is a clean local-to-global theorem for this structure: suitable one-dimensional lossless expansion in the directional graphs implies small-set coboundary expansion of the global level complex. As a consequence, sufficiently imbalanced, two-sided lossless four-dimensional cubical complexes give rise to asymptotically good quantum locally testable codes. We expect the local-to-global principle developed here to have further applications.

^{*}Supported by ERC starting grant 949499.

Email: itay3@mail.tau.ac.il.

[†]Supported by the European Union (ERC, EACTP, 101142020), and the Danish-Israeli Study Foundation in Memory of Josef and Regine Nachemson.

Email: itai.leigh@mail.huji.ac.il.

[‡]Supported by the European Research Council (ERC) under the European Union's Horizon 2020 (N. 882751).

Email: assaf.reiner@mail.huji.ac.il.

[§]Funded by the Israel Science Foundation grant 443/22.

Email: amnon@tauex.tau.ac.il.

[¶]Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

Email: elad.tzalik@weizmann.ac.il.

Any views, opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the European Union or the European Research Council Executive Agency. Neither the European Union nor any granting authority can be held responsible for them.

Contents

1	Introduction	3
1.1	The gradual emergence of expansion and homology in coding theory	4
1.2	The combinatorial framework	6
1.3	The result	8
1.4	The technique	10
1.5	Relation to concurrent work	12
2	Cubical complexes	13
2.1	Simplicial and cubical complexes	13
2.2	The level chain complex	15
2.3	Links	17
2.4	Small-set coboundary expansion	18
2.5	Lossless expansion	18
3	Expansion of the links	19
3.1	Ordered paths and collisions	20
3.2	Embedding the non-colliding part	22
3.3	Transfer of multipartite expansion	24
4	Small-set coboundary expansion from square complexes	26
4.1	Heavy and terminal vertices	26
4.2	The global to local reduction	28
4.3	Excess	29
4.4	Same Type Collisions	32
4.5	Light Squares	33
4.6	Heavy-mass propagation	34
4.7	Choice of parameters and final error bound	35
5	qLTC	38
5.1	Definitions	38
5.2	Construction and Proof	39

1 Introduction

Error-correcting codes enable reliable communication and storage in the presence of noise by encoding information into longer, redundant words. A linear code can be described either by a generating matrix or by a parity-check matrix. Although these two descriptions define the same code, they correspond to fundamentally different computational tasks. The generating matrix maps an information message to its encoded codeword. The parity-check matrix, on the other hand, acts on a received, possibly corrupted, word, producing its syndrome and enabling efficient verification, decoding, and local correction.

These two representations exhibit a fundamental asymmetry. Every column of a generating matrix is itself a codeword, as it is the image of a standard basis vector. Consequently, if the code has linear distance, every nonzero column must have linear Hamming weight, ruling out sparse generating matrices. Sparse parity-check matrices, in contrast, are fully compatible with good code parameters. This observation led Gallager to introduce *low-density parity-check (LDPC) codes* [Gal62], whose parity-check matrices have constant row and column weight. Since their introduction, LDPC codes have become one of the central objects in coding theory, combining excellent asymptotic parameters with efficient decoding algorithms [SS96].

The sparsity of the parity-check matrix turned out to be important for much more than efficient decoding. Because each parity check involves only constantly many symbols, it enables algorithms that inspect only a small portion of the received word. Beginning with Gallager's iterative decoding algorithms [Gal62], this observation evolved into a broad theory of *local* algorithms for error-correcting codes. Depending on the task, one seeks to locally correct individual symbols of a corrupted codeword, locally decode the underlying message, locally test whether a word belongs to or is close to the code, or infer global properties of a word from a small number of queries. These notions have become central throughout coding theory and have found numerous applications in complexity theory, property testing, probabilistically checkable proofs, data structures, and cryptography.

Among these local tasks, local testability occupies a particularly distinguished position. A code is *locally testable* if there exists a randomized algorithm that queries only constantly many coordinates of the input word and accepts every codeword with probability one, while rejecting words with probability proportional to their distance from the code. Local testability is one of the central themes underlying PCP constructions. Goldreich and Sudan isolated this notion in the form of locally testable codes (LTCs) [GS06], and it is a key ingredient in Dinur's gap-amplification proof of the PCP theorem [Din07]. Dinur further showed how to construct constant-query locally testable codes encoding n information bits into $n \text{ polylog}(n)$, leaving open whether it is possible to eliminate the polylogarithmic overhead. This question was resolved only recently by the breakthrough constructions of Panteleev and Kalachev [PK22] and independently Dinur, Evra, Livne, Lubotzky and Mozes [DEL⁺26], who constructed asymptotically good c^3 -locally testable codes (the c^3 indicates all three parameters are constants: rate, relative distance and number of queries). Nevertheless, the existence of linear-length PCPs remains one of the central open problems in complexity theory.

In a seemingly unrelated direction, the quest for scalable quantum computation has motivated the development of quantum error-correcting codes. Since the pioneering works of Shor [Sho96] and Steane [Ste96b] on quantum error correction, and the threshold theorems of Aharonov and Ben-Or [ABO97] and of Knill, Laflamme, and Zurek [KLZ96] establishing the possibility of fault-tolerant quantum computation, quantum error correction has been regarded as one of the

principal ingredients for building scalable quantum computers. A particularly influential milestone was Kitaev’s toric code [Kit03], which demonstrated that quantum information can be protected using local stabilizer measurements arranged on a two-dimensional lattice, introducing both topological quantum codes and geometric locality as guiding principles. The toric code, however, has vanishing rate, motivating the search for quantum codes with sparse stabilizer checks and *asymptotically good* parameters. This led to the study of *quantum low-density parity-check (qLDPC) codes*, the quantum analogue of classical LDPC codes. Following a sequence of breakthrough constructions, it is now known that qLDPC codes can simultaneously achieve constant rate, linear distance, and sparse stabilizer checks, resolving a longstanding open problem in quantum coding theory.

Beyond efficient quantum error correction, one may ask for stronger local robustness properties. Quantum locally testable codes (qLTCs), introduced by Aharonov and Eldar [AE18], require that every low-energy state is close to the code space, providing a quantum analogue of classical local testability. Such robust local detection has attracted considerable attention because of its potential connections to self-correcting quantum memories, the quantum PCP conjecture, and fault-tolerant architectures with passive error suppression. The characterization of quantum local testability developed by Aharonov, Eldar, and collaborators naturally suggests a higher-dimensional homological framework, as we next explain.

1.1 The gradual emergence of expansion and homology in coding theory

Although the classical and quantum theories developed largely independently, a common mathematical picture gradually emerged. Expansion emerged as the combinatorial mechanism that transforms sparse local constraints into global distance and robustness guarantees, while homology emerged as the geometric framework that organizes these constraints and their interactions. We next review how such ideas evolved through the successive development of LDPC codes, LTCs, qLDPC codes, and qLTCs.

Gallager’s original work already suggested that the combinatorial structure of the Tanner graph plays a decisive role in the performance of an LDPC code. A natural first principle was to require the graph to have large girth, ensuring that local neighborhoods remain tree-like and preventing short cycles from creating spurious correlations during iterative decoding. While this viewpoint proved highly successful in practice, it could not by itself explain the global distance and robustness properties exhibited by non-explicit LDPC codes.

The decisive advance came with the work of Sipser and Spielman [SS96], who identified combinatorial expansion as the property underlying good LDPC codes. They showed that if every sufficiently small set of variable nodes has almost the maximum possible number of neighboring check nodes, then the resulting LDPC code has linear distance and admits efficient decoding. This theorem identified a remarkably clean combinatorial mechanism explaining why sparse local constraints can nevertheless enforce strong global properties. At the time, however, such explicit lossless expanders with the required parameters were not known. Sipser and Spielman [SS96] therefore complemented their combinatorial theorem with an elegant analysis based on spectral expanders, showing that comparable coding-theoretic guarantees can also be derived from spectral expansion assumptions.

The work of Sipser and Spielman [SS96] illustrates a recurring theme in the development of coding theory where the identification of a clean combinatorial principle often precedes its

realization. One first isolates the combinatorial property underlying a coding-theoretic phenomenon, and only later learns how to construct objects satisfying this property explicitly, or to derive comparable conclusions from weaker assumptions using more sophisticated techniques. The study of lossless expanders provides a striking illustration of this pattern. Random bipartite graphs establish the existence of optimal lossless expanders by probabilistic arguments, while explicit constructions proved considerably more challenging. Over the following decades, a substantial body of work developed explicit constructions of balanced and unbalanced bipartite lossless expanders, eventually realizing the combinatorial objects that had originally motivated the theory [CRVW02, HLM⁺25].

Local testability introduces a phenomenon that does not arise in ordinary LDPC codes. A simple thought experiment illustrates the difficulty. Consider an LDPC code defined by a full-rank parity-check matrix whose Tanner graph is a good expander. Removing a single parity check enlarges the code by a factor of two, while leaving the expansion, and therefore the linear distance, essentially unchanged. Consequently, every new codeword lies far from the original code, yet violates only the deleted parity check. Such a word is globally far from the code but locally appears almost valid, showing that independent parity checks cannot support local testability. This intuition was formalized by Ben-Sasson, Guruswami, Kaufman, Sudan, and Viderman [BSGK⁺10], who proved that locally testable codes necessarily require substantial redundancy among their checks.

From an algebraic perspective, this redundancy means that the parity-check matrix is no longer an isolated linear map. Besides describing the code as its kernel, one must also understand the linear relations among its rows. Equivalently, the image of the parity-check matrix must itself satisfy another family of linear constraints. Rather than a single linear map, one is naturally led to a sequence of linear maps $C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$, where the image of each map is contained in the kernel of the next. This is precisely the setting of chain complexes.

This homological viewpoint emerged only gradually. Initially, an LDPC code is naturally described by its Tanner graph. From today's perspective, however, the code associated with a Tanner graph is simply a one-dimensional chain complex. This interpretation is easy to overlook, and perhaps rightly so, because in one dimension the homological language adds little beyond the graph itself. The graph was always present; what emerged only later was the realization that it belongs to a hierarchy of increasingly rich chain complexes. Only after two-dimensional constructions entered coding theory did it become clear that classical LDPC codes were already the first instance of a broader homological picture.

This perspective proved remarkably fruitful. Dinur, Evra, Livne, Lubotzky, and Mozes [DEL⁺26] constructed asymptotically good locally testable codes using left-right Cayley complexes, which can be viewed as a special class of two-dimensional cubical complexes. Here, homological expansion in the middle layer of the chain complex provides the redundancy between constraints required for local testability. Around the same time, the breakthrough constructions of Panteleev and Kalachev [PK22], together with subsequent developments, demonstrated that closely related two-dimensional structures also suffice to construct asymptotically good quantum LDPC codes.

Following these developments, Lin and Hsieh [LH22a, LH22b] applied the balanced product construction of Breuckmann and Eberhardt [BE21] to lossless expanders, establishing the homological expansion properties underlying their classical LTCs and quantum LDPC codes. A particularly appealing feature of their approach is that in both settings their results share the same combinatorial object, and the same expansion argument. Viewed in the language of lossless square complexes, their results provide the basis for the unified framework we develop here.

The distinction lies in where the information is stored. In the classical construction, the bits occupy one end of the chain complex, the adjacent boundary map defines the parity checks, and the second boundary map provides the checks on checks required for local testability. In the quantum construction, the two ends of the chain complex correspond to the X - and Z -stabilizer, with the qubits sitting between them. Correspondingly, the classical construction requires one-sided lossless expansion, whereas controlling both quantum distances naturally leads to two-sided lossless expansion.

The natural higher-dimensional analogue of a square complex is a cubical complex. Explicit higher-dimensional cubical complexes with strong spectral expansion are already known and have led to the construction of lossless vertex expanders mentioned before. The characterization of quantum local testability developed by Aharonov, Eldar, and collaborators naturally suggests extending the homological picture one step further. Although quantum local testability is defined for arbitrary stabilizer codes, the CSS setting admits a particularly natural realization in which one introduces an additional layer of relations for the X -stabilizer checks and another one for the Z -stabilizer checks. This naturally points toward higher-dimensional cubical complexes equipped with suitable local expansion properties. Indeed, Dinur, Lin and Vidick [DLV24] constructed quantum locally testable codes with polylogarithmic overhead this way.

In contrast, higher-dimensional cubical complexes satisfying suitable two-sided lossless expansion are not currently known, even by probabilistic arguments. This situation closely parallels the earlier development of LDPC codes and LTCs. In both cases, a clean combinatorial theorem based on lossless expansion preceded the explicit construction of the corresponding combinatorial object. The present work continues this pattern in higher dimension. We show that

**higher-dimensional cubical complexes with suitable two-sided lossless expansion
give rise to asymptotically good quantum locally testable codes**

and leave open the challenge of proving that these combinatorial objects exist and constructing them explicitly.

1.2 The combinatorial framework

The basic combinatorial object in this paper is a *decorated cubical complex*. Such a complex in dimension k consists of vertex sets

$$\{V_I\}_{I \subseteq [k]},$$

indexed by the subsets of $[k]$, together with undirected graphs connecting

$$V_I \quad \text{and} \quad V_{I \cup \{i\}},$$

whenever $i \notin I$. We call each set V_I a block. Thus a decorated cubical complex is obtained from the Boolean cube by replacing each vertex with a block and each edge with a bipartite graph between the corresponding blocks. We refer to the edges between V_I and $V_{I \cup \{i\}}$ as the *direction- i edges*.

A natural chain complex associated with a cubical complex is its *cellular* chain complex, whose degree- i space is spanned by the i -dimensional cubes. Constructions organized around cells, with coefficient spaces supplied by local codes, appear in [DEL⁺26, DLV24, GGJ26]. In

this work, however, we use the *level* chain complex, extending the two-dimensional framework of [PK22, LH22a, LH22b]. We write

$$V_i = \bigcup_{|I|=i} V_I,$$

and call V_i the *level* i of the complex. The level chain complex has $\mathbb{F}_2^{V_i}$ as its i th vector space and its boundary and coboundary maps are induced by adjacency between consecutive levels. To ensure that successive coboundary maps compose to zero, we require the *weak completion property*, which asserts that for each pair of edges in two distinct directions coming out of a vertex, there exists a unique square in the complex containing them. The weak completion property ensures that successive boundary maps compose to zero: each two-step path is paired with the other path around its unique square, and the two contributions cancel over $\mathbb{F}_2$.

The simplest examples of cubical complexes are left-right Cayley complexes. Here there are a (possibly non-Abelian) group Γ and four vertex sets

$$V_\emptyset = V_{\{1\}} = V_{\{2\}} = V_{\{1,2\}} = \Gamma,$$




Figure 1: A left-right Cayley complex satisfying the weak completion property but not the strong one. The blue and red squares share the opposite vertices g and $agb = a'gb'$. The black dashed contour follows a four-cycle which is not a square in the complex.

In the left-right complex, the direction-1 edges are defined by left multiplication by generators from a symmetric set A , while the direction-2 edges are defined by right multiplication by

generators from a symmetric set B . For every $g \in V_\emptyset$, every choice of a direction-1 edge and a direction-2 edge leaving g leads uniquely to

$$ag \in V_{\{1\}}, \quad gb \in V_{\{2\}}, \quad agb \in V_{\{1,2\}},$$

which form a unique square $\{agb; ag, gb; g\}$, and these are all the squares. This is precisely the weak completion property. We remark that unlike this group example, in a general decorated cubical complex the vertex sets V_I need not have the same cardinality.

Throughout this paper we assume that every direction graph between V_I and $V_{I \cup \{i\}}$ is a lossless expander. We call such a decorated cubical complex *direction-wise lossless*. Since every pair of blocks V_I, V_J with $I \subseteq J$ is connected by a monotone sequence of direction graphs, the corresponding lossless expansion estimates compose to give analogous expansion estimates between every such pair of blocks. It is therefore natural to ask whether the resulting cubical complex itself is an expanding object, and if so, in what way.

A first attempt would be to require every sufficiently small cochain to have a large coboundary. This, however, is impossible. Since $\delta^2 = 0$, every coboundary has zero coboundary. Thus any meaningful notion of expansion must first quotient out these trivial obstructions. We say

- A cochain in V_i is *minimal* if it has minimum norm among all cochains in its cohomology class, or equivalently, if adding the coboundary of an arbitrary $(i-1)$ -cochain cannot reduce its norm.
- A cochain is *locally minimal* if no local coboundary operation supported in the neighborhood of a single link can decrease its norm.

Thus minimality is a global notion, whereas local minimality can be verified locally.

We say that a decorated cubical complex is a *small-set level coboundary expander* if every sufficiently small minimal cochain supported on a level below the top has large coboundary. The terminology reflects the fact that ordinary vertex expansion takes place between adjacent blocks, whereas small-set level coboundary expansion measures expansion between adjacent levels, after quotienting out the trivial cohomological obstructions. Similarly, we say the cubical complex is a *robust*¹ small-set level coboundary expander if every sufficiently small *locally* minimal cochain has large coboundary.

1.3 The result

The central combinatorial statement of this paper is the following theorem:

Theorem 1.1 (Informal version of [Theorem 4.10](#)). *Every direction-wise lossless decorated cubical complex satisfying suitable lossless expansion and imbalance assumptions is a robust small-set level coboundary expander.*

This implication is far from obvious. The hypothesis concerns only the direction graphs of the complex, whereas the conclusion is a genuinely higher-dimensional expansion property.

[Theorem 1.1](#) is the higher-dimensional analogue of the two-dimensional expansion theorem underlying the constructions of Lin and Hsieh [[LH22a](#), [LH22b](#)]. We establish the corresponding combinatorial statement in arbitrary dimension k .

¹Outside the introduction, we omit the adjective “robust”: small-set coboundary expansion always refers to expansion of locally minimal cochains.

Robust small-set level coboundary expansion has the following coding-theoretic consequences. In dimension two, robust small-set level coboundary expansion immediately yields LTCs, recovering the square-complex constructions of Lin and Hsieh [LH22a]. The same argument, applied to the very same square complex but interpreted as a CSS chain complex, yields quantum LDPC codes [LH22b]. More generally, following the argument of Aharonov and Eldar [AE18], a four-dimensional robust small-set level coboundary expander gives rise to a quantum locally testable code whenever the corresponding chain complex is interpreted as a CSS code. Thus,

$$\boxed{\text{robust small-set level coboundary expander}} \implies \begin{cases} \text{LTC}, & k = 2, \\ \text{qLDPC}, & k = 2, \\ \text{qLTC}, & k = 4. \end{cases}$$

Combining this implication with Theorem 1.1, we obtain

$$\begin{array}{c} \boxed{\text{Direction-wise lossless decorated cubical complex} + \text{imbalance conditions}} \\ \Downarrow \\ \boxed{\text{Robust small-set level coboundary expander}} \\ \Downarrow \\ \boxed{\text{qLTC}} \end{array}$$

The remaining challenge is therefore purely combinatorial: construct explicit higher-dimensional direction-wise lossless decorated cubical complexes.

A candidate construction, and an obstruction. Explicit constant-degree lossless vertex expanders were constructed by Hsieh et al. [HLM⁺25], but it is unclear how to adapt their construction to obtain higher-dimensional cubical complexes with lossless directional graphs. A natural starting point is the family of balanced cubical complexes underlying their construction, obtained as quotients of products of $(p_i + 1)$ -regular Bruhat–Tits trees by congruence subgroups of quaternionic arithmetic lattices. Their directional graphs inherit the Ramanujan property from the graphs of Lubotzky, Phillips, and Sarnak [LPS88]. Our setting needs imbalance, so we set out to construct the imbalanced analogues of those complexes. The natural input is the family of Ramanujan *bigraphs* recently constructed by Evra, Feigon, Maurischat and Parzanchevski [EFMP23, EFMP25], which are congruence quotients of the $(p^3 + 1, p + 1)$ -biregular Bruhat–Tits tree of a definite unitary group in three variables. Replacing the quaternionic lattice by the unitary one, and the product of regular trees by a product of biregular ones, one obtains for every k a strongly explicit k -dimensional imbalanced decorated cubical complex with complete multipartite links, whose directional graphs are all Ramanujan bigraphs of degrees $(p_i^3 + 1, p_i + 1)$, and whose degrees may be imbalanced as strongly as our imbalance assumption requires. Such a complex satisfies every hypothesis of Theorem 1.1 except one, the lossless expansion of its directional graphs.

The hope that this last hypothesis holds stems from the heuristic that led Kamber and Kaufman [KK22] to conjecture that the LPS graphs are lossless [KK22, Conjecture 1.11]. Pure spectral arguments cannot establish losslessness, as demonstrated by Kahale [Kah95] in the regular case

and by Asherov and Dinur [AD24] in the biregular setting. However, the known counterexamples of Kamber and Kaufman, namely the Morgenstern graphs, fail due to an obstruction with no counterpart over $\mathbb{Q}$, which is the existence of a proper closed subfield of the local field. Because $\mathbb{Q}_p$ has no proper closed subfields, the same reasoning supports both the losslessness of LPS graphs and that of the EFMP bigraphs, which are likewise defined over $\mathbb{Q}$.

It turns out that the opposite is the case. Adapting the closed orbit method of [KK22] to the unitary setting, one finds in every EFMP bigraph a nonempty set that does not expand at all. The set arises as the image of a subtree attached to an algebraic subgroup, rather than to a subfield as in [KK22]. All in all, the EFMP bigraphs, and with them the complexes described above, are not two-sided lossless. These results appear to be of independent interest, and we intend to present them separately rather than develop them here.

The remainder of the introduction is devoted to explaining why the combination of a cubical complex with one-dimensional lossless expansion along the direction graphs suffices to establish this genuinely higher-dimensional expansion theorem.

1.4 The technique

The main technical contribution of this work is a dimension-independent global-to-local reduction, building on ideas introduced by [FK24, EK24, KKL16]. The key idea is that a global cochain should be understood through its local views. The first question is therefore: what is the correct notion of a local view in a decorated cubical complex?

The answer is the notion of a *link*. The link of a vertex x consists of x and all vertices in higher levels that are connected to x by a monotone path, together with their incidence structure. It is therefore the smallest upward closed subcomplex containing the vertex x . Restricting a cochain to the link of a vertex x captures precisely its local behavior around that vertex x . Consequently, links are the natural local structures into which the input cochain should be decomposed.

Then, given a cochain

$$S : V_i \rightarrow \mathbb{F}_2$$

of norm at most α , we decompose S into a collection of cochains $L_t : V_i \rightarrow \mathbb{F}_2$ that are local in the sense that each of them is supported on an individual link. Our global-to-local theorem constructs a local decomposition

$$S \rightsquigarrow \{L_t\}_{t \in T},$$

such that the global coboundary is lower bounded, up to a small error, by the sum of the local coboundaries, namely,

Theorem 1.2 (informal). *Every sufficiently small cochain $S : V_r \rightarrow \mathbb{F}_2$ in a direction-wise lossless cubical complex satisfying the imbalance assumptions admits cochains L_t , each supported in $\text{Lk}(t)$, such that*

$$\bigcup_t \text{supp}(L_t) = \text{supp}(S), \quad \|\delta S\| \geq \sum_t \|\delta L_t\| - C_{\varepsilon,k} \|S\|,$$

where $C_{\varepsilon,k} \rightarrow 0$ as the loss parameters vanish.

For the definition of the norm, see Equation (2.1). For a precise statement, see Proposition 4.2.

Consequently, proving robust level expansion reduces to establishing expansion for the local structures L_t . Thus the proof naturally separates into two independent parts: a dimension-independent global-to-local decomposition and a purely local expansion theorem.

Let us now try to isolate the local components. We start with a cochain $S : V_r \rightarrow \mathbb{F}_2$. A first attempt would be to replace any two (or more) vertices in the support of S that belong to the same link by that link, and then add each of the remaining isolated vertices in the support of S as a link of its own (of size 1). However, this decomposition is too fine to allow a global-to-local reduction.

Instead, following [FK24], we use only those links which have a large intersection with $\text{supp } S$.² We call vertices with such links *heavy*. More accurately, the heavy vertices are organized recursively across the levels of the complex. The set of heavy vertices on level r is exactly $H_r = \text{supp } S$. Then, starting from the support of the input cochain we go down the levels, repeatedly declaring a vertex to be heavy whenever a sufficiently large fraction of its neighbors in the layer above it are already heavy. This process terminates at a collection of *terminal* vertices. These terminal vertices, together with the hierarchy, give rise to the local cochains. Each terminal vertex t gives a local cochain L_t , supported on all the vertices in V_r that are above it in the heavy hierarchy.

The main difficulty is that these local cochains are not disjoint, so a single element can contribute to several local chains L_t , which may cause overcounting. Furthermore, the coboundaries of L_t for different t 's may overlap, leading to cancellations over $\mathbb{F}_2$ in the next level V_{r+1} that do not occur when we count the size of the coboundary of each L_t separately. The heart of our global-to-local reduction is to show that these overlaps are negligible.

The proof shows that the support of δL_t usually does not intersect with the support of $\delta L_{t'}$ of a different link. To show this, we track the number of exceptions to this rule. When they do intersect on a vertex v , we charge this collision to the vertex v , and say that it has *excess*.

Vertices always inherit the excess of their heavy children. In Lemma 4.3 we show that there are only two mechanisms by which a vertex can acquire additional excess. The first comes from collisions in a directional expander graph, and the other comes from a square completion to a non-heavy vertex. See Figure 4 for an illustration of the two mechanisms.

As these two mechanisms are one- and two-dimensional in nature, they were already addressed, in some form, in [LH22a, LH22b]³. Direction-wise lossless expansion completely controls the first mechanism, whereas for the second mechanism an additional imbalance condition is required. In Sections 4.4 to 4.6, we show that requiring both expansion and imbalance implies that the total overlap between distinct local structures is small.

Crucially, local minimality *plays no role* in the global-to-local reduction. It enters only in the local expansion theorem, where its purpose is to eliminate the local cohomological obstructions to expansion. Thus links and local minimality arise together: once links are identified as the local structures, local minimality becomes the natural notion of local irreducibility. Accordingly, we define a cochain to be locally minimal if no coboundary operation supported inside a single link decreases its norm. This is precisely the hypothesis under which the local expansion theorem is proved. In a sense, this approach mirrors the conceptual picture developed in the previous subsections. There, expansion and homology emerged as the two organizing principles of the theory. Both are inherently global notions. Here, they reappear in local form. Links identify the correct notion of locality, while local minimality is the corresponding notion of minimality on

²The precise definition uses the direction-dependent factors encoded by the norm, which we suppress in the introduction.

³We stress that the proof of small-set coboundary expansion by [LH22a, LH22b] differs from the one we give here, *already in dimension 2*. In their work, the second mechanism bounds *all* squares (light or not), while we apply a light/heavy distinction, and explicitly decompose the cochain. The distinction is required to handle the richer structure of cancellations over $\mathbb{F}_2$ that may occur in the coboundary of high dimensional cochains.

a single link and provides the corresponding local homological notion. A central insight of this work is that in *direction-wise lossless* decorated cubical complexes, the expansion in the links is lifted automatically to the entire complex through the global-to-local reduction. In [Section 3](#), we establish the required relaxed form of coboundary expansion for the links. Combining these results yields the following theorem:

Theorem 1.3 (Robust small-set coboundary expansion at an arbitrary level). *Let X be a k -dimensional lossless square complex, with parameters $(D_i, D'_i, c_i, \varepsilon_i)$. If for every $s < t$, $\varepsilon_t D'_s \leq \sum \varepsilon_i$ then there exist constants c, a independent of $|X(0)|$ such that X is a robust $(c, a - \varepsilon)$ small-set coboundary expander, for $\varepsilon = O_k \left((\sum \varepsilon_i)^{\Omega_k(1)} \right)$*

1.5 Relation to concurrent work

While this work was being completed, Gay and Jeronimo [[GGJ26](#)] gave the first construction of asymptotically good quantum locally testable codes. Their construction, like ours, is based on higher-dimensional cubical complexes, but the algebraic structures associated with the complexes and the mechanisms underlying the two approaches are rather different. Gay and Jeronimo work with a cosheaf complex whose cochain groups are built from the cubical faces. In contrast, we work with the level complex of a decorated cubical complex: all cochains are supported on vertices, with their degree determined by their level in the Boolean-cube decoration.

The two approaches also make a different tradeoff between geometry and local codes. Gay and Jeronimo use arithmetic Ramanujan cubical complexes together with carefully chosen Reed–Solomon codes satisfying product-expansion properties. Thus, spectral expansion of the geometry is supplemented by strong expansion properties of the local codes. In our framework there are no nontrivial local codes: the burden is instead placed entirely on the geometry. We show that sufficiently strong lossless expansion of the one-dimensional directional graphs implies the higher-dimensional small-set level coboundary expansion needed for quantum local testability. Thus [[GGJ26](#)] establishes the existence of asymptotically good qLTCs, while the present work isolates a different, purely combinatorial route to such codes.

There is a useful historical parallel with the development of expander codes. Sipser and Spielman first identified lossless expansion as a clean combinatorial property through which sparse local constraints can enforce strong global coding properties, even though the explicit constructions available at the time were spectral expanders and their coding results were therefore also developed under spectral hypotheses. Subsequent work on lossless expanders showed that this stronger combinatorial notion is important in its own right, and eventually led to explicit constructions with increasingly strong parameters. We believe that a similar phenomenon may occur here. Lossless graph expansion has repeatedly proved useful beyond what can be obtained from spectral expansion alone, and our results suggest that its higher-dimensional cubical analogue is likewise a natural object to study. We therefore believe that suitably imbalanced cubical complexes with two-sided lossless directional expansion exist. Just as lossless expanders became fundamental objects independently of their original coding-theoretic motivation, we expect lossless cubical complexes to have applications beyond the particular qLTC construction considered here. We therefore view their construction as a natural higher-dimensional expansion problem in its own right.

Acknowledgements This work grew out of a summer reading group on qLTCs. We thank Aseel Sakhnini, Sergey Komech, and Yali Aharon Holzman for many helpful discussions. We

also thank Shai Evra and Alex Lubotzky for illuminating conversations about constructions of lossless cubical complexes.

Use of AI We used ChatGPT 5.4-6 and Codex to assist with drafting and editing the manuscript, refining the exposition, and checking mathematical calculations. A preliminary version of [Lemma 4.3](#) was obtained in a human-guided conversation with GPT-5.5, asking to extend our low-dimensional recursion using the same/light mechanism we proposed. All main ideas of this work are human-generated. All AI-assisted material was reviewed by the authors, who take full responsibility for the content of this work.

2 Cubical complexes

We define cubical complexes alongside multipartite simplicial complexes, highlighting their similarities and connections. Along the way, we develop the notation and gradually build toward the statements of our main results.

Throughout, all complexes and vertex blocks are finite and nonempty. We write $[k] = \{1, \dots, k\}$ and take all chains and cochains over $\mathbb{F}_2$.

2.1 Simplicial and cubical complexes

Definition 2.1 (Simplicial complex). *A simplicial complex Y consists of a vertex set $Y(0)$ and a downward-closed collection of nonempty subsets of $Y(0)$. For $i \geq 0$, let*

$$Y(i) = \{\sigma \in Y : |\sigma| = i + 1\}$$

be the set of i -dimensional faces.⁴ In particular, $Y(1)$ is the edge set of the underlying graph $(Y(0), Y(1))$. For $i \geq 0$, let

$$\text{Cl}_Y(i) = \left\{ \sigma \in \binom{Y(0)}{i+1} : \binom{\sigma}{2} \subseteq Y(1) \right\}.$$

Thus $\text{Cl}_Y(i)$ is the collection of $(i+1)$ -vertex cliques in this graph, and $Y(i) \subseteq \text{Cl}_Y(i)$.

Equivalently, a simplicial complex is specified by its underlying graph and a collection of declared cliques from $\{\text{Cl}_Y(i)\}_i$, closed under taking subsets. We now define multipartite complexes. To draw the analogy with decorated cubical complexes appearing later, we call them decorated simplicial complexes.

Definition 2.2 (Decorated simplicial complex). *A simplicial complex Y is k -decorated if $Y(0) = V_1 \sqcup \dots \sqcup V_k$ and every edge has endpoints in distinct blocks. Consequently, every face contains at most one vertex from each block.*

A decorated simplicial complex is modeled by the $(k-1)$ -simplex: the decoration map $Y(0) \rightarrow [k]$, which sends every vertex in V_j to j , constrains the i -faces in Y (sets of size $i+1$) only to those whose image under the decoration map is an i -face. In that sense, the decoration gives a global model dictating which faces are allowed.

⁴We adjoin the empty face, of dimension -1 , and set $Y(-1) = \{\emptyset\}$.

For cubical complexes, we use the k -dimensional Boolean cube Q_k as our global model. Its vertices are the subsets of $[k]$, where I and J are adjacent precisely when they differ by one element $|I \triangle J| = 1$.

Definition 2.3 (Cube-decorated graph). *A graph G is k -cube-decorated if its vertex set is partitioned into 2^k blocks V_I , indexed by $I \subseteq [k]$, and every edge joins blocks V_I, V_J with $\{I, J\} \in E(Q_k)$.*

For $v \in V_I$, write $\text{type}(v) = I$. An embedded i -cube in a k -cube-decorated graph G is a set $Q \subseteq V(G)$ satisfying the following conditions: (1) the vertices of Q have distinct types and $\{\text{type}(v) : v \in Q\}$ is the vertex set of an i -face of Q_k ; and (2) any two vertices of Q with adjacent types in Q_k are adjacent in G . We write $\text{Cu}_G(i)$ for the set of all embedded i -cubes in G . An embedded cube $Q' \subseteq Q$ is called a subcube of Q .

Definition 2.4 (Decorated cube complex). *A decorated k -cube complex X is a k -cube-decorated graph $G = (X(0), X(1))$ with sets $X(i) \subseteq \text{Cu}_G(i)$, $0 \leq i \leq k$, closed under subcubes, meaning that every j -subcube of a cube in $X(i)$ belongs to $X(j)$.*

As with the cube Q_k , we regard G as edge-colored by $[k]$: an edge between V_I and V_J has color, or direction, i when $I \triangle J = \{i\}$. An edge set is *rainbow* if its directions are pairwise distinct.

We now describe completion properties, which are crucial to this work. A completion property is an additional assumption guaranteeing that certain configurations extend to a unique cube in the complex, of the corresponding dimension.

Definition 2.5 (Completion properties). *Let $1 \leq r \leq k$.*

1. *X has r -joint completion if, for every $1 \leq t \leq r$, every vertex together with t incident edges in distinct directions is contained in exactly one member of $X(t)$.*
2. *X has r -path completion if, for every $1 \leq t \leq r$, every path of t edges in distinct directions is contained in exactly one member of $X(t)$.*
3. *X has r -rainbow-tree completion if, for every $1 \leq t \leq r$, every rainbow tree in $X(1)$ with t edges is contained in exactly one member of $X(t)$.*

In each case, uniqueness is among the cubes in $X(t)$, and not with respect to $\text{Cu}_G(t)$.

Lemma 2.6. *For every decorated k -cube complex and $1 \leq r \leq k$, the completion properties are equivalent.*

$$r\text{-joint completion} \iff r\text{-rainbow-tree completion} \iff r\text{-path completion}.$$

Proof. We prove that the three properties imply one another cyclically. Both nontrivial implications use induction on the number $t \leq r$ of edges, with $t = 1$ immediate.

Joint $\Rightarrow$ rainbow. Let T be a rainbow tree with $t \geq 2$ edges, and delete a leaf u adjacent to v . By induction, $T \setminus \{u\}$ has a unique completion $Q' \in X(t-1)$. Complete the t -joint formed by uv and the edges of Q' at v to $Q \in X(t)$. Its face through v in the directions of Q' equals Q' by joint completion, so $T \subseteq Q$. Any member of $X(t)$ containing T contains Q' by induction, hence this joint, and therefore equals Q .

Rainbow $\Rightarrow$ path. Every rainbow path is a rainbow tree.

Path $\Rightarrow$ joint. For a t -joint centered at v , remove an edge e and complete the remaining joint to $Q' \in X(t-1)$ by induction. Append e to a rainbow path of length $t-1$ in Q' ending at v , and complete it to $Q \in X(t)$. Its face through v in the directions of Q' equals Q' by path completion, so Q contains the joint. Any member of $X(t)$ containing the joint contains Q' by induction, hence this path, and therefore equals Q . $\square$

Definition 2.7 ((k, r) -cubical complex). *A decorated (k, r) -cubical complex is a decorated k -cube complex satisfying any of the equivalent r -completion properties. We subsequently omit the word decorated. A $(k, 2)$ -cubical complex is called a k -dimensional square complex.*

The left–right Cayley complex. A standard example of a $(2, 2)$ -cubical complex is obtained from a finite group Γ and symmetric subsets $A, B \subseteq \Gamma$, denoted $X_{A, \Gamma, B}$. For $I \subseteq \{1, 2\}$, put $V_I = \Gamma \times \{I\}$ and write $g_I = (g, I)$. Direction one edges correspond to left multiplication by elements of A ; direction two edges correspond to right multiplication by elements of B . The squares in $X_{A, \Gamma, B}(2)$ are precisely the squares

$$Q(g, a, b) = \{g_\emptyset, (ag)_{\{1\}}, (gb)_{\{2\}}, (agb)_{\{1, 2\}}\}, \quad g \in \Gamma, a \in A, b \in B.$$

It satisfies the completion axiom. Indeed, the joint at $g_\emptyset$ with neighbors $(ag)_{\{1\}}$ and $(gb)_{\{2\}}$ has the unique completion $Q(g, a, b)$; joints at the other vertex types are treated similarly. See [Figure 1](#) for an illustration.

Remark 2.8. *An embedded cube $Q \in \text{Cu}_X(i)$ need not be in $X(i)$. For example, when $A = B = \Gamma$, $\text{Cu}_{X_{A, \Gamma, B}}(2)$ contains $|\Gamma|^4$ four-cycles with one vertex in each block, whereas only the $|\Gamma|^3$ squares $Q(g, a, b)$ are in $X_{A, \Gamma, B}(2)$. Thus our definition does not impose a “flag” condition requiring every embedded cube to be declared.*

Remark 2.9. *One may also impose stronger uniqueness properties: whenever $x \in V_I$ and $y \in V_J$ are joined by a rainbow path of length $d = |I \triangle J|$, they belong to a unique declared d -cube. This requires uniqueness from the endpoints alone. The LPS-based cubical complexes constructed in [HLM⁺25, Lemma 3.3 and Theorem 3.5] satisfy this condition, but left–right Cayley complexes need not, as depicted in [Figure 1](#).*

2.2 The level chain complex

Let X be a $(k, 2)$ -cubical complex. Our chains and cochains are functions on the vertices of X , rather than on its cubical faces. Taking $\emptyset$ as the origin of the model cube, we write $V_s = \bigsqcup_{|I|=s} V_I$ for the s -th level, $0 \leq s \leq k$. Unless specified otherwise, subscripts on vertex sets and functions indicate restriction to the corresponding type or level. Thus, for $A \subseteq X(0)$, we write $A_I = A \cap V_I$ and $A_s = A \cap V_s$; similarly, $f_I = f|_{V_I}$ and $f_s = f|_{V_s}$. When viewing a scalar-valued function on a subset of vertices as an ambient function, we extend it by zero.

For $i \notin I$, orient every edge between V_I and $V_{I \cup \{i\}}$ toward $V_{I \cup \{i\}}$. Write $x \triangleleft y$ when an edge is oriented from x to y , and $x \leq y$ when $x = y$ or there is an upward path from x to y . If b, b' are neighbors of a in distinct directions, denote the fourth vertex of their unique square completion by $b \overset{a}{\wedge} b'$. In particular, if $b, b' \triangleleft a$, then $b \overset{a}{\wedge} b' \triangleleft b, b'$. For readability, we write $\{a; b, b'; c\}$ for a square with $a \in V_i$, $b, b' \in V_{i-1}$, and $c \in V_{i-2}$.

For $0 \leq s \leq k$, define $C_s(X) = C^s(X) = \mathbb{F}_2^{V_s}$. The boundary map $\partial_s : C_s(X) \rightarrow C_{s-1}(X)$ and the coboundary map $\delta_s : C^s(X) \rightarrow C^{s+1}(X)$ are given by

$$(\partial_s f)(u) = \sum_{u \triangleleft v} f(v), \quad (\delta_s f)(w) = \sum_{v \triangleleft w} f(v),$$

for $1 \leq s \leq k$ and $0 \leq s < k$, respectively. All sums are over $\mathbb{F}_2$, and $\partial_s = \delta_{s-1}^T$. The following lemma shows that square completion gives the identities $\partial^2 = \delta^2 = 0$. The following lemma and its proof are from [LH22a, LH22b].

Lemma 2.10. *For every $(k, 2)$ -cubical complex X and $1 \leq s \leq k-1$,*

$$\partial_s \partial_{s+1} = 0, \quad \delta_s \delta_{s-1} = 0.$$

Proof. Fix $x \in V_{s+1}$ and $z \in V_{s-1}$. The (z, x) -entry of $\partial_s \partial_{s+1}$ counts, modulo two, the vertices $y \in V_s$ satisfying $z \triangleleft y \triangleleft x$. For each such y , the unique square completing along this path has another middle vertex

$$y' = z \overset{y}{\wedge} x.$$

Observe that completing $z \triangleleft y' \triangleleft x$ gives the same square, i.e. $z \overset{y'}{\wedge} x = y$, so $y \mapsto y'$ is a fixed-point-free involution, since $\text{type}(y) \neq \text{type}(y')$. Therefore the number of vertices $y \in V_s$ satisfying $z \triangleleft y \triangleleft x$ is even, proving $\partial_s \partial_{s+1} = 0$. Taking transposes yields $\delta_s \delta_{s-1} = 0$. $\square$

For $I \subseteq [k]$ and $i \notin I$, let $X_{I,i}$ be the bipartite graph between V_I and $V_{I \cup \{i\}}$. We assume throughout that $X_{I,i}$ is (D_i, D'_i) -biregular, with $D_i, D'_i > 0$: vertices in V_I have degree D_i , and vertices in $V_{I \cup \{i\}}$ have degree D'_i . In particular, throughout, degrees depend only on the direction i , but not on I .

Norms. For $I \subseteq [k]$, put $D_I = \prod_{i \in I} D_i$ and $D'_I = \prod_{i \in I} D'_i$, with $D_\emptyset = D'_\emptyset = 1$. Assign each vertex $v \in V_I$ weight $\|v\|_X = \|I\|_X = D_I^{-1}$. For $f : X(0) \rightarrow \mathbb{R}$, define

$$\|f\|_X = \sum_{I \subseteq [k]} \sum_{v \in V_I} \frac{|f(v)|}{D_I}. \quad (2.1)$$

When the ambient complex is clear from context we omit the subscript. For an $\mathbb{F}_2$ -valued chain or cochain f , its norm is the total weight of its support: $\|f\| = \sum_{v \in \text{supp } f} \|v\|$. We identify a set with its indicator function, writing $\|H\| = \|\mathbf{1}_H\|$. With addition over $\mathbb{F}_2$, we have $\|f + g\| \leq \|f\| + \|g\|$. This norm extends to k -partite simplicial complexes. A face of type I is assigned the norm D_I^{-1} , and the norm of a cochain is the total weight of its support. Throughout, we use the normalized distance. For a cochain $f_r \in C^r$ and $C \subseteq C^r(X)$, write

$$\text{dist}_X(f_r, C) := \min_{c \in C} \|f_r + c\|, \quad \text{dist}_X(f_r) = \text{dist}(f_r, \text{im}(\delta^{r-1}))$$

Here too, the subscript is omitted when the ambient complex is clear from context.

Remark 2.11. *For every $f : X(0) \rightarrow \mathbb{F}_2$,*

$$\frac{|\text{supp } f|}{D_{[k]}} \leq \|f\| \leq |\text{supp } f|.$$

Thus the weighted norm and Hamming weight are equivalent up to a factor depending only on the degrees.

For $v \in V_I$ and $i \notin I$, the total weight of its neighbors in direction i is $D_i/D_{I \cup \{i\}} = \|v\|$. Summing over the upward directions gives

$$\|\delta_{|I|} \mathbf{1}_{\{v\}}\| = \sum_{v \prec w} \|w\| = \sum_{i \notin I} \frac{D_i}{D_{I \cup \{i\}}} = (k - |I|) \|v\|. \quad (2.2)$$

Finally, choosing $[k]$ as the origin reverses every edge orientation. The resulting *reversed complex* X^{rev} has the same vertices, edges, and declared cubes, with $V_I^{\text{rev}} = V_{[k] \setminus I}$. The following lemma relates its coboundary maps to the boundary maps of X .

Lemma 2.12. *For $0 \leq s \leq k$, the identity on functions identifies $C_s(X)$ with $C^{k-s}(X^{\text{rev}})$. Under this identification,*

$$\partial_s^X = \delta_{k-s}^{X^{\text{rev}}} \quad (1 \leq s \leq k).$$

Moreover, $D_i^{\text{rev}} = D'_i$ and $(D'_i)^{\text{rev}} = D_i$.

Proof. By definition, $V_s(X) = V_{k-s}(X^{\text{rev}})$, and $x \prec_X y$ if and only if $y \prec_{X^{\text{rev}}} x$. Thus the corresponding boundary and coboundary maps sum over the same neighbors. Reversal also exchanges the two sides of every directional graph, giving the degree identities. $\square$

This identification need not preserve norms. For $f \in C_s(X)$, the reversed norm is

$$\|f\|_{X^{\text{rev}}} = \sum_{\substack{I \subseteq [k] \\ |I|=s}} \sum_{v \in V_I} \frac{|f(v)|}{D'_{[k] \setminus I}}.$$

2.3 Links

Definition 2.13 (Link). *Let $v \in V_I$. For $J \subseteq [k] \setminus I$, define*

$$\text{Lk}_J(v) = \{x \in V_{I \cup J} : v \leq x\}.$$

The link of v is $\text{Lk}(v) = \{x \in X(0) : v \leq x\}$, with levels

$$\text{Lk}_s(v) = \bigsqcup_{\substack{J \subseteq [k] \setminus I \\ |J|=s+1}} \text{Lk}_J(v), \quad -1 \leq s \leq k - |I| - 1.$$

In particular, $\text{Lk}_{-1}(v) = \text{Lk}_\emptyset(v) = \{v\}$.

The chain and cochain groups of the link are $C_s(\text{Lk}(v)) = C^s(\text{Lk}(v)) = \mathbb{F}_2^{\text{Lk}_s(v)}$. Their boundary and coboundary maps use the same incidence formulas as above, with sums restricted to vertices in the link. Thus δ_{-1} sends a value at v to the constant cochain with that value on $\text{Lk}_0(v)$.

For $h \in C^s(\text{Lk}(v))$, let $\iota_v h \in C^{|I|+s+1}(X)$ be its extension by zero. The link norm is

$$\|h\|_v = \sum_{\substack{J \subseteq [k] \setminus I \\ |J|=s+1}} \sum_{x \in \text{Lk}_J(v)} \frac{|h(x)|}{D_J}. \quad (2.3)$$

Since $D_{I \cup J} = D_I D_J$, we have $\|\iota_v h\|_X = D_I^{-1} \|h\|_v$.

Lemma 2.14. *If X is a $(k, 2)$ -cubical complex, then the restricted link maps satisfy $\partial^2 = 0$ and $\delta^2 = 0$.*

Proof. The involution of [Lemma 2.10](#) preserves the link. $\square$

2.4 Small-set coboundary expansion

Definition 2.15 (Minimality). *For $1 \leq r \leq k$, a cochain $f \in C^r(X)$ is minimal if*

$$\|f\| \leq \|f + \delta_{r-1}h\| \quad \text{for every } h \in C^{r-1}(X).$$

It is locally minimal if, for every $0 \leq s \leq r-1$, every $v \in V_s$, and every $h \in C^{r-s-2}(\text{Lk}(v))$,

$$\|f\| \leq \|f + \delta_{r-1}(\iota_v h)\|.$$

All cochains in $C^0(X)$ are minimal.

Every minimal cochain is locally minimal, since each $\iota_v h$ in the definition belongs to $C^{r-1}(X)$.

Definition 2.16 (Small-set coboundary expansion). *For $0 \leq r \leq k-1$, a cubical complex X is an (α, β) -small-set coboundary expander at level r if every locally minimal $f \in C^r(X)$ satisfying $\|f\| \leq \alpha$ also satisfies $\|\delta_r f\| \geq \beta \|f\|$.*

2.5 Lossless expansion

For a graph H , write $\Gamma(v)$ for the neighborhood of v and $\Gamma(A) = \bigcup_{v \in A} \Gamma(v)$. For $v \in V_I$ and $i \in [k]$, let $\Gamma_i(v)$ denote the neighbors of v in direction i . Put $\Gamma_i(A) = \bigcup_{v \in A} \Gamma_i(v)$. Thus $\Gamma_i(v) \subseteq V_{I \Delta \{i\}}$.

Definition 2.17 (One-sided lossless bigraph). *Let $H = (L, R, E)$ be a (D, D') -biregular bigraph. For $c, \epsilon \in (0, 1)$, we say that H is one-sided (c, ϵ) -lossless from L to R if every $A \subseteq L$ with $|A| \leq c|L|$ satisfies*

$$|\Gamma(A)| \geq (1 - \epsilon)D|A|.$$

Definition 2.18 (Lossless cubical complex). *A cubical complex X is lossless with parameters $(D_i, D'_i, c_i, \epsilon_i)_{i=1}^k$ if, for every $I \subseteq [k]$ and $i \notin I$, the directional graph $X_{I,i}$ is one-sided (c_i, ϵ_i) -lossless from V_I to $V_{I \cup \{i\}}$. It is two-sided lossless if X^{rev} is also lossless, with D_i, D'_i interchanged and the same c_i, ϵ_i .*

For a bigraph $H = (L, R, E)$, let $e(A, B)$ denote the number of edges between $A \subseteq L$ and $B \subseteq R$. For a vertex v and a vertex set S , put $\deg_S(v) = |\Gamma(v) \cap S|$. Denote $t_+ = \max\{t, 0\}$.

Lemma 2.19. *Let $H = (L, R, E)$ be a one-sided (c, ϵ) -lossless (D, D') -biregular bigraph.*

1. **Excess bound.** *If $A \subseteq L$ satisfies $|A| \leq c|L|$, then*

$$\sum_{y \in R} (\deg_A(y) - 1)_+ = D|A| - |\Gamma(A)| \leq \epsilon D|A|.$$

2. **Lossless edge bound.** *If $A \subseteq L$ satisfies $|A| \leq c|L|$, then, for every $B \subseteq R$,*

$$e(A, B) \leq |B| + \epsilon D|A|.$$

3. **Large-degree bound.** *For $B \subseteq R$, suppose that at most $c|L|$ vertices $x \in L$ satisfy $\deg_B(x) > \epsilon D$. Then*

$$\sum_{x \in L} \left(\frac{\deg_B(x)}{D} - \epsilon \right)_+ \leq \frac{|B|}{D}.$$

Proof. Excess bound (1). Fix $A \subseteq L$ with $|A| \leq c|L|$. Every vertex in $\Gamma(A)$ has at least one neighbor in A , and vertices outside $\Gamma(A)$ have none. Consequently, by regularity and lossless expansion,

$$\sum_{y \in R} (\deg_A(y) - 1)_+ = \sum_{y \in \Gamma(A)} (\deg_A(y) - 1) = D|A| - |\Gamma(A)| \leq \varepsilon D|A|.$$

Lossless edge bound (2). Fix $A \subseteq L$ with $|A| \leq c|L|$ and $B \subseteq R$. At each vertex of $B \cap \Gamma(A)$, count one incident edge from A separately from the remaining edges. This gives

$$e(A, B) = |B \cap \Gamma(A)| + \sum_{y \in B} (\deg_A(y) - 1)_+ \leq |B| + \sum_{y \in R} (\deg_A(y) - 1)_+ \leq |B| + \varepsilon D|A|.$$

The last inequality follows from part (1).

Large-degree bound (3). Fix $B \subseteq R$ as in the statement, and put $A = \{x \in L : \deg_B(x) > \varepsilon D\}$, so $|A| \leq c|L|$. The summand $(\deg_B(x)/D - \varepsilon)_+$ is positive precisely when $x \in A$. Therefore,

$$\sum_{x \in L} \left(\frac{\deg_B(x)}{D} - \varepsilon \right)_+ = \sum_{x \in A} \left(\frac{\deg_B(x)}{D} - \varepsilon \right) = \frac{e(A, B)}{D} - \varepsilon|A|.$$

The bound $|A| \leq c|L|$ allows us to apply part (2), obtaining

$$e(A, B) \leq |B| + \varepsilon D|A|.$$

Substituting this bound into the preceding identity gives the required upper bound $|B|/D$. $\square$

3 Expansion of the links

In this section we show that links in lossless cubical complexes can be approximated by complete multipartite simplicial complexes. Recall that we define the norm on such complexes similarly to the norm on cubical complexes. If the parts have sizes D_i for $i \in [d]$, then a face of type I has weight D_I^{-1} , and the norm of a cochain is the sum of the weights of the faces in its support. Under this norm, complete d -partite complexes are known to expand minimal cochains at every level $-1 \leq r \leq d-2$, with a constant factor $m_{d,r}$ depending on d and r but not on the part sizes:

$$m_{d,r} := \left(\binom{d}{r+1} \binom{d+r+1}{r+2} \right)^{-1}. \quad (3.1)$$

Theorem 3.1 ([LMM16, Corollary 3.1]). *Let $\mathbf{M} = K(D_1, \dots, D_d)$ be the complete d -partite complex with block sizes $D_1, \dots, D_d$. For $-1 \leq r \leq d-2$, every $f \in C^r(\mathbf{M})$ satisfies*

$$\|\delta_r f\| \geq m_{d,r} \text{dist}(f).$$

We prove that lossless square complexes inherit this expansion in their links, up to an additive error depending on the loss.

Theorem 3.2. *Let X be a k -dimensional lossless square complex with parameters $(D_i, D'_i, c_i, \varepsilon_i)_{i=1}^k$. Fix $I \subseteq [k]$ with $|I| < k$ and a vertex $v \in V_I$, and put $d = k - |I|$. Then, for every $-1 \leq r \leq d - 2$ and every cochain $f : \text{Lk}_r(v) \rightarrow \mathbb{F}_2$,*

$$\|\delta_r f\|_v \geq m_{d,r} \text{dist}(f) - O_k \left(\sum_{j=1}^k \varepsilon_j \right). \quad (3.2)$$

To simplify notations, we first treat the link of a vertex $v \in V_\emptyset$ and reduce to this case in the end.

3.1 Ordered paths and collisions

Let $v \in V_\emptyset$ and let $I \subseteq [k]$, with $|I| = i$. Consider the set of monotone paths from v to $\text{Lk}_I(v)$,

$$\mathcal{P} := \{(x_0, \dots, x_i) : x_0 = v, x_i \in \text{Lk}_I(v), x_{j-1} \leq x_j \text{ for every } j \in [i]\}.$$

Every path in $\mathcal{P}$ uses each direction in I exactly once, though the order can vary. Define

$$\begin{aligned} \text{Ord} : \mathcal{P} &\longrightarrow \{\pi : [i] \rightarrow I\}, & \{\text{Ord}((x_0, \dots, x_i))(j)\} &:= \text{type}(x_j) \setminus \text{type}(x_{j-1}), \\ \text{End} : \mathcal{P} &\longrightarrow \text{Lk}_I(v), & \text{End}((x_0, \dots, x_i)) &:= x_i. \end{aligned}$$

That is, Ord records the order in which the directions appear along the path, and End stands for the other endpoint of the path. For an ordering π of I , let

$$\mathcal{P}_\pi := \{P \in \mathcal{P} : \text{Ord}(P) = \pi\}, \quad \text{End}_\pi := \text{End}|_{\mathcal{P}_\pi}.$$

At the j th step of a path in $\mathcal{P}_\pi$ there are $D_{\pi(j)}$ choices, so

$$|\mathcal{P}_\pi| = \prod_{j=1}^i D_{\pi(j)} = D_I. \quad (3.3)$$

Lemma 3.3 (Switching ordered paths). *For every $x \in \text{Lk}_I(v)$, the fiber size $|\text{End}_\pi^{-1}(x)|$ is independent of the ordering π .*




Figure 2: The link above v in a three-dimensional cubical complex. The solid path P has order $(1, 2, 3)$, while the dashed path P' has order $(2, 1, 3)$. The pale shaded quadrilateral marks the square 2-cell witnessing the adjacent switch.

Proof. Given $P = (x_0, \dots, x_i) \in \text{End}_\pi^{-1}(x)$, fix $1 \leq j < i$. The segment $x_{j-1} \triangleleft x_j \triangleleft x_{j+1}$ is a two-edge joint at x_j . The square-completion axiom supplies a unique square with other intermediate vertex $x'_j = x_{j-1} \wedge^{x_j} x_{j+1}$. Since $v \leq x_{j-1} \triangleleft x'_j$, this vertex also belongs to $\text{Lk}(v)$. Replacing x_j by x'_j gives the path

$$(x_0, \dots, x_{j-1}, x'_j, x_{j+1}, \dots, x_i),$$

which has the same endpoint and swaps the j th and $(j+1)$ st directions, as illustrated in Figure 2. Repeating the switch recovers the original path. It therefore gives a bijection between $\text{End}_\pi^{-1}(x)$ and $\text{End}_{\pi \circ \sigma_{j,j+1}}^{-1}(x)$, where $\sigma_{j,j+1}$ is the adjacent transposition. Since adjacent transpositions generate the symmetric group, all fibers have the same size. $\square$

Corollary 3.4. *For every ordering π of I , the endpoint map End_π is surjective.*

Proof. Let $x \in \text{Lk}_I(v)$. Some monotone path P ends at x ; if $\rho = \text{Ord}(P)$, then $|\text{End}_\rho^{-1}(x)| \geq 1$. By Lemma 3.3, the same holds for every ordering π . $\square$

Corollary 3.5. *For every $I \subseteq [k]$ and every ordering π of I ,*

$$\|\text{Lk}_I(v)\| \leq \frac{1}{D_I} |\mathcal{P}_\pi| = 1.$$

Lossless expansion gives the reverse inequality up to the loss. Fix an ordering $\pi(j) = s_j$ and put $I_t := \{s_1, \dots, s_t\}$, with $I_0 = \emptyset$. By Corollary 3.4, the endpoints of all length- t paths with ordering $\pi|_{[t]}$ are precisely $\text{Lk}_{I_t}(v)$, and hence

$$\text{Lk}_{I_t}(v) = \Gamma_{s_t}(\text{Lk}_{I_{t-1}}(v)).$$

Successive lossless expansion gives

$$\|\text{Lk}_I(v)\| \geq \prod_{j \in I} (1 - \varepsilon_j) \geq 1 - \sum_{j \in I} \varepsilon_j. \quad (3.4)$$

For each type I , fix a reference ordering and define, for $x \in \text{Lk}_I(v)$,

$$\mu_I(x) := |\text{End}_\pi^{-1}(x)| - 1, \quad B_I := \{x \in \text{Lk}_I(v) : \mu_I(x) > 0\}.$$

The function μ_I measures the endpoint excess and B_I is the set of bad, or colliding, endpoints. By (3.3) and Corollary 3.5

$$\|\mu_I\| = \frac{1}{D_I} \sum_{x \in \text{Lk}_I(v)} (|\text{End}_\pi^{-1}(x)| - 1) = \frac{|\mathcal{P}_\pi| - |\text{Lk}_I(v)|}{D_I} = 1 - \|\text{Lk}_I(v)\| \quad (3.5)$$

Since the indicator of B_I is pointwise at most μ_I , and using (3.4), we get

$$\|B_I\| \leq \|\mu_I\| \leq \sum_{j \in I} \varepsilon_j. \quad (3.6)$$

Put $B := \bigcup_I B_I$. The set B is upward closed. Indeed, two distinct paths with the same ordering ending at x remain distinct after appending any fixed monotone path from x to an ancestor y . By Lemma 3.3, a duplicated fiber for one ordering is duplicated for every ordering, including the reference ordering used to define the bad set containing y . Thus $x \in B$ and $x \leq y$ imply $y \in B$. Equivalently, defining the non-colliding set by $G := \text{Lk}(v) \setminus B$, we obtain that G is downward closed within $\text{Lk}(v)$.

3.2 Embedding the non-colliding part

Let $\mathbf{M}$ be the complete k -partite simplicial complex whose j th part is $\text{Lk}_{\{j\}}(v) = \Gamma_j(v)$.

Claim 3.6. *For every $x_I \in G$, the interval $[v, x_I]$ ⁵ satisfies the following:*

- (1) *For every $J \subseteq I$, there is a unique vertex $x_J \in [v, x_I] \cap \text{Lk}_J(v)$.*
- (2) *For every $J, K \subseteq I$, we have $x_J \leq x_K$ if and only if $J \subseteq K$. In particular, the interval is an $|I|$ -dimensional Boolean cube.*
- (3) *Every four-cycle in this interval is a square in $X(2)$.*

Proof. (1) Fix $J \subseteq I$ and an ordering π of I that lists J first. By Corollary 3.4, there is a path P of order π from v to x_I . Let x_J be its vertex after $|J|$ steps. Any $y \in [v, x_I] \cap \text{Lk}_J(v)$ lies on a monotone path from v to x_I . Applying the square-switch argument of Lemma 3.3 separately before and after y gives a path of order π through y . Since $x_I \in G$, this path must be P , so $y = x_J$.

(2) Suppose $J \subseteq K \subseteq I$. Choose an ordering of I that lists first J , then $K \setminus J$, and finally $I \setminus K$. By Corollary 3.4, there is a path of this order from v to x_I . The vertices reached after $|J|$ and $|K|$ steps have types J and K , respectively. Both lie in $[v, x_I]$, so by (1) they are x_J and x_K .

⁵As usual for a poset, $[v, x_I] = \{y : v \leq y \leq x_I\}$.

The part of the path from x_J to x_K gives $x_J \leq x_K$. Conversely, if $x_J \leq x_K$, an upward path from x_J to x_K can only add directions to the type, so $J \subseteq K$. By (1), the vertices of the interval are precisely $\{x_J : J \subseteq I\}$. The equivalence just proved identifies their order with inclusion of subsets of I , giving an $|I|$ -dimensional Boolean cube.

(3) By (2), every four-cycle in $[v, x_I]$ is a two-dimensional face of the Boolean cube. Thus, for some $J \subseteq I$ and distinct $j_1, j_2 \in I \setminus J$, its vertices can be labeled

$$\begin{aligned} a &= x_{J \cup \{j_1, j_2\}}, & b &= x_{J \cup \{j_1\}}, \\ b' &= x_{J \cup \{j_2\}}, & c &= x_J. \end{aligned}$$

Complete the path $c \leq b \leq a$ to a square in $X(2)$, and let b'' be its fourth vertex. The path through b uses directions j_1, j_2 , so the other path $c \leq b'' \leq a$ uses directions j_2, j_1 . Hence b'' has type $J \cup \{j_2\}$, the same as b' . Also, $v \leq c \leq b'' \leq a \leq x_I$, so $b'' \in [v, x_I]$. By the uniqueness in (1), $b'' = b'$, and hence $\{a, b, b', c\} \in X(2)$. $\square$

For $x_I \in G$, define

$$\varphi(x_I) := \{x_{\{j\}} \in \text{Lk}_{\{j\}}(v) : x_{\{j\}} \leq x_I, j \in I\}.$$

By [Claim 3.6](#), this selects exactly one vertex from each part indexed by I , so $\varphi(x_I)$ is a face of $\mathbf{M}$.




Figure 3: For a non-colliding vertex x_{123} , the interval below it is a Boolean cube. The map φ identifies it with the face poset of the simplex on its three level-one vertices, with v corresponding to the empty face.

Lemma 3.7. *The map $\varphi : G \rightarrow \mathbf{M}$ is a type-preserving poset isomorphism between G and its image.*

Proof. We need to show that φ preserves types, is injective on G , and preserves the order in both directions.

(1) *Type preservation.* If $x_I \in G$ has type I , then $\varphi(x_I)$ contains exactly one vertex from each part indexed by I and no other vertices. Thus $\varphi(x_I)$ also has type I .

(2) *Injectivity on G .* Suppose that two distinct vertices have the same image. By (1), they have the same type I ; write them as $x_I, x'_I \in G$. For every $J \subseteq I$, let x_J and x'_J be the unique type- J vertices in $[v, x_I]$ and $[v, x'_I]$, respectively, given by [Claim 3.6](#). Choose J of minimum cardinality for which $x_J \neq x'_J$. The two intervals agree at v and, since the images are equal, at every singleton type, so $|J| \geq 2$. Choose distinct $j, j' \in J$ and put

$$\begin{aligned} b &:= x_{J \setminus \{j\}} = x'_{J \setminus \{j\}}, \\ b' &:= x_{J \setminus \{j'\}} = x'_{J \setminus \{j'\}}, \\ c &:= x_{J \setminus \{j, j'\}} = x'_{J \setminus \{j, j'\}}, \end{aligned}$$

where the equalities follow from the choice of J . By [Claim 3.6](#), both $\{x_J; b, b'; c\}$ and $\{x'_J; b, b'; c\}$ are squares in $X(2)$. They complete the same joint formed by the edges cb and cb' , so uniqueness of square completion gives $x_J = x'_J$, a contradiction. Thus φ is injective on G .

(3) *Order preservation in both directions.* If $x_J \leq x_I$, every level-one vertex below x_J also lies below x_I , so $\varphi(x_J) \subseteq \varphi(x_I)$. Conversely, suppose $\varphi(x_J) \subseteq \varphi(x_I)$. By (1), $J \subseteq I$. Let $x'_J \in [v, x_I]$ be the unique type- J vertex given by [Claim 3.6](#). Since G is downward closed, $x'_J \in G$. By the forward implication and (1), both $\varphi(x_J)$ and $\varphi(x'_J)$ are contained in $\varphi(x_I)$ and have type J . Since $\varphi(x_I)$ contains at most one vertex in each part, these two faces are equal. By (2), $x_J = x'_J \leq x_I$. Therefore φ is a type-preserving poset isomorphism onto its image, as depicted in [Figure 3](#). $\square$

3.3 Transfer of multipartite expansion

Theorem 3.8. *Let X be a k -dimensional square complex and let $v \in V_\emptyset$. Fix $0 \leq r \leq k - 2$. If $\|\text{Lk}_{[k]}(v)\| \geq 1 - \eta$, then every cochain $f_r : \text{Lk}_r(v) \rightarrow \mathbb{F}_2$ satisfies*

$$\|\delta f_r\| \geq m_{k,r} \text{dist}(f_r) - O_k(\eta). \quad (3.7)$$

Proof. We start by propagating the norm assumption on the top part of the link to all its parts. Fix $I \subseteq [k]$, $j \in I$, and put $J := I \setminus \{j\}$. Choose an ordering of I whose first $|J|$ directions are the elements of J . By [Corollary 3.4](#),

$$\text{Lk}_I(v) = \Gamma_j(\text{Lk}_J(v)).$$

Since every vertex has D_j upward neighbours in direction j ,

$$|\text{Lk}_I(v)| \leq D_j |\text{Lk}_J(v)|.$$

Dividing by $D_I = D_J D_j$, we get that $\|\text{Lk}_I(v)\| \leq \|\text{Lk}_J(v)\|$. Iterating from $I = [k]$ down to an arbitrary type gives

$$\|\text{Lk}_I(v)\| \geq \|\text{Lk}_{[k]}(v)\| \geq 1 - \eta. \quad (3.8)$$

For $I \subseteq [k]$, let B_I and G_I be the bad and good vertices of type I , respectively, and set

$$\widehat{G}_I := \varphi(G_I) \subseteq \mathbf{M}(I), \quad \widehat{B}_I := \mathbf{M}(I) \setminus \widehat{G}_I.$$

The exact endpoint count in [\(3.5\)](#) gives

$$\|B_I\| \leq \|\mu_I\| = 1 - \|\text{Lk}_I(v)\| \leq \eta.$$

The map φ is type preserving, so it preserves the norm on G_I . Thus $\|G_I\| = \|\varphi(G_I)\|$, and since $\|\mathbf{1}_{\mathbf{M}(I)}\| = 1$, we also have

$$\|\widehat{B}_I\| = 1 - \|\widehat{G}_I\| = (1 - \|\mathbf{1}_{\text{Lk}_I(v)}\|) + \|B_I\| = 1 - \|\text{Lk}_I(v)\| + \|B_I\| \leq 2\eta.$$

For $-1 \leq s \leq k - 1$, let $B_s, G_s, \widehat{B}_s, \widehat{G}_s$ be the corresponding unions over all types of size $s + 1$. Thus $B_s, G_s \subseteq \text{Lk}_s(v)$ and $\widehat{B}_s, \widehat{G}_s \subseteq \mathbf{M}(s)$, and we obtain

$$\|B_s\| \leq \binom{k}{s+1} \eta, \quad \|\widehat{B}_s\| \leq 2 \binom{k}{s+1} \eta. \quad (3.9)$$

Let $\widehat{h}_{r-1} : \mathbf{M}(r-1) \rightarrow \mathbb{F}_2$. Define the pushforward $\widehat{f}_r : \mathbf{M}(r) \rightarrow \mathbb{F}_2$ of f_r and the pullback $h_{r-1} : \text{Lk}_{r-1}(v) \rightarrow \mathbb{F}_2$ of $\widehat{h}_{r-1}$ by

$$\widehat{f}_r(y) := \begin{cases} f_r(\varphi^{-1}(y)), & y \in \widehat{G}_r, \\ 0, & y \in \widehat{B}_r. \end{cases}, \quad h_{r-1}(x) := \begin{cases} \widehat{h}_{r-1}(\varphi(x)), & x \in G_{r-1}, \\ 0, & x \in B_{r-1}. \end{cases}$$

We now compare $\text{dist}_{\mathbf{M}} \widehat{f}_r$ with $\text{dist}_{\text{Lk}(v)}(f_r)$. If $x \in G_r$, then all children of x in $\text{Lk}(v)$ are good, and [Claim 3.6](#) and [Lemma 3.7](#) show that φ bijects them with the children of $\varphi(x)$. Thus, for $x \in G$ and for $\widehat{\delta}$ the coboundary map on $\mathbf{M}$,

$$\widehat{\delta} \widehat{h}_{r-1}(\varphi(x)) = \delta h_{r-1}(x).$$

Consequently,

$$\begin{aligned} \|\widehat{f}_r + \widehat{\delta} \widehat{h}_{r-1}\| &\geq \left\| \left(\widehat{f}_r + \widehat{\delta} \widehat{h}_{r-1} \right) \Big|_{\widehat{G}_r} \right\| \\ &= \|(f_r + \delta h_{r-1})|_{G_r}\| \\ &\geq \|f_r + \delta h_{r-1}\| - \|B_r\| \\ &\geq \text{dist}(f_r) - \|B_r\|. \end{aligned}$$

Taking the minimum over $\widehat{h}_{r-1}$ gives

$$\text{dist}_{\mathbf{M}}(\widehat{f}_r) \geq \text{dist}(f_r) - \|B_r\|. \quad (3.10)$$

Therefore [Theorem 3.1](#) implies

$$\|\widehat{\delta} \widehat{f}_r\| \geq m_{k,r} (\text{dist}(f_r) - \|B_r\|). \quad (3.11)$$

Likewise, if $x \in G_{r+1}$, then its children in $\text{Lk}(v)$ are good and correspond under φ to the children of $\varphi(x)$, so

$$\widehat{\delta} \widehat{f}_r(\varphi(x)) = \delta f_r(x).$$

Hence,

$$\|\delta f_r\| \geq \|(\delta f_r)|_{G_{r+1}}\| = \left\| \left(\widehat{\delta} \widehat{f}_r \right) \Big|_{\widehat{G}_{r+1}} \right\| \geq \|\widehat{\delta} \widehat{f}_r\| - \|\widehat{B}_{r+1}\|$$

Combining this with (3.11) and then using (3.9), we get

$$\begin{aligned} \|\delta f_r\| &\geq m_{k,r} \text{dist}(f_r) - m_{k,r} \|B_r\| - \|\widehat{B}_{r+1}\| \\ &\geq m_{k,r} \text{dist}(f_r) - \left(m_{k,r} \binom{k}{r+1} + 2 \binom{k}{r+2} \right) \eta \\ &= m_{k,r} \text{dist}(f_r) - O_k(\eta), \end{aligned}$$

as claimed. □

Proof of Theorem 3.2. Let $I \subseteq [k]$ and let $v \in V_I$. Put $d = k - |I|$. Consider the sub-complex X^I whose relative type- J part is

$$V_J^I := V_{I \cup J}, \quad J \subseteq [k] \setminus I.$$

This is a d -dimensional square complex, and v lies in its relative empty part. The link of v and its levels are unchanged, and the norm in the relative slice, restricted to this link, agrees with $\|\cdot\|_v$.

Suppose first that $0 \leq r \leq d - 2$. Applying Theorem 3.8 to X^I , with dimension d and link level r , together with (3.4) gives

$$\|\delta f\|_v \geq m_{d,r} \operatorname{dist}_{\operatorname{Lk}(v)}(f) - O_k \left(\sum_{j \in [k] \setminus I} \varepsilon_j \right). \quad (3.12)$$

Since the sum over $[k] \setminus I$ is at most $\sum_{j=1}^k \varepsilon_j$ this yields (3.2).

For link level -1 , the cochain f is just a value at v , giving

$$\|\delta f\|_v = d \|f\|_v \geq m_{d,-1} \|f\|_v.$$

□

4 Small-set coboundary expansion from square complexes

We establish small-set coboundary expansion at every level below the top, arising from imbalanced lossless square complexes. Throughout this section X is a k -dimensional square complex with parameters $(D_i, D'_i, c_i, \varepsilon_i)_{i=1}^k$. The proof follows the heavy-branch strategy of [FK24].

Throughout this section, we assume that every set $A \subseteq V_J$ to which lossless expansion is applied in direction $j \notin J$ satisfies $\|A\| \leq c_j \|V_J\|$. The choice of c in (4.26) ensures this whenever the initial cochain H_r satisfies $\|H_r\| \leq c \|V_r\|$.

4.1 Heavy and terminal vertices

Fix a level $0 \leq r \leq k - 1$ and a cochain,

$$H_r = \bigsqcup_{|I|=r} H_I \subseteq V_r.$$

We regard H_r as the initial set of *heavy* vertices. Following [FK24], we recursively pull H_r down through the complex. Suppose that the heavy vertices at level i have already been defined. We declare x at level $i - 1$ heavy, if it has many heavy parents. Concretely, for $x \in V_I$ and $j \notin I$, we denote the *direction j heavy degree* and *normalized heavy degree* of x by

$$d_{j, H_{I \cup \{j\}}}(x) := |\{y \in H_{I \cup \{j\}} : x \leq y\}|, \quad u_{j, H_{I \cup \{j\}}}(x) := \frac{d_{j, H_{I \cup \{j\}}}(x)}{D_j}. \quad (4.1)$$

When I and $H_{I \cup \{j\}}$ are clear from context, we will omit $H_{I \cup \{j\}}$ from the notation. The construction uses thresholds $0 < \gamma_i < k - i$ for $0 \leq i < r$, which will be fixed later. For $|I| < r$, a vertex

$x \in V_I$ is declared heavy if its total heavy normalized degree is at least the threshold of its level, $\gamma_{|I|}$.

$$H_I := \left\{ x \in V_I : \sum_{j \notin I} u_j(x) \geq \gamma_{|I|} \right\}, \quad H_i := \bigsqcup_{|I|=i} H_I. \quad (4.2)$$

We also define heavy vertices in the level *above* the initial cochain,

$$H_{r+1} := \{v \in V_{r+1} : a \triangleleft v \text{ for some } a \in H_r\}, \quad (4.3)$$

Notice that H_{r+1} is not obtained from another threshold test - it is simply the full upward neighborhood of H_r . Thus every $a \in H_r$ has all its upward neighbors in H_{r+1} , and hence for every $j \notin \text{type}(a)$, $d_j(a) = D_j$ and $\sum_{j \notin \text{type}(a)} u_j(a) = k - r$. Let

$$H := \bigcup_{i=0}^{r+1} H_i$$

be the set of heavy vertices. All other vertices are called *light*, and are denoted by $L := X(0) \setminus H$. We view the heavy vertices as an oriented graph whose edges point downward to the level below. For $v \in H_i$, define its *heavy children* by

$$\text{Ch}(v) := \begin{cases} \{u \in H_{i-1} : u \triangleleft v\}, & i \geq 1, \\ \emptyset, & i = 0. \end{cases}$$

Definition 4.1 (Terminal vertices). *A heavy vertex $v \in H$ is terminal if $\text{Ch}(v) = \emptyset$. We write*

$$\text{Ter} := \{v \in H : \text{Ch}(v) = \emptyset\}.$$

Observe that by definition, every vertex in H_{r+1} has a heavy child, and so it cannot be terminal and

$$\text{Ter} \subseteq \bigcup_{i=0}^r H_i.$$

For every heavy vertex $v \in H$, recursively define its set of terminal descendants by

$$\text{Ter}(v) := \begin{cases} \{v\}, & v \in \text{Ter}, \\ \bigcup_{u \in \text{Ch}(v)} \text{Ter}(u), & v \notin \text{Ter}. \end{cases} \quad (4.4)$$

In words, we start at v , follow every possible downward heavy path, and record the terminal vertices at which those paths stop. From the other perspective, for a terminal vertex $t \in \text{Ter}$ we let its *terminal cluster* $H(t)$ be

$$H(t) := \{a \in H : t \in \text{Ter}(a)\}. \quad (4.5)$$

Equivalently, $a \in H(t)$ precisely when t is a descendant of a along a downward path of heavy vertices. The clusters need not be disjoint: a vertex a belongs to a cluster for each terminal in

$\text{Ter}(a)$, hence to exactly $|\text{Ter}(a)|$ terminal clusters. The quantity $|\text{Ter}(a)| - 1$ measures *excess* coverage by terminal clusters, we call it the *excess* of a and we denote it by $E(a)$, i.e.,

$$E(a) := |\text{Ter}(a)| - 1. \quad (4.6)$$

We observe that every heavy vertex has at least one terminal descendant, so that $E(v)$ is non-negative. $E(v) = 0$ exactly when v has a unique terminal descendant; every additional terminal endpoint contributes one unit of excess. We measure the total ambiguity at level i by

$$\|E_i\| = \sum_{a \in H_i} E(a) \|a\|.$$

These quantities play a central role in what follows. We refer the reader to [Figure 4](#) for an illustration of the setup described in this subsection. It turns out that in the small-set range, when the initial set H_r is small enough for lossless vertex expansion to apply, almost all vertices in H are assigned to a unique terminal vertex, and the norm of the excess function is small. Proving this is a main goal of this section.

4.2 The global to local reduction

To prove small-set coboundary expansion we need to show that $\|\delta H_r\|$ is large relative to $\|H_r\|$. We decompose the global function δH_r into a sum of local functions, associated with terminal vertices.

We now compare $\|\delta H_r\|$ to $\sum_{t \in \text{Ter}} \|\delta H_r(t)\|$. Notice that if

1. Every vertex in H_r has only one terminal (which, in turn, implies that vertices in H_r that have different terminals have disjoint heavy descendants at all levels below r), and,
2. Vertices in H_r that have different terminals have different ancestors at level $r + 1$

then $\|\delta H_r\| = \sum_{t \in \text{Ter}} \|\delta H_r(t)\|$. Thus, there are two reasons why $\|\delta H_r\|$ might not equal $\sum_{t \in \text{Ter}} \|\delta H_r(t)\|$. First, a vertex of H_r may belong to several terminal clusters. Second, even after this overcounting is corrected in level r , the cluster coboundaries may overlap and cancel in level $r + 1$. We prove:

Proposition 4.2 (Local-to-global comparison).

$$\|\delta H_r\| \geq \sum_{t \in \text{Ter}} \|\delta H_r(t)\| - (k - r) \|E_r\| - 2 \|E_{r+1}\|. \quad (4.7)$$

Proof. Notice that

$$H_r = \sum_{t \in \text{Ter}} H_r(t) + \sum_{a \in H_r} E(a) \{a\}.$$

where the sum is over $\mathbb{F}_2$, because every $a \in H_r$ is counted $|\text{Ter}(a)|$ times in the first argument, and $E(a) = |\text{Ter}(a)| - 1$ in the second argument.

We first deal with the problem that a vertex $a \in H_r$ may have several terminals.

$$\|\delta H_r\| \geq \left\| \delta \sum_{t \in \text{Ter}} H_r(t) \right\| - \sum_{a \in H_r} E(a) \|\delta \{a\}\| = \left\| \sum_{t \in \text{Ter}} \delta H_r(t) \right\| - (k - r) \|E_r\|, \quad (4.8)$$

where the first inequality is the triangle inequality, and the equality follows from $\|\delta\{a\}\| = (k-r)\|a\|$ for $a \in H_r$.

We now deal with the problem that neighborhoods of vertices in H_r may overlap. For $b \in H_{r+1}$ and a terminal $t \in \text{Ter}$, we say t *contributes* to b , if $b \in \delta H_r(t)$, i.e., if b belongs to the coboundary of the local piece $H_r(t)$. Let

$$n_b := |\{t \in \text{Ter} : b \in \delta H_r(t)\}|.$$

Thus n_b is the number of terminal-cluster coboundaries that contain b . Notice that

$$n_b \leq |\text{Ter}(b)|,$$

and equality may not hold, because it is possible that b is connected through heavy vertices to a terminal t , but b cancels out in $\delta H_r(t)$.

In this terminology,

$$\sum_{t \in \text{Ter}} \|\delta H_r(t)\| = \sum_{b \in H_{r+1}} n_b \|b\|, \quad \text{and} \quad \left\| \sum_{t \in \text{Ter}} \delta H_r(t) \right\| = \sum_{b \in H_{r+1}} (n_b \bmod 2) \|b\|.$$

Hence,

$$\sum_{t \in \text{Ter}} \|\delta H_r(t)\| - \left\| \sum_{t \in \text{Ter}} \delta H_r(t) \right\| = \sum_{b \in H_{r+1}} (n_b - (n_b \bmod 2)) \|b\| \leq \sum_{b \in H_{r+1}} 2(n_b - 1)_+ \|b\|.$$

where we used the inequality $n - (n \bmod 2) \leq 2(n - 1)_+$ which holds for every $n \in \mathbb{N}$. Now, substituting $n_b \leq |\text{Ter}(b)| = E(b) + 1$ and rearranging, we get

$$\sum_{t \in \text{Ter}} \|\delta H_r(t)\| \leq \left\| \sum_{t \in \text{Ter}} \delta H_r(t) \right\| + 2 \sum_{b \in H_{r+1}} E(b) \|b\| = \left\| \sum_{t \in \text{Ter}} \delta H_r(t) \right\| + 2 \|E_{r+1}\|. \quad (4.9)$$

Combining [Equations \(4.8\)](#) and [\(4.9\)](#) proves the proposition. $\square$

4.3 Excess

We want to control the overall excess E_i through *local* structures, where by local we mean that we only consider vertices that are one layer or two away. We define:

- For $a \in H_i$, $i \geq 1$, define the same-type excess

$$\text{Same}(a) := \sum_{|J|=i-1} (|\text{Ch}(a) \cap H_J| - 1)_+. \quad (4.10)$$

- If b and b' are children of a of distinct types, the cubical completion property supplies a lower vertex completing their square. We say that the square is *light* if its lower vertex is light, $b \overset{a}{\wedge} b' \in L$. We denote the number of light squares below a by $\text{Light}(a)$, i.e.,

$$\text{Light}(a) := |\{\{a; b, b'; c\} \in X(2) \mid b, b' \in \text{Ch}(a), \quad c \in L\}|. \quad (4.11)$$

We also set $\text{Same}(a) = \text{Light}(a) = 0$ for $a \in H_0$, and $\text{Light}(a) = 0$ for $a \in H_1$. These two quantities describe the two ways in which heavy branches can fail to merge. The crux of the proof is that there are no other sources of terminal excess.

Lemma 4.3 (Recursive excess bound). *For every heavy vertex $a \in H$,*

$$E(a) \leq \text{Same}(a) + \text{Light}(a) + \sum_{x \in \text{Ch}(a)} E(x). \quad (4.12)$$

Proof. If $\text{Ch}(a) = \emptyset$, then a is terminal and $E(a) = 0$. Otherwise, fix one heavy child $b^* \in \text{Ch}(a)$, and partition the remaining children into three sets:

- Vertices in $\text{Ch}(a)$ that have the same type as b^* ,
- Vertices in $\text{Ch}(a)$ that have a different type than b^* , such that $b \overset{a}{\wedge} b^* \in H$,
- Vertices in $\text{Ch}(a)$ that have a different type than b^* , such that $b \overset{a}{\wedge} b^* \in L$.

Accordingly, we define,

$$\begin{aligned} \text{SameType}_{a,b^*} &:= \{b \in \text{Ch}(a) : b \neq b^*, \text{ type}(b) = \text{type}(b^*)\}, \\ \text{LightSquare}_{a,b^*} &:= \left\{b \in \text{Ch}(a) : \text{type}(b) \neq \text{type}(b^*), b \overset{a}{\wedge} b^* \in L\right\}. \end{aligned}$$

We claim that for every $b^* \neq b \in \text{Ch}(a)$,

$$|\text{Ter}(b) \setminus \text{Ter}(b^*)| \leq (|\text{Ter}(b)| - 1) + \mathbf{1}_{b \in \text{SameType}_{a,b^*}} + \mathbf{1}_{b \in \text{LightSquare}_{a,b^*}}. \quad (4.13)$$

To see this notice that $|\text{Ter}(b) \setminus \text{Ter}(b^*)| \leq |\text{Ter}(b)|$, and so if $b \in \text{SameType}_{a,b^*}$, or $b \in \text{LightSquare}_{a,b^*}$ the inequality immediately holds. Otherwise, the fourth vertex $c = b \overset{a}{\wedge} b^*$ of the unique square completing a, b, b^* is heavy. Observe that $\text{Ter}(c) \subseteq \text{Ter}(b) \cap \text{Ter}(b^*)$. It follows that

$$|\text{Ter}(b) \setminus \text{Ter}(b^*)| \leq |\text{Ter}(b)| - |\text{Ter}(c)| \leq |\text{Ter}(b)| - 1.$$

Again the inequality in (4.13) immediately holds. As

$$\text{Ter}(a) = \text{Ter}(b^*) \cup \bigcup_{b^* \neq b \in \text{Ch}(a)} (\text{Ter}(b) \setminus \text{Ter}(b^*)),$$

we have

$$|\text{Ter}(a)| \leq |\text{Ter}(b^*)| + \sum_{b^* \neq b \in \text{Ch}(a)} E(b) + |\text{SameType}_{a,b^*}| + |\text{LightSquare}_{a,b^*}|. \quad (4.14)$$

Clearly, $|\text{SameType}_{a,b^*}| \leq \text{Same}(a)$ and $|\text{LightSquare}_{a,b^*}| \leq \text{Light}(a)$. Plugging in these bounds and subtracting one from both sides of (4.14) gives

$$E(a) \leq \text{Same}(a) + \text{Light}(a) + \sum_{b \in \text{Ch}(a)} E(b)$$

as required. $\square$




Figure 4: The four vertex sets of a two-dimensional square complex, arranged by level. An initial set of vertices H_1 is marked heavy (red) on level 1. In the layer below, only vertices that have many heavy parents are heavy. In the layer above, all upward neighbors of the heavy vertices are marked heavy. Terminal vertices (heavy vertices without heavy children) are enlarged. Red vertices shaped like diamonds, stars, and triangles have excess. The excess of the triangular vertex is accounted for by Same_2 , whereas the excess of the diamond shaped vertex is accounted for by Light_2 . The excess of the star shaped vertex is accounted for by both functions.

In words, the recursive lemma says that ambiguity can only come from ambiguity already present one layer below, together with new same-type and light-square charges created at level i . We now unroll the recursion.

Theorem 4.4 (Excess recursion). *For every $1 \leq i \leq r + 1$,*

$$\|E_i\| \leq (k - i + 1) \|E_{i-1}\| + \|\text{Same}_i\| + \|\text{Light}_i\| \quad (4.15)$$

Proof. By Lemma 4.3,

$$\sum_{a \in H_i} \|a\| E(a) \leq \sum_{a \in H_i} \|a\| (\text{Same}(a) + \text{Light}(a)) + \sum_{a \in H_i} \|a\| \sum_{b \in \text{Ch}(a)} E(b).$$

Rearranging the first sum, we have

$$\sum_{a \in H_i} \|a\| (\text{Same}(a) + \text{Light}(a)) = \|\text{Same}_i\| + \|\text{Light}_i\|.$$

For the second sum, changing the order of summation, we get

$$\sum_{a \in H_i} \|a\| \sum_{b \in \text{Ch}(a)} E(b) = \sum_{b \in H_{i-1}} E(b) \sum_{a: b \in \text{Ch}(a)} \|a\| \leq \sum_{b \in H_{i-1}} E(b) \sum_{b \leq a} \|a\|,$$

where the inequality holds because on the left-hand side we sum only over heavy parents, whereas on the right-hand side we sum over all parents. Next, we use the norm-identity on [Equation \(2.2\)](#) to obtain

$$\sum_{b \in H_{i-1}} E(b) \sum_{b \leq a} \|a\| = (k - i + 1) \sum_{b \in H_{i-1}} E(b) \|b\| = (k - i + 1) \|E_{i-1}\|.$$

Combining the equations above concludes the proof. $\square$

Upon recalling that every $a \in H_0$ is terminal and has no excess, we get $E_0 = 0$ and we can immediately deduce the following:

Corollary 4.5 (Iterated excess bound). *For every $0 \leq i \leq r + 1$,*

$$\|E_i\| \leq \sum_{j=1}^i \frac{(k-j)!}{(k-i)!} (\|\text{Same}_j\| + \|\text{Light}_j\|). \quad (4.16)$$

4.4 Same Type Collisions

The previous subsection reduced terminal overlap to the level charges $\|\text{Same}_i\|$ and $\|\text{Light}_i\|$. Our next task is to show that every one of these charges is small compared to the norm of the set H_{i-1} . There are two ingredients: lossless expansion controls same-type collisions and additional imbalance assumptions control light squares. In this subsection we deal with the first type of error.

Lemma 4.6. *For every $1 \leq i \leq r + 1$,*

$$\|\text{Same}_i\| \leq \left(\sum_{j=1}^k \varepsilon_j \right) \|H_{i-1}\|. \quad (4.17)$$

Proof. Expand the definition of $\|\text{Same}_i\|$ and group the sum by the type I of the heavy vertices at level i :

$$\begin{aligned} \sum_{a \in H_i} \|a\| \text{Same}(a) &= \sum_{|I|=i} \|I\| \cdot \sum_{a \in H_I} \sum_{j \in I} \left(\deg_{H_I \setminus \{j\}}(a) - 1 \right)_+ \\ &= \sum_{|I|=i} \|I\| \cdot \sum_{j \in I} \sum_{a \in H_I} \left(\deg_{H_I \setminus \{j\}}(a) - 1 \right)_+ \\ &\leq \sum_{|I|=i} \|I\| \cdot \sum_{j \in I} \varepsilon_j D_j |H_I \setminus \{j\}|, \end{aligned}$$

where we have used [Item 1](#) in [Lemma 2.19](#).

Each set J of size $i - 1$ appears in the summands corresponding to $I = J \cup \{j\}$ for $j \notin J$. Therefore, we can rewrite the sum as

$$\|\text{Same}_i\| \leq \sum_{|J|=i-1} \sum_{j \notin J} \|J \cup \{j\}\| \varepsilon_j D_j |H_J| = \sum_{|J|=i-1} \sum_{j \notin J} \varepsilon_j \|H_J\| \leq \left(\sum_{j=1}^k \varepsilon_j \right) \|H_{i-1}\|$$

$\square$

4.5 Light Squares

Each light square contributing to Light_i lies between levels $i-2$, $i-1$, and i . Instead of counting these squares from their heavy upper vertices at level i , we count them from their light lower vertices at level $i-2$. Throughout, we assume that every set to which lossless expansion is applied is small enough to be in the appropriate lossless range.

Lemma 4.7. *For every $2 \leq i \leq r+1$,*

$$\|\text{Light}_i\| \leq (i-1)(k-i+1) \max_{s < t} \{\varepsilon_t D'_s, \gamma_{i-2}\} \|H_{i-1}\|. \quad (4.18)$$

Proof. Fix $2 \leq i \leq r+1$. By definition,

$$\|\text{Light}_i\| = \sum_{a \in H_i} |\{\{a; b, b'; c\} \in X(2) : b, b' \in \text{Ch}(a), c \in L\}| \cdot \|a\| = \sum_{a \in H_i} \sum_{\substack{\{b, b'\} \subseteq \text{Ch}(a) \\ \text{type}(b) \neq \text{type}(b')}} \mathbf{1}_{b \wedge b' \in L} \|a\|.$$

Instead of counting the light squares from above, we change the order of summation and count them from below.

$$\|\text{Light}_i\| = \sum_{c \in L_{i-2}} \sum_{\substack{\{b, b'\} \subseteq H_{i-1} \\ \text{type}(b) \neq \text{type}(b') \\ c \leq b, b'}} \mathbf{1}_{b \wedge b' \in H_i} \|b \overset{c}{\wedge} b'\| \leq \sum_{c \in L_{i-2}} \sum_{\substack{\{b, b'\} \subseteq H_{i-1} \\ \text{type}(b) \neq \text{type}(b') \\ c \leq b, b'}} \|b \overset{c}{\wedge} b'\|.$$

For a fixed $c \in L_J$, unique square completion gives one square for every choice of heavy parents in two directions $s < t$ outside J . Consequently,

$$\begin{aligned} \sum_{\substack{\{b, b'\} \subseteq H_{i-1} \\ \text{type}(b) \neq \text{type}(b') \\ c \leq b, b'}} \|b \overset{c}{\wedge} b'\| &= \sum_{\substack{s < t \\ s, t \notin J}} \|J \cup \{s, t\}\| d_s(c) d_t(c) \\ &= \|J\| \sum_{\substack{s < t \\ s, t \notin J}} u_s(c) u_t(c). \end{aligned}$$

Substituting this identity into the preceding inequality gives

$$\|\text{Light}_i\| \leq \sum_{|J|=i-2} \|J\| \sum_{c \in L_J} \sum_{\substack{s < t \\ s, t \notin J}} u_s(c) u_t(c). \quad (4.19)$$

For a fixed J , we analyze the inner two sums of (4.19). List the directions outside J as $m_1 < m_2 < \dots < m_{k-i+2}$ and rewrite the pair-product sum as

$$\sum_{1 \leq s < t \leq k-i+2} u_{m_s}(c) u_{m_t}(c) = \sum_{t=2}^{k-i+2} u_{m_t}(c) U_{< m_t}(c), \quad \text{for} \quad U_{< m_t}(c) := \sum_{s < t} u_{m_s}(c).$$

Following [PK22], we supply a bound for the product $u_{m_t}(c) U_{< m_t}(c)$, which is linear in the upward heavy-degrees. We claim

$$u_{m_t}(c) U_{< m_t}(c) \leq \gamma_{i-2} (u_{m_t}(c) - \varepsilon_{m_t})_+ + \varepsilon_{m_t} U_{< m_t}(c). \quad (4.20)$$

Indeed, for $u_{m_t}(c) \leq \varepsilon_{m_t}$, the second term on the right-hand side already dominates the left-hand side. For $u_{m_t}(c) > \varepsilon_{m_t}$, the inequality follows by plugging in $\gamma_{i-2} \geq U_{<m_t}(c)$, which follows from the fact that c is light. Applying this bound pointwise for every c in the summation and changing the order of summation, we obtain

$$\sum_{c \in L_J} \sum_{\substack{s < t \\ s, t \notin J}} u_s(c) u_t(c) \leq \sum_{t=2}^{k-i+2} \left[\gamma_{i-2} \sum_{c \in L_J} (u_{m_t}(c) - \varepsilon_{m_t})_+ + \varepsilon_{m_t} \sum_{c \in L_J} U_{<m_t}(c) \right].$$

Now, once the expression is linear in the upward-degrees, we translate the summation to edge counts. [Item 3](#) in [Lemma 2.19](#) and biregularity give, for every $m \notin J$,

$$\sum_{c \in L_J} (u_m(c) - \varepsilon_m)_+ \leq \frac{|H_{J \cup \{m\}}|}{D_m}, \quad \sum_{c \in L_J} u_m(c) \leq \frac{D'_m |H_{J \cup \{m\}}|}{D_m}.$$

Combining these two estimates with [\(4.20\)](#), we get

$$\sum_{c \in L_J} \sum_{\substack{s < t \\ s, t \notin J}} u_s(c) u_t(c) \leq \sum_{t=2}^{k-i+2} \left(\gamma_{i-2} \frac{|H_{J \cup \{m_t\}}|}{D_{m_t}} + \sum_{s < t} \varepsilon_{m_t} D'_{m_s} \frac{|H_{J \cup \{m_s\}}|}{D_{m_s}} \right).$$

Rearranging the summation, the coefficient of $|H_{J \cup \{m_s\}}|$ is

$$\frac{1}{D_{m_s}} \left(\mathbf{1}_{s \geq 2} \gamma_{i-2} + \sum_{t=s+1}^{k-i+2} \varepsilon_{m_t} D'_{m_s} \right) \leq \frac{k-i+1}{D_{m_s}} \max_{1 \leq p < q \leq k} \{ \varepsilon_q D'_p, \gamma_{i-2} \}.$$

Indeed, the expression in parentheses has at most $k-i+1$ nonzero terms, each at most $\max_{1 \leq p < q \leq k} \{ \varepsilon_q D'_p, \gamma_{i-2} \}$. Substituting this bound into [\(4.19\)](#), we obtain

$$\begin{aligned} \|\text{Light}_i\| &\leq (k-i+1) \max_{s < t} \{ \varepsilon_t D'_s, \gamma_{i-2} \} \sum_{|J|=i-2} \|J\| \sum_{m \notin J} \frac{|H_{J \cup \{m\}}|}{D_m} \\ &= (k-i+1) \max_{s < t} \{ \varepsilon_t D'_s, \gamma_{i-2} \} \sum_{|J|=i-2} \sum_{m \notin J} \|H_{J \cup \{m\}}\|. \end{aligned}$$

Every type K of size $i-1$ arises from exactly $i-1$ pairs (J, m) , one for each $m \in K$. Consequently,

$$\|\text{Light}_i\| \leq (i-1)(k-i+1) \max_{s < t} \{ \varepsilon_t D'_s, \gamma_{i-2} \} \|H_{i-1}\|,$$

as claimed. □

4.6 Heavy-mass propagation

In [Sections 4.4](#) and [4.5](#) we showed that $\|\text{Same}_i\|$ and $\|\text{Light}_i\|$ are small relative to $\|H_{i-1}\|$. However, we want to compare the error with $\|H_r\|$ rather than $\|H_{i-1}\|$. In this section we relate the two quantities.

Lemma 4.8 (Heavy-mass propagation). *For every $1 \leq i \leq r$, if $\gamma_{i-1} > \sum_{j \in [k]} \varepsilon_j$, then*

$$\|H_{i-1}\| \leq \frac{i}{\gamma_{i-1} - \sum_{j \in [k]} \varepsilon_j} \|H_i\|. \quad (4.21)$$

The same holds for $i = r + 1$ under the convention $\gamma_r = k - r$.

Proof. Fix $1 \leq i \leq r$ and $I \subseteq [k]$ with $|I| = i - 1$. The definition of heaviness gives $\sum_{j \notin I} u_j(a) \geq \gamma_{i-1}$ for every $a \in H_I$. Summing over $a \in H_I$, we obtain

$$\gamma_{i-1} |H_I| \leq \sum_{j \notin I} \frac{1}{D_j} \sum_{a \in H_I} d_j(a). \quad (4.22)$$

For each $j \notin I$, the lossless edge estimate gives

$$\sum_{a \in H_I} d_j(a) = e(H_I, H_{I \cup \{j\}}) \leq |H_{I \cup \{j\}}| + \varepsilon_j D_j |H_I|.$$

Substituting these estimates into (4.22) and using $\sum_{j \notin I} \varepsilon_j \leq \sum_{j \in [k]} \varepsilon_j$ yields

$$\left(\gamma_{i-1} - \sum_{j \in [k]} \varepsilon_j \right) |H_I| \leq \sum_{j \notin I} \frac{|H_{I \cup \{j\}}|}{D_j}. \quad (4.23)$$

Divide both sides by D_I and sum over all types I of size $i - 1$. Every type K of size i occurs exactly i times on the right, once for each choice of $j \in K$. Consequently,

$$\left(\gamma_{i-1} - \sum_{j \in [k]} \varepsilon_j \right) \|H_{i-1}\| \leq i \|H_i\|.$$

Dividing both sides by $\gamma_{i-1} - \sum_{j \in [k]} \varepsilon_j$ gives the desired inequality. $\square$

The iterated one-step comparison tells us the total price of pulling heaviness down from level r to level $i - 1$. Thus, for every $1 \leq i \leq r + 1$, including the empty-product case $i = r + 1$, we obtain

Corollary 4.9. *Assume that $\gamma_{s-1} > \sum_{j \in [k]} \varepsilon_j$ for every $i \leq s \leq r$. Then,*

$$\|H_{i-1}\| \leq \prod_{s=i}^r \frac{s}{\gamma_{s-1} - \sum_{j \in [k]} \varepsilon_j} \|H_r\| \quad (4.24)$$

4.7 Choice of parameters and final error bound

Imbalance Throughout, we assume that the complex is imbalanced enough, so that

$$\max_{1 \leq s < t \leq k} \{\varepsilon_t D'_s\} \leq \varepsilon_\star, \quad \text{for} \quad \varepsilon_\star := \sum_{j=1}^k \varepsilon_j. \quad (4.25)$$

Initial norm bound Lossless expansion in direction $j \notin J$ applies to sets $A \subseteq V_J$ with $|A| \leq c_j |V_J|$, equivalently $\|A\| \leq c_j \|V_J\|$. To ensure that all sets used in the lossless estimates lie in this range, put

$$c = \frac{\min_j c_j}{2^k k! D'_{[k]} D_{[k]}}. \quad (4.26)$$

The reader may verify that this conservative choice of c guarantees that all sets used in the lossless expansion estimates satisfy $\|A\| \leq c_j \|V_J\|$. One can also improve this constant by following the heavy-mass propagation through the levels of the complex, but we do not pursue this here.

Heaviness thresholds For $1 \leq s \leq r$, define

$$\gamma_{s-1} := \varepsilon_\star + \varepsilon_\star^{1/2^s}. \quad (4.27)$$

The thresholds satisfy $\gamma_{s-1} > \varepsilon_\star$ for $1 \leq s \leq r$, as required for the heavy-mass propagation corollary, [Corollary 4.9](#).

The main theorem

Theorem 4.10. *Let X be a lossless k -dimensional square complex with parameters $(D_i, D'_i, c_i, \varepsilon_i)_{i \in [k]}$. Assume that the complex is imbalanced enough, so that [Equation \(4.25\)](#) holds. For $0 \leq r \leq k-1$, let $H_r : V_r \rightarrow \mathbb{F}_2$ be a locally minimal cochain, of norm at most $c \|V_r\|$, for c as in [Equation \(4.26\)](#). Then,*

$$\|\delta H_r\| \geq (m_{k,r-1} - O_k(\varepsilon_\star^{1/2^r})) \|H_r\|.$$

Proof. Set the thresholds as in (4.27). Invoking [Proposition 4.2](#), we get

$$\|\delta H_r\| \geq \sum_{t \in \text{Ter}} \|\delta H_r(t)\| - (k-r) \|E_r\| - 2 \|E_{r+1}\|. \quad (4.28)$$

Then, the excess norm is bounded using [Corollary 4.5](#), giving us

$$(k-r) \|E_r\| + 2 \|E_{r+1}\| \leq O_k(\|\text{Same}\| + \|\text{Light}\|). \quad (4.29)$$

The imbalance assumption gives us that for all $2 \leq i \leq r+1$,

$$\max_{1 \leq s < t \leq k} \{\varepsilon_t D'_s\} \leq \varepsilon_\star \leq \gamma_{i-2}.$$

Furthermore, we may assume that $\varepsilon_\star$ is sufficiently small in terms of k , since otherwise the claimed lower bound is nonpositive. Hence, all conditions of [Lemmas 4.6](#) and [4.7](#) are satisfied. We get

$$\begin{aligned} \|\text{Same}_i\| &= O_k(\varepsilon_\star) \|H_{i-1}\|, & 1 \leq i \leq r+1, \\ \|\text{Light}_i\| &= O_k(\gamma_{i-2}) \|H_{i-1}\|, & 2 \leq i \leq r+1. \end{aligned}$$

Also, $\text{Light}_1 = 0$, since there are no squares with upper vertex at level 1. Plugging the chosen thresholds into [Corollary 4.9](#), we translate the above bounds to depend on the initial set H_r , rather than on H_{i-1} . For $1 \leq i \leq r+1$, including the empty-sum case $i = r+1$, we have

$$\|H_{i-1}\| = O_k\left(\varepsilon_\star^{-\sum_{s=i}^r 2^{-s}}\right) \|H_r\| = O_k\left(\varepsilon_\star^{-1/2^{i-1} + 1/2^r}\right) \|H_r\|.$$

Consequently, for $1 \leq i \leq r+1$,

$$\begin{aligned}\|\text{Same}_i\| &= O_k \left(\varepsilon_\star^{1-1/2^{i-1}+1/2^r} \right) \|H_r\| \leq O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r\|, & 1 \leq i \leq r+1, \\ \|\text{Light}_i\| &= O_k \left(\gamma_{i-2} \varepsilon_\star^{-1/2^{i-1}+1/2^r} \right) \|H_r\| = O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r\|, & 2 \leq i \leq r+1.\end{aligned}$$

Summing over these ranges gives a factor of at most $r+1$, which is absorbed into the O_k notation, so that

$$\|\text{Same}\| = O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r\| \quad \text{and} \quad \|\text{Light}\| = O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r\|.$$

Combining these bounds with [Equations \(4.28\) and \(4.29\)](#) then gives us

$$\|\delta H_r\| \geq \sum_{t \in \text{Ter}} \|\delta H_r(t)\| - O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r\|. \quad (4.30)$$

Fix a terminal vertex $t \in V_I$ and put $i = |I|$. Viewed as a cochain on $\text{Lk}(t)$, the cluster $H_r(t)$ lies at link level $r-i-1$. This cochain is minimal in the link. Indeed, a link coboundary extends to an ambient coboundary of a cochain supported in $\text{Lk}(t)$. Since $H_r(t) \subseteq H_r$, any such correction decreasing the norm of $H_r(t)$ would also decrease the norm of H_r , contradicting local minimality. Applying [Theorem 3.2](#) at this level and converting link norms to ambient norms by the factor D_I^{-1} gives

$$\|\delta H_r(t)\| \geq m_{k-i,r-i-1} \|H_r(t)\| - D_I^{-1} O_k(\varepsilon_\star) \geq m_{k,r-1} \|H_r(t)\| - D_I^{-1} O_k(\varepsilon_\star).$$

We now show that the additive term is negligible compared with the mass of the terminal cluster. The heavy mass propagates, and [Corollary 4.9](#) invoked with $H(t)$ instead of H gives

$$\|H_r(t)\| \geq \frac{i!}{r!} \varepsilon_\star^{1/2^i-1/2^r} \|t\|. \quad (4.31)$$

As $\|t\| = D_I^{-1}$, the error term in $D_I^{-1} O_k(\varepsilon_\star)$ is bounded by

$$D_I^{-1} O_k(\varepsilon_\star) \leq O_k \left(\varepsilon_\star^{1-1/2^i+1/2^r} \right) \|H_r(t)\| \leq O_k \left(\varepsilon_\star^{1/2^r} \right) \|H_r(t)\|.$$

Hence, we can replace the additive error by a multiplicative error and obtain

$$\|\delta H_r(t)\| \geq \left(m_{k,r-1} - O_k \left(\varepsilon_\star^{1/2^r} \right) \right) \|H_r(t)\|.$$

Summing over all terminals yields

$$\sum_{t \in \text{Ter}} \|\delta H_r(t)\| \geq \left(m_{k,r-1} - O_k \left(\varepsilon_\star^{1/2^r} \right) \right) \sum_{t \in \text{Ter}} \|H_r(t)\| \geq \left(m_{k,r-1} - O_k \left(\varepsilon_\star^{1/2^r} \right) \right) \|H_r\|,$$

where the last inequality follows because every $x \in H_r$ belongs to the cluster of at least one terminal. Plugging this into [Equation \(4.30\)](#) concludes the proof. $\square$

Proof of Theorem 1.3. With c as in [\(4.26\)](#), [Theorem 4.10](#) implies that, at every level $0 \leq r < k$, every locally minimal cochain H_r of norm at most $c \|V_r\|$ expands by a factor of at least $m_{k,r-1} - O_k \left(\varepsilon_\star^{\Omega_k(1)} \right)$. Choosing $a = \min_r m_{k,r-1}$ concludes the proof. $\square$

5 qLTC

In this section we use [Theorem 4.10](#) to show that a sufficiently imbalanced four-dimensional decorated square complex gives rise to a quantum locally testable code. We start with the required definitions and then describe the construction and prove its qLTC properties. Throughout this section, the degree and lossless-expansion parameters are fixed independently of the code length.

5.1 Definitions

One way to characterize a linear code is by a parity-check matrix for it, $\ker(H) = \mathcal{C}$. Such a matrix is not unique. Any matrix having the same row-span over $\mathbb{F}_2$ characterizes the same code. Consider the following probabilistic procedure for testing whether a given word $w \in \mathbb{F}_2^n$ is in the code: choose a row R uniformly at random from H and check if $R \cdot w^t = \sum_{i=1}^n R_i w_i$ is 0 over $\mathbb{F}_2$. If it is – accept, announcing $w \in \mathcal{C}$. If it is 1 – reject, announcing $w \notin \mathcal{C}$. A codeword will never be rejected, but the rejection probability of a non-codeword depends on H . Ideally the distance of the given w from $\mathcal{C}$ should be reflected by the rejection probability of the testing procedure. Codes which allow for such a choice of H are called locally testable codes.

For the codes arising from our complexes, we use the following weighted formulation of local testability. Words and syndromes carry the norms of their respective levels in the underlying complex; for the reversed complex, we use its reversed norms.

Definition 5.1 (Classical Locally Testable Code (LTC) [[GS06](#), [LLZ22](#), Definition 11]). *For $s > 0$, a classical linear code $\mathcal{C} \leq \mathbb{F}_2^n$ is s -locally testable if it has a parity-check matrix $H \in \mathbb{F}_2^{m \times n}$ ($\ker(H) = \mathcal{C}$) with rows and columns of weight bounded by a constant such that for every word $w \in \mathbb{F}_2^n$,*

$$\|Hw\| \geq s \cdot \text{dist}(w, \mathcal{C})$$

where, recall, $\text{dist}(x, \mathcal{C}) = \min_{c \in \mathcal{C}} \|x - c\|$.

The parameter s is the soundness of the test, while its locality is the maximum row weight of H .

The weighted norms in [Definition 5.1](#) are comparable to Hamming weight by [Remark 2.11](#). For our codes, biregularity also bounds the ratio m/n above and below by positive constants depending only on k, D_i, D'_i . Consequently, the weighted formulation is equivalent, up to constant factors in soundness, to the usual inequality $\frac{1}{m} \|Hw\| \geq \frac{s}{n} \text{dist}_H(w, \mathcal{C})$, where dist_H denotes Hamming distance. These factors depend only on k, D_i, D'_i .

Definition 5.2 (Quantum CSS code [[CS96](#), [Ste96a](#)]). *A quantum CSS code $\mathcal{C}(\mathcal{C}_X, \mathcal{C}_Z)$ is specified by two classical linear codes $\mathcal{C}_X, \mathcal{C}_Z \leq \mathbb{F}_2^n$ satisfying $\mathcal{C}_X^\perp \leq \mathcal{C}_Z$ (and thus also $\mathcal{C}_Z^\perp \leq \mathcal{C}_X$).*

Using parity-check matrices $H_Z \in \mathbb{F}_2^{m_Z \times n}, H_X \in \mathbb{F}_2^{m_X \times n}$ for $\mathcal{C}_Z, \mathcal{C}_X$ respectively, with code length n and number of Z -checks m_Z and X -checks m_X , the definition is equivalent to the following being a cochain complex (composition of maps is zero): $\mathbb{F}_2^{m_X} \xrightarrow{H_X^t} \mathbb{F}_2^n \xrightarrow{H_Z} \mathbb{F}_2^{m_Z}$. We say the middle space $\mathbb{F}_2^n$ hosts the qubits.

Using this notation, the distance of the CSS code is the minimum $\min\{d_Z, d_X\}$ with $d_Z := \min\{|w| : w \in \ker(H_Z) \setminus \text{im}(H_X^t)\}$ and $d_X = \min\{|w| : w \in \ker(H_X) \setminus \text{im}(H_Z^t)\}$. The rate ρ is the ratio of log-dimension of the code to the number of qubits n (its length). It is lower bounded by $\rho \geq \frac{n - m_Z - m_X}{n} = 1 - \frac{m_Z + m_X}{n}$.

Definition 5.3 (Quantum Locally Testable CSS Code (qLTC) [AE18, EH17, Fact 17]⁶). A quantum CSS code $(\mathcal{C}_X, \mathcal{C}_Z)$ is s -locally testable if both $\mathcal{C}_X$ and $\mathcal{C}_Z$ are.

5.2 Construction and Proof

For the quantum locally testable code we will use a $(4, 2)$ -cubical complex. The code uses the middle three layers as a CSS cochain complex, i.e. $C^1(X) \xrightarrow{\delta_1} C^2(X) \xrightarrow{\delta_2} C^3(X)$, with $\mathcal{C}_Z = \ker(\delta_2)$ and $\mathcal{C}_X = \ker(\delta_1^t)$ and we use the coboundary and boundary maps as the parity-check matrices $H_Z := \delta_2$ and $H_X := \delta_1^t$. To show this is indeed a qLTC we follow Lin and Hsieh's arguments in [LH22a, LH22b] which used and generalized Sipser and Spielman's decoder of [SS96]. With the adjustments described below, their argument applies in our setting. At a high level, in [LH22b] they show that, using necessary imbalance of the blocks, a $(2, 2)$ -cubical complex as a CSS code constitutes an asymptotically good quantum LDPC code, and in [LH22a] they show that using such a complex (with different imbalance) and putting the bits in the bottom layer constitutes an asymptotically good classical LTC. We use the former for the sub-complex $C^1(X) \xrightarrow{\delta_1} C^2(X) \xrightarrow{\delta_2} C^3(X)$ to show the code has constant relative distance, and the latter for the sub-complexes $C^2(X) \xrightarrow{\delta_2} C^3(X) \xrightarrow{\delta_3} C^4(X)$ and $C^0(X) \xleftarrow{\delta_0^t} C^1(X) \xleftarrow{\delta_1^t} C^2(X)$ to show each of the codes $\mathcal{C}_Z$ and $\mathcal{C}_X$ has a constant local testability soundness. Constant rate is shown by another imbalance condition.

The main generalization that requires an adjustment for our setting is the notion of locality: local minimality in this work requires considering any cochain in the previous level that is supported in any one link of the bottom layer. This is the adjustment we introduce into the decoder algorithm [LH22a, Definition 16].

Recall that for $v \in V_0$, a cochain in its link $h \in C^{i-2}(\text{Lk}(v))$ corresponds to the layer V_{i-1} of X , and that $\iota_v h \in C^{i-1}(X)$ is its corresponding cochain.

Algorithm 1 Greedy flipping, adapted from [LH22a, Def. 16]

Input: $H_i \in C^i(X)$, with $1 \leq i \leq k$.

- 1: **while** there exist $v \in V_0$ and $h \in C^{i-2}(\text{Lk}(v))$ such that $\|H_i + \delta(\iota_v h)\| < \|H_i\|$ **do**
 - 2: $H_i \leftarrow H_i + \delta(\iota_v h)$
 - 3: **end while**
 - 4: **return** H_i
-

Lemma 5.4. For an input cochain H_i , *Algorithm 1* stops after at most $D_{[k]} \cdot \|H_i\|$ iterations and returns a locally-minimal cochain.

Proof. By the stopping condition, no correction supported in a bottom-vertex link decreases the output's norm. Every vertex v has a bottom ancestor $u \in V_0$, with $\text{Lk}(v) \subseteq \text{Lk}(u)$. Thus every

⁶The definition in [AE18] is more general. With the conventional normalization, [EH17, Fact 17] showed that a CSS code $(\mathcal{C}_X, \mathcal{C}_Z)$ for which both codes are s -locally testable is itself s -locally testable in the general definition. On the other hand they showed that if the CSS code is s -locally testable, then both $\mathcal{C}_Z, \mathcal{C}_X$ must be $s/2$ -locally testable. In this work we only care about the existence of a constant soundness for quantum local testability so we define the CSS code's local testability soundness directly in terms of the soundness parameters of the two classical codes.

correction allowed by Definition 2.15 is also supported in a bottom-vertex link, so the output is locally minimal.

Every cochain norm is an integer multiple of $\alpha := D_{[k]}^{-1}$, since D_I divides $D_{[k]}$ for every $I \subseteq [k]$. Each iteration strictly decreases the norm, hence decreases it by at least α . Thus the algorithm stops after at most $\frac{1}{\alpha} \|H_i\| = D_{[k]} \|H_i\|$ iterations for an input cochain H_i . $\square$

To obtain constant rate, we require the combined size of the two adjacent layers to be smaller than that of the middle layer by a constant fraction. For that, we will assume a minimal stretch/shrinkage of the direction-graphs in the complex. Define the stretch/shrinkage factor of direction i by $\mu_i := \frac{D_i}{D'_i}$.

Theorem 5.5 (qLTC from a $(4, 2)$ -cubical complex). *Let X be a $(4, 2)$ -cubical complex. Assume X and X^{rev} satisfy the assumptions in Theorem 4.10 with $\varepsilon_\star$ sufficiently small⁷ and with stretch/shrinkage factors $\mu_1, \mu_2 \geq 4$, and $\mu_3, \mu_4 \leq \frac{1}{4}$.*

Then $\mathcal{C}(\ker \delta_1^t, \ker \delta_2)$ is a quantum error-correcting code with constant rate, relative distance, and local testability soundness.




Figure 5: An illustration of the size of the blocks in the four-dimensional Boolean cube for $\mu_1 = \mu_2 = 4$ and $\mu_3 = \mu_4 = \frac{1}{4}$. Each block shows the ratio $|V_I| / |V_\emptyset|$.

Proof. Every vertex in V_3 and in V_1 has at most a constant number of neighbors in V_2 . Specifically, every row of $H_Z = \delta_2$ has weight at most $3 \max_{i \in [4]} D'_i$, and every row of $H_X = \delta_1^t$ has weight at

⁷ $\varepsilon_\star$ needs to be small so that the inequality in Theorem 4.10 is meaningful.

most $3 \max_{i \in [4]} D_i$. Thus the checks have constant weight, at most $3 \max_{i \in [4]} \{D_i, D'_i\}$. Similarly, every vertex in V_2 has at most $\sum_{i=1}^k D_i + D'_i$ neighbors in both V_1, V_3 , meaning H_Z, H_X have constant weight columns, or equivalently, every qubit participates in a constant number of checks.

For the rate, by biregularity, $|V_{I \cup \{i\}}| = \mu_i |V_I|$ for $i \notin I$, so $|V_I| = (\prod_{i \in I} \mu_i) |V_\emptyset|$, as illustrated in Figure 5. To bound the rate, count qubits and checks

$$\rho \geq 1 - \frac{|V_1| + |V_3|}{|V_2|} > 0.1.$$

To bound both the local testability soundness and the relative distance, we apply the Greedy Flipping Algorithm, Algorithm 1, using both Lemma 5.4 and Theorem 4.10, to show the parameters are bounded below by positive constants. We give the full arguments for both parameters considering the Z -codes, $\mathcal{C}_Z = \ker(\delta_2)$ and $\mathcal{C}_Z/\mathcal{C}_X^\perp = \ker(\delta_2)/\text{im}(\delta_1)$, respectively. The same arguments applied to the reversed complex give the local testability soundness and relative distance of the X -codes $\mathcal{C}_X = \ker(\delta_2^{\text{rev}}) = \ker(\delta_1^t)$ and $\mathcal{C}_X/\mathcal{C}_Z^\perp = \ker(\delta_2^{\text{rev}})/\text{im}(\delta_1^{\text{rev}}) = \ker(\delta_1^t)/\text{im}(\delta_2^t)$.

For distance, let $w \in \ker(\delta_2) \setminus \text{im}(\delta_1)$ represent a nonzero element in $\mathcal{C}_Z/\mathcal{C}_X^\perp$. Run the greedy flipping algorithm on it. Its output is $w' = w + \sum_{i=1}^r \delta_1(v_i)$, with $\|w'\| \leq \|w\|$, and is locally minimal by Lemma 5.4. Since $w' - w \in \text{im}(\delta_1) = \mathcal{C}_X^\perp$, the cochains w and w' represent the same nonzero element, so $w' \neq 0$. As $\delta_2(w') = 0$, Theorem 4.10 gives $\|w\| \geq \|w'\| \geq c \|V_2\|$.

For the local testability of $\mathcal{C}_Z = \ker(\delta_2)$, look at its syndrome layer V_3 . Let $w \in C^2(X)$ be some word (not necessarily a codeword in the kernel of δ_2). By adding the closest codeword to it with no change to its syndrome we may assume its closest codeword is 0. We wish to show that its norm is bounded above by a constant multiple of the norm of its syndrome. Let $H_{\text{syn}} := \delta_2(w)$ be its syndrome. Run the greedy flipping algorithm on H_{syn} and obtain $H'_{\text{syn}} = \delta_2(w) + \sum_{i=1}^r \delta_2(v_i)$ which has at most the same norm and is locally minimal. As $C^2(X) \rightarrow C^3(X) \rightarrow C^4(X)$ is a cochain complex, $\delta_3(H'_{\text{syn}}) = \delta_3(\delta_2(w + \sum_{i=1}^r v_i)) = 0$. By Theorem 4.10 either $H'_{\text{syn}} = 0$ or it has large norm $\|H'_{\text{syn}}\| \geq c \|V_3\|$. Divide into these two cases:

1. $\|H'_{\text{syn}}\| \geq c \|V_3\|$. The algorithm never increases the norm, so $\|H_{\text{syn}}\| \geq \|H'_{\text{syn}}\|$. Also, $\text{dist}(w, \mathcal{C}_Z) \leq \|w\| \leq \|V_2\|$. By biregularity, $\beta := \|V_3\| / \|V_2\|$ is a positive constant depending only on the degree parameters. Thus,

$$\|H_{\text{syn}}\| \geq \|H'_{\text{syn}}\| \geq c \|V_3\| = c\beta \|V_2\| \geq c\beta \cdot \text{dist}(w, \mathcal{C}_Z).$$

This gives the desired soundness bound with $s = c\beta > 0$ in this case.

2. $H'_{\text{syn}} = 0$. Set $u = \sum_{i=1}^r v_i$. Since $\delta_2(w + u) = H'_{\text{syn}} = 0$, the corrected word $w + u$ belongs to $\mathcal{C}_Z$. Each change $v_i \in C^2(X)$ is supported on $\text{Lk}_1(v)$ for some $v \in V_0$. This link level has six blocks, each of ambient norm at most 1; hence $\|v_i\| \leq 6$. By Lemma 5.4, $r \leq D_{[k]} \cdot \|H_{\text{syn}}\|$. Therefore,

$$\text{dist}(w, \mathcal{C}_Z) \leq \|u\| \leq \sum_{i=1}^r \|v_i\| \leq 6r \leq 6D_{[k]} \|H_{\text{syn}}\|.$$

Rearranging gives $\|H_{\text{syn}}\| \geq s \cdot \text{dist}(w, \mathcal{C}_Z)$ with the constant $s = \frac{1}{6D_{[k]}} > 0$.

Together, the two cases give soundness $s_Z = \min\{c\beta, \frac{1}{6D_{[k]}}\} > 0$. $\square$

The construction generalizes to any dimension $k \geq 4$ and any three consecutive layers $V_{\ell-1}, V_\ell, V_{\ell+1}$ where $2 \leq \ell \leq k-2$. To obtain constant rate, it is enough to assume $\mu_i \geq 2k$ for $i \leq \ell$ and $\mu_i \leq 1/(2k)$ for $i > \ell$. This makes $V_{[\ell]}$, a block in the middle layer, larger than all other blocks combined by a constant factor.

References

- [ABO97] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error. In *Proceedings of STOC*, 1997. [1](#)
- [AD24] Ron Asherov and Irit Dinur. Bipartite unique neighbour expanders via Ramanujan graphs. *Entropy*, 26(4):348, 2024. [1.3](#)
- [AE18] Dorit Aharonov and Lior Eldar. Quantum locally testable codes. *SIAM Journal on Computing*, 47(4):1552–1581, 2018. [1](#), [1.3](#), [5.3](#), [6](#)
- [BE21] Nikolas P. Breuckmann and Jens N. Eberhardt. Balanced product quantum codes. *IEEE Transactions on Information Theory*, 67(10):6653–6674, 2021. [1.1](#)
- [BSGK⁺10] Eli Ben-Sasson, Venkatesan Guruswami, Tali Kaufman, Madhu Sudan, and Michael Viderman. Locally testable codes require redundant testers. *SIAM Journal on Computing*, 39(7):3230–3247, 2010. [1.1](#)
- [CRVW02] Michael Capalbo, Omer Reingold, Salil Vadhan, and Avi Wigderson. Randomness conductors and constant-degree expansion beyond the degree / 2 barrier. In *STOC*, pages 659–668, 2002. [1.1](#)
- [CS96] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Phys. Rev. A*, 54:1098–1105, Aug 1996. [5.2](#)
- [DEL⁺26] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Good locally testable codes. *Annals of Mathematics*, 203(2):511–553, 2026. [1](#), [1.1](#), [1.2](#)
- [Din07] Irit Dinur. The pcg theorem by gap amplification. *Journal of the ACM*, 54(3):12:1–12:44, 2007. [1](#)
- [DLV24] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of high-dimensional cubical complexes: with application to quantum locally testable codes. In *Proceedings of the 65th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 379–385, 2024. [1.1](#), [1.2](#)
- [EFMP23] Shai Evra, Brooke Feigon, Kathrin Maurischat, and Ori Parzanchevski. Ramanujan bigraphs. *arXiv preprint arXiv:2312.06507*, 2023. [1.3](#)
- [EFMP25] Shai Evra, Brooke Feigon, Kathrin Maurischat, and Ori Parzanchevski. Ramanujan bigraphs and applications. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2088–2096, 2025. [1.3](#)
- [EH17] Lior Eldar and Aram W. Harrow. Local hamiltonians whose ground states are hard to approximate. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 427–438, 2017. [5.3](#), [6](#)
- [EK24] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. *Journal of the American Mathematical Society*, 37(1):39–68, 2024. [1.4](#)

- [FK24] Uriya A. First and Tali Kaufman. Cosystolic expansion of sheaves on posets with applications to good 2-query locally testable codes and lifted codes, 2024. [1.4](#), [1.4](#), [4](#), [4.1](#)
- [Gal62] Robert G. Gallager. Low-density parity-check codes. *IRE Transactions on Information Theory*, 8(1):21–28, 1962. [1](#)
- [GGJ26] William Gay and Fernando Granha Jeronimo. Asymptotically good quantum locally testable codes. <https://arxiv.org/abs/2609.20780>, 2026. [1.2](#), [1.5](#)
- [GS06] Oded Goldreich and Madhu Sudan. Locally testable codes and pcps of almost-linear length. *Journal of the ACM*, 53(4):558–655, 2006. [1](#), [5.1](#)
- [HLM⁺25] Jun-Ting Hsieh, Alexander Lubotzky, Sidhanth Mohanty, Assaf Reiner, and Rachel Yun Zhang. Explicit lossless vertex expanders. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 894–911, 2025. [1.1](#), [1.3](#), [2.9](#)
- [Kah95] Nabil Kahale. Eigenvalues and expansion of regular graphs. *Journal of the ACM*, 42(5):1091–1106, 1995. [1.3](#)
- [Kit03] Alexei Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003. [1](#)
- [KK22] Amitay Kamber and Tali Kaufman. Combinatorics via closed orbits: number theoretic Ramanujan graphs are not unique neighbor expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 2022. Full version: arXiv:2103.04311. [1.3](#)
- [KKL16] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for Ramanujan complexes and topological expanders. *Geometric and Functional Analysis*, 26(1):250–287, 2016. [1.4](#)
- [KLZ96] Emanuel Knill, Raymond Laflamme, and Wojciech Zurek. Threshold accuracy for quantum computation. *arXiv preprint quant-ph/9610011*, 1996. [1](#)
- [LH22a] Ting-Chun Lin and Min-Hsiu Hsieh. c^3 -Locally Testable Codes from Lossless Expanders. *IEEE International Symposium on Information Theory (ISIT)*, pages 1175–1180, 2022. [1.1](#), [1.2](#), [1.3](#), [1.4](#), [3](#), [2.2](#), [5.2](#), [1](#)
- [LH22b] Ting-Chun Lin and Min-Hsiu Hsieh. Good quantum LDPC codes with linear time decoder from lossless expanders, 2022. [1.1](#), [1.2](#), [1.3](#), [1.4](#), [3](#), [2.2](#), [5.2](#)
- [LLZ22] Anthony Leverrier, Vivien Londe, and Gilles Zémor. Towards local testability for quantum coding. *Quantum*, 6:661, 02 2022. [5.1](#)
- [LMM16] Alexander Lubotzky, Roy Meshulam, and Shahar Mozes. Expansion of building-like complexes. *Groups, Geometry, and Dynamics*, 10(1):155–175, 2016. [3.1](#)
- [LPS88] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. [1.3](#)

- [PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, pages 375–388, 2022. [1](#), [1.1](#), [1.2](#), [4.5](#)
- [Sho96] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 56–65, 1996. [1](#)
- [SS96] Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Transactions on Information Theory*, 42(6):1710–1722, 1996. [1](#), [1.1](#), [5.2](#)
- [Ste96a] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 452(1954):2551–2577, 11 1996. [5.2](#)
- [Ste96b] Andrew M. Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793–797, 1996. [1](#)