

# Explicit Lower Bounds Against $\Omega(n)$ -Rounds of Sum-of-Squares

Max Hopkins\*

Ting-Chun Lin<sup>†</sup>

May 12, 2022

## Abstract

We construct an explicit family of 3-XOR instances hard for  $\Omega(n)$ -levels of the Sum-of-Squares (SoS) semi-definite programming hierarchy. Not only is this the first explicit construction to beat brute force search (beyond low-order improvements (Tulsiani 2021, Pratt 2021)), combined with standard gap amplification techniques it also matches the (optimal) hardness of random instances up to imperfect completeness (Grigoriev TCS 2001, Schoenebeck FOCS 2008).

Our result is based on a new form of *small-set* high dimensional expansion (SS-HDX) inspired by recent breakthroughs in locally testable and quantum LDPC codes. Adapting the recent framework of Dinur, Filmus, Harsha, and Tulsiani (ITCS 2021) for SoS lower bounds from the Ramanujan complex to this setting, we show any (bounded-degree) SS-HDX can be transformed into a highly unsatisfiable 3-XOR instance that cannot be refuted by  $\Omega(n)$ -levels of SoS. We then show Leverrier and Zémor's (Arxiv 2022) recent qLDPC construction gives the desired explicit family of bounded-degree SS-HDX. Incidentally, this gives the strongest known form of bi-directional high dimensional expansion to date.

---

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

<sup>†</sup>Department of Physics, UCSD, CA 92092, and Hong Hai Research Institute, Taipei. Email: til022@ucsd.edu.

# 1 Introduction

The Sum-of-Squares (SoS) semi-definite programming (SDP) hierarchy is one of the most powerful and widely studied algorithmic frameworks for approximating constraint satisfaction problems (CSPs) in theoretical computer science, yet very little is known about the structure of instances that are *hard* for the paradigm. Indeed, while it has long been known that *random* instances of CSPs are hard for Sum-of-Squares [Gri01b, Sch08, Tul09, BCK15, Cha16, KMOW17], there are essentially no *explicit* constructions of hard instances better than brute force search [DFHT20, Tul21, Pra22]. Leveraging recent breakthroughs in locally testable [DEL<sup>+</sup>21, LH22a] and quantum low-density parity-check (qLDPC) codes [PK21a, LZ22], we resolve this problem, giving the first explicit family of highly unsatisfiable CSPs that cannot be refuted by  $\Omega(n)$ -rounds of Sum-of-Squares.

**Theorem 1.1** (Main Result: Explicit 3-XOR Instances Hard for SoS). *There exist constants  $\mu_1, \mu_2 \in (0, 1)$  and an infinite family of 3-XOR instances constructable in deterministic polynomial time such that:*

1. *No assignment satisfies more than a  $1 - \mu_1$  fraction of constraints*
2. *No instance can be refuted by  $\mu_2 n$  levels of the corresponding Sum-of-Squares SDP Relaxation.*

Though Theorem 1.1 only exhibits an ‘integrality gap’ of 1 v.s  $1 - \mu_1$  (meaning the instance are  $(1 - \mu_1)$ -satisfiable but *look* fully satisfiable to SoS), combined with standard PCP-like reductions in the SoS hierarchy this gap can be amplified to  $1 - \varepsilon$  v.s  $\frac{1}{2} + \varepsilon$  for any  $\varepsilon > 0$  [Tul09, DFHT20], which matches the hardness of random 3-XOR instances up to imperfect completeness [Gri01b, Sch08].<sup>1</sup> In fact, it is worth noting that Theorem 1.1 is the first explicit family of CSPs to even beat more than  $O(\log(n))$  levels of the SoS hierarchy, which can be done either by unique neighbor expanders [Pra22, AC02] or (up to lower order factors) simply by brute force search [Tul21]. While explicit constructions against  $\Omega(n)$ -rounds of SoS were known in proof complexity (e.g. Tseitin formulas [Gri98], knapsack [Gri01a]), these examples do not lead to inapproximability since their satisfiability is not bounded away from 1.

Thus, at a high level, Theorem 1.1 provides the first example of an approximation problem with short witnesses of unsatisfiability that cannot be captured by the Sum-of-Squares proof system, settling (in the negative) the completeness of SoS in this setting. Furthermore, it is worth noting that 3-XOR is not somehow ‘special’ in this sense. As observed in [DFHT20] (who showed an analogous result for  $O(\sqrt{\log(n)})$ -levels of SoS), Theorem 1.1 also gives explicit hard instances across many types of CSPs by standard reduction techniques [Tul09], including instances with optimal integrality gaps for CSPs with approximation resistant predicates based on pairwise independent subgroups [Cha16, DFHT20].

## 1.1 High Dimensional Small-Set Expanders

Theorem 1.1 is based on a new form of *high dimensional expansion* (HDX), a nascent area of computer science and math that has already seen an impressive array of breakthrough results across areas such as coding theory [JST21, DEL<sup>+</sup>21, PK21a, LH22a], approximate sampling [KO20, ALOV19, AL20, ALO20], approximation algorithms [AJT19, BHKL22], analysis of boolean functions [DDFH18, BHKL21, GLL21], agreement testing [DK17, DD19], and, recently, Sum-of-Squares lower bounds [DFHT20]. While most of these works consider notions of expansion on *hypergraphs* (often called *simplicial complexes* in this setting), we take inspiration from recent breakthroughs on LTCs [DEL<sup>+</sup>21, LH22a] and quantum codes [PK21a, LZ22] and consider expansion on the more general class of *chain complexes*:

$$X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}.$$

<sup>1</sup>Indeed one can see such a gap is essentially optimal, as a random assignment to any 3-XOR instance will satisfy 1/2 the constraints in expectation.

Here  $X(0)$ ,  $X(1)$ , and  $X(2)$  are sets,  $\delta_0$  and  $\delta_1$  are linear maps (called the *co-boundary operators*),  $\partial_2$  and  $\partial_1$  are their transposes (called the *boundary operators*), and both satisfy  $\partial_1\partial_2 = 0$ ,  $\delta_1\delta_0 = 0$ .

Chain complexes admit a natural analog of boundary (edge) expansion in graphs called high-dimensional (co)-boundary expansion [LM06]. To see this, we first note an important inherent structural property of chain complexes: any function  $f \in \text{im}(\delta_0)$  (called a *co-boundary*) satisfies  $|\delta_1 f| = 0$ . A complex is called a  *$\rho$ -co-boundary expander* essentially when this is the only obstruction to  $|\delta_1 f|$  being large:

$$\forall f \in \mathbb{F}_2^{X(1)} : |\delta_1 f| \geq \rho \cdot d(f, \text{im}(\delta_0)).$$

For intuition, it is worth briefly discussing why this generalizes boundary expansion on graphs. Any graph  $G = (V, E)$  (or indeed hypergraph, see Section 4.2) can be written as a chain complex:

$$X : \mathbb{F}_2^\emptyset \begin{array}{c} \xrightarrow{\delta_0} \\ \xleftarrow{\partial_1} \end{array} \mathbb{F}_2^V \begin{array}{c} \xrightarrow{\delta_1} \\ \xleftarrow{\partial_2} \end{array} \mathbb{F}_2^E,$$

where  $\delta_0 f(v) = f(\emptyset)$ ,  $\delta_1 f((u, v)) = f(u) \oplus f(v)$ , and it is easily checked that  $\delta_1\delta_0 = 0$ . Notice that in this setting the only co-boundaries are  $\text{im}(\delta_0) = \{\emptyset, V\}$ , and furthermore that for any  $S \subset V$  and  $e \in E$ , the value of  $\delta_1 1_S$  on  $e$  is 1 iff  $e$  crosses the cut defined by  $S$ . This implies the ratio  $\frac{|\delta_1 1_S|}{d(1_S, \text{im}(\delta_0))} = \frac{E(S, V \setminus S)}{\min\{|S|, |V \setminus S|\}}$ , which is just the standard boundary expansion of  $G$ !

Unfortunately, while standard boundary expansion on (random) graphs has been quite useful for proving SoS lower bounds in the past [BSW99, Gri01b, Sch08], high dimensional co-boundary expansion seems to be too strong a notion for this setting: good (co)-boundary expanders are not known to exist (even probabilistically), and their structure is prohibitively restrictive in other senses as well.<sup>2</sup> We avoid these issues by introducing a simple relaxation of boundary expansion to *small-sets*:

**Definition 1.2** (Small-set (Co)-Boundary Expansion). *We call  $X$  a  $(\rho_1, \rho_2)$ -small-set boundary expander if the weight of any ‘small’ function  $f \in \mathbb{F}_2^{X(1)}$  satisfying  $|f| \leq \rho_1 |X(1)|$  expands:*

$$|\partial_1 f| \geq \rho_2 \cdot d(f, \text{im}(\partial_2)).$$

*Similarly,  $X$  is a  $(\rho_1, \rho_2)$ -small-set co-boundary expander if all  $f \in \mathbb{F}_2^{X(1)}$  s.t.  $|f| \leq \rho_1 |X(1)|$  satisfy:*

$$|\delta_1 f| \geq \rho_2 \cdot d(f, \text{im}(\delta_0)).$$

*We call  $X$  a  $(\rho_1, \rho_2)$ -small-set HDX (SS-HDX) if it satisfies both the above conditions.*

Small-set (co)-boundary expansion is a direct generalization of small-set expansion on graphs, a notion that lies at the heart of many problems in hardness of approximation (especially with respect to Khot’s unique games conjecture [Kho02, RST12, KMS18]). Similar notions of high dimensional small-set expansion have also been considered on simplicial complexes in work towards building good co-systolic expanders (a different weakening of co-boundary expansion) [KKL14, EK16, KM18, KM21]. We discuss these notions and their relation to our definition in more depth in Section 3.1.

In the next section, we will show how SS-HDX naturally lead to hard instances of XOR for Sum-of-Squares (largely following a similar result of Dinur, Filmus, Harsha, and Tulsiani [DFHT20] for the LSV complex [LSV05]), giving the first connection between hardness of approximation and *high dimensional* small-set expanders. Thus, Theorem 1.1 boils down to constructing an infinite family of SS-HDX on a growing number of vertices, each of which can be constructed in deterministic polynomial time. While this may seem hopelessly strong, a weaker variant of these requirements was very recently achieved in breakthrough constructions of qLPDC codes by [PK21a, LZ22]. Indeed, it turns out these known constructions are already enough: we show Leverrier and Zémor’s [LZ22] recent qLPDC codes are in fact small-set HDX as well.

<sup>2</sup>We’ll discuss this issue in Section 2, but in brief co-boundary expansion implies  $\ker(\delta_1) = \text{im}(\delta_0)$ . Like [DFHT20], our instances will rely on a function in  $\ker(\delta_1) \setminus \text{im}(\delta_0)$  to enforce global structure on the CSP that cannot be detected through local algorithms like Sum-of-Squares.

**Theorem 1.3** (Small-Set HDX Exist (informal Theorem 8.1)). *There exist constants  $\rho_1, \rho_2 \in (0, 1)$  and an explicit (polynomial time constructable) infinite family of bounded-degree<sup>3</sup> (3-term) chain complexes  $\{X_i\}$  satisfying:*

1.  $X_i$  has non-trivial ‘co-homology,’ i.e.  $\text{im}(\delta_0) \neq \ker(\delta_1)$
2.  $X_i$  is a  $(\rho_1, \rho_2)$ -SS HDX.

The guarantees of Theorem 1.3 are stronger than those originally proved by Leverrier and Zémor [LZ22] (see Section 3.1 for discussion), and give the strongest known form of bi-directional high dimensional expansion to date.<sup>4</sup> Indeed the expansion is so strong that if one could remove the small-set requirement<sup>5</sup> or prove similar bounds for a 5-term chain complex, it would resolve the qLTC conjecture [KKL14, EH17, LH22a], a major open problem in quantum computation.

## 2 Proof Overview

We now overview the constructions and proof techniques underlying our main result (Theorem 1.1). Broadly speaking, this breaks into two main steps:

1. Show any SS-HDX implies a hard instance of 3-XOR
2. Construct an explicit infinite family of SS-HDX.

To start, it will be useful to cover some basic background on CSPs, Sum-of-Squares, and chain complexes in a bit more detail. A more formal treatment is given in Section 4 and Section 7.

### 2.1 Background

In this work, we study the limitations of the Sum-of-Squares proof system for refuting MAX- $k$ -XOR, a widely studied class of constraint satisfaction problems (CSPs). An instance of MAX- $k$ -XOR  $\mathcal{I}$  consists of a set of variables  $\{x_i\}_{i \in [n]}$  and constraints  $\{C_i\}_{i \in [m]}$ , where each  $C_i$  is a boolean function of the form:

$$C_i(x) = \mathbf{1} \{x_{i_1} \oplus \dots \oplus x_{i_j} = b_i\},$$

where  $j = j(i) \leq k$  and  $\{i_1, \dots, i_j\} \subset [n]$ . If all constraints have exactly  $k$  variables, we say  $\mathcal{I}$  is an instance of  $k$ -XOR. We will usually omit the indicator  $\mathbf{1}$  from notation when clear from context. The *value* of  $\mathcal{I}$  is the maximum fraction of constraints that can be satisfied by any assignment, and we say  $\mathcal{I}$  is  $(1 - \mu)$ -satisfiable if there exists an assignment satisfying at least a  $(1 - \mu)$  fraction of constraints. We call an infinite family of instances  $\{\mathcal{I}_i\}$  *explicit* if each instance can be constructed in deterministic polynomial time in the number of variables.

The Sum-of-Squares semi-definite programming hierarchy is a powerful algorithmic framework for approximating the value of any CSP (or more generally for solving constrained polynomial optimization problems). The hierarchy consists of *rounds* or *levels* of progressively stronger SDP relaxations (see Algorithm 1). For the moment, it is enough to know that the round- $t$  SoS relaxation is local<sup>6</sup> in the sense that

<sup>3</sup>A complex is bounded degree roughly if each element in  $X(i)$  only has constantly many neighbors with respect to the boundary and co-boundary operators. See Section 4 for an exact definition.

<sup>4</sup>In fact it’s worth mentioning we actually prove a stronger guarantee regarding *local* functions. See Remark 8.9 and discussion in Section 3.1.

<sup>5</sup>Though it is worth noting one must be careful that the dimension of the cohomology stays large, which requires weakening the expansion guarantee to a related notion called (co)-systolic expansion (the correct notion for qLTC regardless) [EH17].

<sup>6</sup>We note the relaxation does have (low-degree) global consistency checks, so it is not fully a local algorithm in this sense.

it ranges over subsets of variables of size at most  $t$ . We will cover more details on the SoS framework as they arise.

Finally, it will be useful to have some basic terminology corresponding to chain complexes. Recall that a chain complex is a sequence  $X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}$  such that  $\partial_1 \partial_2 = 0$ ,  $\delta_1 \delta_0 = 0$ . Functions in the image of  $\partial_2$  and  $\delta_0$  are called *boundaries* and *co-boundaries* respectively, and are denoted:

$$\text{im}(\partial_2) = B_1, \quad \text{im}(\delta_0) = B^1.$$

Functions in the kernel of  $\partial_1$  and  $\delta_1$  are called *cycles* and *co-cycles* respectively, and are denoted:

$$\text{ker}(\partial_1) = Z_1, \quad \text{ker}(\delta_1) = Z^1.$$

The structure of a chain complex promises that  $B_1 \subset Z_1 \subset \mathbb{F}_2^{X(1)}$  and  $B^1 \subset Z^1 \subset \mathbb{F}_2^{X(1)}$ . This leads to notions of *homology* and *co-homology* given by (co)-cycles mod (co)-boundary and respectively denoted:

$$H_1 = Z_1/B_1, \quad H^1 = Z^1/B^1,$$

where  $G/H$  denotes the quotient group. A complex has non-trivial co-homology if  $B^1 \neq Z^1$ .

## 2.2 From SS-HDX to Hardness

With notation out of the way, we can now discuss how to transform an expanding chain complex into a hard instance of 3-XOR. Before we give an informal theorem statement to this effect, it is instructive to overview how one even relates a CSP to a chain complex at all. To this end, let's first recall the classical construction of CSPs (also frequently seen in coding theory) based upon a bipartite graph  $B = (L, R, E)$ . In this setting, elements in  $L$  correspond to variables  $\{x_v\}_{v \in L}$ , and elements in  $R$  correspond to the set of constraints  $\{C_r\}_{r \in R}$ . Fixing some assignment  $\beta \in \{0, 1\}^R$  to constraints, the XOR instance classically associated with the graph  $B$  is characterized by ensuring the (mod 2) sum across neighbors of each  $r \in R$  is given by  $\beta(r)$ :

$$C_r := \left\{ \sum_{v \in N(r)} x_v = \beta(r) \pmod{2} \right\}. \quad (1)$$

In prior hardness constructions,  $B$  is typically picked at random in order to satisfy strong expansion properties, while  $\beta$  is typically chosen at random to ensure un-satisfiability (see e.g. [Gri01b, Sch08, KMOW17]). While it is sometimes possible to de-randomize the choice of  $B$  and retain good inapproximability guarantees, no de-randomization of  $\beta$  better than brute force search over  $\log(n)$ -size instances was known up until this point.

The basic form of our XOR instances from chain complexes is actually very similar to Equation (1) (indeed they can be viewed as a special instantiation of this framework). Recall that a chain complex is a sequence:

$$X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)},$$

and in particular that the co-boundary operator  $\delta_0 : \mathbb{F}_2^{X(0)} \rightarrow \mathbb{F}_2^{X(1)}$  is a linear map. To define an instance of XOR on  $X$ , we simply move to the *graph representation* of  $\delta_0$ . Namely, recall that any linear operator mapping from  $\mathbb{F}_2^{X(0)}$  to  $\mathbb{F}_2^{X(1)}$  can be written as an  $(|X(1)| \times |X(0)|)$ -dimensional matrix over  $\mathbb{F}_2$ . We can think of this matrix as the bipartite adjacency matrix of a graph on left vertex set  $L = X(0)$  and right vertex set  $R = X(1)$ . Thus given a function  $\beta \in \mathbb{F}_2^{X(1)}$ , we construct the associated XOR instance, denoted  $\mathcal{I}_{X,\beta}$  as

in Equation (1) by adding the constraint for each  $r \in X(1)$ :

$$C_r := \left\{ \sum_{\substack{v \in X(0): \\ e_r^T \delta_0 e_v = 1}} x_v = \beta(r) \pmod{2} \right\}, \quad (2)$$

where  $e_v \in \mathbb{F}_2^{X(0)}$  and  $e_r \in \mathbb{F}_2^{X(1)}$  are the standard basis vectors associated to  $v \in X(0)$  and  $r \in X(1)$ . Note that  $e_r^T \delta_0 \in \mathbb{F}_2^{X(0)}$  is just the list of neighbors of  $r$ , so this is indeed an instantiation of the standard bipartite framework. We note that this construction also generalizes the recent approach of [DFHT20] who built XOR instances via a 3-dimensional simplicial complex (4-uniform hypergraph) by letting *triangles* correspond to constraints, and *edges* correspond to variables. This is exactly the result of the above construction when applied to the natural chain complex associated with a 3-dimensional simplicial complex (see Section 4.2 for further details).

So far, we have not used the fact that  $\delta_0$  is part of a chain complex, or even the fact that the higher dimensional component  $X(2)$  exists at all. This structure comes into play in the choice of  $\beta$ . Notice that by construction, the instance corresponding to  $X$  and a choice of  $\beta$  is satisfiable exactly when  $\beta$  is a co-boundary. Following the framework laid out in [DFHT20], the idea is to choose  $\beta \in Z^1 \setminus B^1$ , a function which is a co-cycle, but not a co-boundary. On a sufficiently expanding complex, this choice induces *global structure* on the XOR instance that cannot be captured by *local views* of the complex, where both the homology and co-homology look trivial. Since Sum-of-Squares only looks over local views in this sense, this leads to the following direct translation between SS-HDX and hard instances of XOR.

**Theorem 2.1** (SS-HDX  $\implies$  Hard XOR Instance (Informal Theorem 6.4)). *Let  $X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}$  be an SS-HDX with non-trivial co-homology. Then there exist  $\mu_1, \mu_2 \in (0, 1)$  such that for any  $\beta \in Z^1 \setminus B^1$ , the associated XOR instance  $\mathcal{I}_{X,\beta}$  satisfies:*

1. *Soundness:  $\mathcal{I}_{X,\beta}$  is at most  $(1 - \mu_1)$ -satisfiable,*
2. *Completeness:  $\mathcal{I}_{X,\beta}$  cannot be refuted by  $\mu_2|X(0)|$  levels of the SoS hierarchy.*

Before moving on to the construction of SS-HDX, let's discuss how small-set expansion implies soundness and completeness for these instances. Soundness, the simpler of the two, intuitively comes from the fact that small-set co-boundary expansion promises that any element in  $Z^1 \setminus B^1$  must be *far from the co-boundary*.<sup>7</sup> Recall that by construction, the instance  $\mathcal{I}_{X,\beta}$  is satisfiable exactly when  $\beta \in \mathbb{F}_2^{X(1)}$  is a co-boundary. Intuitively one might then expect that functions which are far from the co-boundary would therefore be far from satisfiable. Indeed this intuition holds true—it is easy to show this robust version of the statement holds for small-set co-boundary expanders, and therefore that our instances are far from satisfiable as well.

Completeness is somewhat trickier and, unlike soundness, does actually require the full power of small-set boundary expansion. We stated earlier that the completeness of our instances, much like those of [DFHT20], comes from the fact that the global structure of (co)-homology cannot be detected through local views of the complex. This is formalized by observing that small-set boundary expansion can be equivalently re-stated as the following *isoperimetric inequality* (see Lemma 5.4): “small, minimal<sup>8</sup> functions have large boundaries.” Largely following [DFHT20] (who use a much weaker isoperimetric inequality for the LSV complex due

<sup>7</sup>It is worth noting that this property, called *co-systolic distance*, is quite well studied. Indeed as we will soon discuss it is exactly the property needed (in both directions) to build good qLDPC codes [PK21a], and was also used directly by [DFHT20] to prove soundness of their 3-XOR instances by the same argument stated here.

<sup>8</sup>A function  $f \in \mathbb{F}_2^{X(1)}$  is said to be minimal if adding any boundary can only increase its size (Hamming weight).

to Gromov [Gro83]), the idea is then to combine this fact with the classical arguments of Ben-Sasson and Wigderson [BSW99] to show that the width<sup>9</sup> of any refutation of  $\mathcal{I}_{X,\beta}$  in the  $\oplus$ -resolution proof system<sup>10</sup> is large. Since Schoenebeck [Sch08] showed any such bound transfers to a completeness lower bound for Sum-of-Squares, this completes the proof.

In slightly more detail, a refutation in the  $\oplus$ -resolution system can be viewed as an (in-degree two) DAG where leaves correspond to the original XOR constraints, internal nodes correspond to the XOR of their two parents (as in the  $\oplus$ -resolution derivation rule), and the root derives the contradiction  $0 = 1$ . Recall that each element  $s \in X(1)$  corresponds to a constraint in our XOR instance. Following [DFHT20], the idea is to assign a function in  $h_v \in \mathbb{F}_2^{X(1)}$  for each node  $v$  in the DAG that tracks which XOR constraints are being used at that node. The *boundary* of this function,  $\partial_1 h_v \in \mathbb{F}_2^{X(0)}$ , is exactly the set of variables appearing in the equation corresponding to node  $v$ . Thus lower bounding the width of the refutation boils down to finding a node with large boundary.

This is where small-set boundary expansion (namely the isoperimetric formulation) finally comes into play. In particular, the corresponding inequality states that it is enough to find a node  $v$  of ‘medium’ weight:<sup>11</sup> small enough that one can apply the inequality, but large enough to result in a large boundary. This can be done by fairly standard potential arguments (see e.g. [BSW99, DFHT20]) where one sets up a potential function tracking this weight throughout the DAG, and argues that the leaves have small potential, the root has large potential, and that potential is sub-additive. This implies the existence of an interior node with medium potential and completes the proof. The details are given in Section 6.

Finally, before moving on to overiewing our construction of SS-HDX, we note that except in very special cases (e.g. the simplicial complexes considered in [DFHT20]), the CSPs given by Equation (2) (and therefore also Theorem 2.1) are actually instance of MAX- $k$ -XOR, not 3-XOR, where  $k$  is given by the maximum degree of the complex. As it turns out, this is not a significant issue because the SS-HDX we construct in the next section are *bounded degree*, meaning not only that every constraint in the XOR has a constant number of variables, but also that every variable only appears in a constant number of constraints. This observation allows us to move to hard instances of 3-XOR by standard NP-reduction type arguments within the SoS hierarchy [Sch08, Tul09] while only losing constant factors in the soundness and levels of hardness for SoS.

## 2.3 Constructing SS-HDX

Now that we know how to transform an expanding chain complex into a hard instance of 3-XOR, we turn our attention to the construction of such complexes. Our method relies on recent breakthroughs on LTCs [DEL<sup>+</sup>21, LH22a] and quantum LDPC codes [PK21a, LZ22]. As such, we’ll split this section into three parts: a review of the connection between quantum LDPC codes and expanding chain complexes, the recent qLDPC construction of Leverrier and Zémor [LZ22], and our proof of small-set (co)-boundary expansion.

### 2.3.1 Quantum LDPC Codes and Chain Complexes

A classical error correcting code is a method of encoding  $k$  classical bits into  $n > k$  classical bits such that it is possible to recover the original bit string even if the encoded string becomes corrupted. We will consider *linear codes*, which are defined by a linear operator  $M : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{n-k}$  called the parity check matrix,<sup>12</sup> where the corresponding code  $\mathcal{C} := \ker M$ .

<sup>9</sup>The width of a refutation is the largest number of variables appearing in any equation.

<sup>10</sup>In this proof system, one is allowed to combine linear equations (equivalently XOR constraints)  $\ell_1 = b_1$  and  $\ell_2 = b_2$  to derive the equation  $\ell_1 \oplus \ell_2 = b_1 \oplus b_2$ . A refutation is a proof based on this rule deriving a contradiction ( $0 = 1$ ), which is equivalent in our setting to showing the XOR instance is unsatisfiable.

<sup>11</sup>We note that weight here is not just the standard Hamming weight, but must take into account distance from the boundary as well. See Section 6.

<sup>12</sup>We note the parity check matrix is traditionally denoted by ‘ $H$ ,’ but this conflicts with the notation for homology.

Similar to the classical setting, a *quantum* code encodes quantum bits into a larger number of quantum bits, but is resistant to two types of corruption: the  $X$ -type errors (bit flips) and the  $Z$ -type errors (phase flips). In this work, we will focus on a popular notion of quantum codes called *CSS-codes* [CS96, Ste96], which come with the benefit of having an entirely classical interpretation. In particular, a length  $n$  CSS-code is made up of two classical codes  $\mathcal{C}_0 := \ker M_0 \subset \mathbb{F}_2^n$  and  $\mathcal{C}_1 := \ker M_1 \subset \mathbb{F}_2^n$  such that  $\mathcal{C}_0^\perp \subset \mathcal{C}_1$ , or equivalently  $M_1 \cdot M_0^T = 0$ .<sup>13</sup> The dimension of the code is defined as  $k = \dim \mathcal{C}_0 - \dim \mathcal{C}_1^\perp$ , and its *distance* (which measures how much corruption it can handle) is defined as  $d = \min(d_x, d_z)$  where

$$d_x = \min_{v \in \mathcal{C}_0 \setminus \mathcal{C}_1^\perp} |v|, d_z = \min_{v \in \mathcal{C}_1 \setminus \mathcal{C}_0^\perp} |v|$$

and  $d_x$  ( $d_z$ ) is called the  $X$ -distance ( $Z$ -distance). The quantum low-density parity-check (LDPC) conjecture, recently resolved by [PK21a], states that there exists a family of quantum CSS codes with linear dimension and distance,  $k = \Theta(n)$  and  $d = \Theta(n)$ , where  $M_0$  and  $M_1$  have at most some constant number of ones in any row or column (and thus are ‘low-density’ parity check matrices).

Since we are promised by definition that  $M_1 \cdot M_0^T = 0$ , it is easy to see that any CSS-code induces the following chain complex:

$$X : \mathbb{F}_2^{m_0} \begin{array}{c} \xrightarrow{M_0^T} \\ \xleftarrow{M_0} \end{array} \mathbb{F}_2^n \begin{array}{c} \xrightarrow{M_1} \\ \xleftarrow{M_1^T} \end{array} \mathbb{F}_2^{m_1},$$

where  $m_i = \dim(\text{im}(M_i))$ . Indeed the same holds in reverse as well, given a chain complex

$$X : \mathbb{F}_2^{X(0)} \begin{array}{c} \xrightarrow{\delta_0} \\ \xleftarrow{\partial_1} \end{array} \mathbb{F}_2^{X(1)} \begin{array}{c} \xrightarrow{\delta_1} \\ \xleftarrow{\partial_2} \end{array} \mathbb{F}_2^{X(2)},$$

one obtains a quantum CSS code by letting  $M_0 := \partial_1$ , and  $M_1 := \delta_1$ .

In fact, it turns out this equivalence between quantum CSS codes and chain complexes runs deeper: all of the discussed properties (e.g. distance, LDPC) have analogs in the homological language we developed in the previous section. The classical codes  $\mathcal{C}_0$  and  $\mathcal{C}_1$ , for instance, correspond to the cycles and co-cycles of the chain complex ( $\mathcal{C}_0 = Z_1, \mathcal{C}_1 = Z^1$ ), while the dual codes  $\mathcal{C}_0^\perp$  and  $\mathcal{C}_1^\perp$  correspond to the co-boundaries and boundaries ( $\mathcal{C}_0^\perp = B^1, \mathcal{C}_1^\perp = B_1$ ). The dimension of the code  $k$  corresponds to the dimension of the co-homology ( $k = \dim H^1$ ), and the maximum degree of the complex corresponds to the maximum density of the parity check codes (so the bounded-degree and LDPC conditions are equivalent). Finally, the  $X$ -distance and  $Z$ -distance of the code correspond to what is known as the *(co)-systolic distance* of the chain complex, the minimum weight of any (co)-cycle that is not a (co)-boundary:

$$d_x = \min_{v \in \mathcal{C}_0 \setminus \mathcal{C}_1^\perp} |v| = \min_{v \in Z_1 \setminus B_1} |v|,$$

$$d_z = \min_{v \in \mathcal{C}_1 \setminus \mathcal{C}_0^\perp} |v| = \min_{v \in Z^1 \setminus B^1} |v|.$$

In [PK21a] and [LZ22], the authors construct two different explicit families of good quantum LDPC codes. This partially solves our problem since the codes correspond to a family of bounded-degree chain complexes with non-trivial co-homology and linear co-systolic distance (which is enough to imply soundness of our XOR construction). We will show these complexes in fact satisfy the stronger small-set (co)-boundary expansion condition, which as discussed in the previous section further implies completeness and (up to reduction to 3-XOR) finishes the proof of Theorem 1.1.

<sup>13</sup>Here  $\mathcal{C}_0^\perp$  denotes the dual code, consisting of all elements orthogonal to  $\mathcal{C}_0$ . This code is generated by the transpose of the parity check matrix  $M_0^T$ .

### 2.3.2 Leverrier and Zémor’s qLDPC Codes

Before discussing the proof, we need to overview the original construction of [LZ22]. A significantly more detailed description of the construction and its associated components is given in Section 7 and Section 8.

Leverrier and Zémor’s qLDPC codes are based on a classical object called a *Tanner code* [Tan81]. Given an  $n_0$ -regular graph  $\mathcal{G} = (V, E)$  and a linear code  $C$  of length  $n_0$ , the Tanner code  $T(\mathcal{G}, C) \subset \mathbb{F}_2^E$  is

$$\{c \in \mathbb{F}_2^E : \forall v \in V, c|_{E(v)} \in C\},$$

where  $c|_{E(v)} \in \mathbb{F}_2^{n_0}$  is the vector formed by the values on the edges incident to  $v$ . Tanner codes have long been used in coding theory. The main insight of [LZ22] was to observe that one can construct a quantum CSS code via two Tanner codes coming from a higher-dimensional object called the *left-right Cayley complex*, recently developed in [DEL<sup>+</sup>21] to construct c3-LTCs.

The left-right Cayley complex corresponding to a group  $G$  and two sets of generators  $A = A^{-1}$  and  $B = B^{-1}$  consists of a vertex set  $V = G$ , edges given by (left) Cayley graph  $C(G, A)$  and (right) Cayley graph  $C(G, B)$ , and higher-dimensional ‘squares’ of the form  $\{g, ag, gb, agb\}$  for  $g \in G, a \in A, b \in B$ . More formally, [LZ22] consider the *double cover* of this complex where:

- The vertices are  $V = V_0 \cup V_1$  where  $V_0 = G \times \{0\}$  and  $V_1 = G \times \{1\}$ .
- The ‘ $A$ -edges’ and ‘ $B$ -edges’ are respectively:

$$E_A = \{(g, 0), (ag, 1)\} : g \in G, a \in A\}, E_B = \{(g, 0), (gb, 1)\} : g \in G, b \in B\}.$$

- The squares are

$$F = \{(g, 0), (ag, 1), (gb, 1), (agb, 0)\} : g \in G, a \in A, b \in B\}.$$

Notice each square contains exactly two vertices in  $V_0$  and two vertices in  $V_1$ . This allows us to think of each square as an *edge* between two vertices in  $V_0$  (or  $V_1$ ) and to define corresponding graphs  $\mathcal{G}_0^\square = (V_0, F)$  and  $\mathcal{G}_1^\square = (V_1, F)$ . The local view around each vertex in  $(g, i) \in \mathcal{G}_i^\square$  then corresponds to the squares  $\{(g, i), (ag, 1 - i), (gb, 1 - i), (agb, i)\}$  for  $a \in A, b \in B$ . Assuming  $|A| = |B| = \Delta$  for some constant  $\Delta$ , we will always think about these local views as square matrices with rows indexed by  $A$  and columns indexed by  $B$ .

Leverrier and Zémor [LZ22] observed that the Tanner codes associated to these graphs,  $\mathcal{C}_0 = T(\mathcal{G}_0^\square, C_0^\perp)$  and  $\mathcal{C}_1 = T(\mathcal{G}_1^\square, C_1^\perp)$ , give a quantum CSS code (i.e. satisfy  $C_0^\perp \subset C_1$ ) whenever the associated local codes  $C_0 = C_A \otimes C_B$  and  $C_1 = C_A^\perp \otimes C_B^\perp$  are tensors<sup>14</sup> of linear codes  $C_A \subseteq \mathbb{F}_2^A$  and  $C_B \subseteq \mathbb{F}_2^B$ . Furthermore, they showed that whenever  $C_A, C_B, C_A^\perp, C_B^\perp$  have linear distance and the codes  $C_1^\perp = C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$ , and  $C_0^\perp = C_A^\perp \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B^\perp$  satisfy certain robustness properties (see Section 7.5), then the associated quantum code has linear distance. [LZ22] complete their construction by showing random base codes  $C_A, C_B$  satisfy these properties with high probability. Note that because these base codes are constant size, this final step can be brute-forced to maintain explicitness of the construction.

### 2.3.3 Proving Small-Set (Co)-Boundary Expansion

With [LZ22]’s construction in hand, we can now sketch the proof of small-set (co)-boundary expansion. As mentioned previously, all other major requirements (e.g. non-trivial homology, bounded-degree) already follow from the fact that the complex corresponds to a good qLDPC code. We will focus here on proving

<sup>14</sup>The tensor code  $C_A \otimes C_B$  is the set of matrices whose rows are given by elements of  $C_B$  and columns are given by elements of  $C_A$ .

small-set  $co$ -boundary expansion in particular, but we note that small-set boundary expansion follows the same argument by symmetry of [LZ22]’s construction.

With this in mind, recall that small-set co-boundary expansion can equivalently be phrased as an isoperimetric inequality for small, minimal functions (see Lemma 5.4). In particular, to show small-set co-boundary expansion for the chain complex

$$X : \mathbb{F}_2^{m_0} \xrightarrow{\delta_0 := C_0^T} \mathbb{F}_2^n \xrightarrow{\delta_1 := C_1} \mathbb{F}_2^{m_1},$$

it is enough to show there exist constants  $\rho_1, \rho_2 \in (0, 1)$  such that any minimal  $x \in \mathbb{F}_2^n$  with weight  $|x| \leq \rho_1 n$  has large boundary:  $|\delta_1 x| \geq \rho_2 |x|$ . We proceed by contradiction. Assuming  $|\delta_1 x| < \rho_2 |x|$ , we will show  $x$  is not minimal by finding  $y \in B^1$  such that  $|x + y| < |x|$ .

The proof of this fact largely follows the technique of [LZ22] for proving the weaker co-systolic distance property. The main difference is that while [LZ22] only consider functions  $x \in \mathbb{F}_2^n$  that are co-cycles, we consider arbitrary functions. In particular, recall that the co-cycles in our construction correspond to codewords in the Tanner code  $T(\mathcal{G}_1^n, C_1^\perp)$ , or equivalently to functions  $x \in \mathbb{F}_2^n$  whose ‘local view’ around each vertex  $(g, 1) \in V_1$  is given by a codeword of  $C_1^\perp$ . Since our functions do not a priori have this structure, we will need to track the set of ‘violations’ coming from local views that are not codewords (this essentially corresponds to where  $\delta_1 x$  is non-zero).

To this end, recall  $x$  is a bit string indexed by the squares of the double-covered Cayley complex, and let  $S \subset V_1$  denote the set of vertices incident to any square in  $x$ . We partition  $S$  into three parts: the violated vertices  $S_v$ , the normal vertices  $S_n$ , and the exceptional vertices  $S_e$ . A vertex is *violated* if the local view of  $x$  around the vertex does not form a codeword in  $C_1^\perp$ . When the local view does form a codeword, if the codeword has weight less than  $w := \Delta^{3/2-\varepsilon}$  we call it *normal*, and otherwise call it *exceptional*. This weight-based distinction comes from the robustness condition of the local tensor code. We cover this in detail in Section 7.5, but for the moment it is sufficient to think of robustness as a structural condition forcing codewords with weight less than  $w$  to be zero outside of a small number of rows and columns. In particular, this promises that each column (respectively row) in the local view of a normal vertex is at most  $O(\Delta^{1/2-\varepsilon})$  away from a codeword in  $C_A$  (respectively  $C_B$ ).

Following [LZ22], our goal is now to find a vertex  $v \in V_0$  that shares  $\Omega(\Delta)$  columns or rows with  $S_n$ . As long as  $S_e$  and  $S_v$  are not too large compared to  $S_n$ , robustness of the code then implies the local view of  $v$  is within  $O(\Delta^{3/2+\varepsilon})$  of a codeword  $c \in C_A \otimes C_B$ , but also has total weight  $\Omega(\Delta^2)$ .<sup>15</sup> This means we can construct a vector  $y \in B^1$  by defining  $y$  to be  $c$  on the local view of  $v$  and 0 everywhere else. Since  $x + y$  and  $x$  match outside the local view (where  $x$  has weight  $\Omega(\Delta^2)$  and  $x + y$  has weight  $O(\Delta^{3/2+\varepsilon})$ ), this implies  $|x + y| < |x|$  as desired.

It therefore remains to find such a vertex  $v \in V_0$ , which is the main technical component of the proof. Let  $T \subset V_0$  be the vertices that share at least one ‘heavy’ column or row with a normal vertex (that is one with many 1s). One can equivalently think of this as an edge between  $V_0$  and  $V_1$  that is ‘heavy’ in the sense that it is contained in many squares in  $x$ . The idea is then to show that there are many such heavy edges passing between  $S$  and  $T$ . Using expansion of the underlying graph and our assumption  $|\delta_1 x| < \rho_2 |x|$ , one can prove that  $T$ ,  $S_e$ , and  $S_v$  are small compared to  $S_n$ . This implies that a typical vertex in  $T$  has not just one, but  $\Omega(\Delta)$  heavy edges to  $S_n$ , which in turn corresponds to sharing  $\Omega(\Delta)$  rows and columns with normal vertices and completes the proof.

<sup>15</sup>We note  $C_A$  and  $C_B$  can be chosen to have linear distance to ensure this.

## 3 Discussion

### 3.1 Related Work

**Sum-of-Squares Lower Bounds:** At a conceptual level, our work fits into a long line of research on the limitations of Sum-of-Squares and related proof systems (e.g. Nullstellensatz [BIK<sup>+</sup>96, Gri98], Polynomial Calculus [CEI96, BGIP01]), and LP/SDP hierarchies (e.g. Sherali-Adams [CMM09, MS09, BGMT12], Lovász-Schrijver [AAT05, STT07, GMPT10]). Most relevant to our setting is the line of work on Sum-of-Squares lower bounds initiated by Grigoriev [Gri01b] (and later independently Schoenebeck [Sch08]), who used boundary expansion to prove random 3-XOR instances cannot be refuted by  $\Omega(n)$  levels of SoS. This led to a number of works improving integrality gaps for more general classes of random  $k$ -CSPs [Tul09, BCK15, Cha16, KMOW17] along with a number of other combinatorial optimization problems by reduction [Sch08, Tul09].

In a sense, these prior works on SoS lower bounds for random instances can be viewed as increasingly strong and general formulations of the statement: ‘Sum-of-Squares fails to capture the probabilistic method.’ In contrast, Dinur, Filmus, Harsha, and Tulsiani [DFHT20] recently exhibited the first *explicit* families of CSPs hard for Sum-of-Squares based on an algebraic, highly structured family of objects called Ramanujan (or LSV) complexes [LSV05], suggesting a new paradigm of hardness for structured instances. Due to the poor systolic expansion of the Ramanujan complex, [DFHT20]’s bounds only hold up to  $O(\sqrt{\log(n)})$  rounds of SoS as compared to  $\Omega(n)$  levels for random instances. Nevertheless, the authors conjectured it might be possible to use such anti-random objects to fool  $\Omega(n)$  levels as well. Our work can be viewed as a confirmation of this general hypothesis: anti-random structure (in particular certain *algebraic* structure) is indeed as hard as random for Sum-of-Squares.

**High Dimensional Expansion:** High dimensional expansion in the form we consider (i.e. topological expansion) was originally introduced by Linial and Meshulam [LM06] to study the vanishing of cohomology on random simplicial complexes, and independently by Gromov [Gro10] to study the topological overlapping principle. While our particular notion of small-set (co)-boundary expansion has not been studied in the literature, a stronger isoperimetric inequality for small, *locally* minimal<sup>16</sup> functions was used by Kaufman, Kazhdan, and Lubotzky [KKL14] to show the existence of bounded degree co-systolic expanders (another weakening of co-boundary expansion that replaces distance from  $B^1$  with distance from  $Z^1$ ), and later in [EK16, KM18, KM21] as well. A similar strategy was recently employed by Lin and Hsieh to construct c3-LTCs [LH22a] and later (conditional) qLDPC codes [LH22b]. It is worth noting that this stronger condition actually holds for our construction as well (see Remark 8.9).

**Quantum Codes and LTCs:** Quantum LDPC and locally testable codes have long been known to share a close connection with topological notions of high dimensional expansion (see e.g. [EH17]). Indeed it was qLDPC constructions based on the Ramanujan complex [EKZ20, KT21] that first broke the  $\sqrt{n}$  distance barrier and started the race to good qLDPCs [EKZ20, BE21, PK21b, KT21, HHO21, PK21a, JMO<sup>+</sup>21, LZ22]. As discussed in Section 2.3.1, qLDPC codes satisfy a weaker variant of expansion called (co)-systolic distance, but must do so *in both directions*. This is in strong contrast to typical constructions in the HDX literature which, due to the inherent asymmetry of simplicial complexes, typically have very poor boundary expansion (indeed this is also why we avoid simplicial complexes in this work). Such a guarantee was only recently achieved by Panteleev and Kalachev [PK21a] using refined products of chain complexes, and very recently simplified through a more geometric lens by Leverrier and Zémor [LZ22]. Since small-set

---

<sup>16</sup>A function is locally minimal if its weight cannot be decreased by adding the image of any standard basis vector  $\partial_2(e_v)$ . Any minimal function is also locally minimal (and the converse does not in general hold), so this is a strictly stronger notion of expansion than we study.

(co)-boundary expansion is a stronger notion than (co)-systolic distance (see Section 5), our analysis provides the strongest form of two-sided topological expansion to date. Further, this stronger form of two-sided expansion also gives some hope for a positive resolution of the famous qLTC conjecture. If, for instance, one can construct a 5-term chain complex satisfying similar bi-directional small set expansion guarantees, qLTC would follow by the arguments of [KKL14, EH17, LH22a].

## 3.2 Further Directions

**Improved Integrality Gaps:** We prove the existence of an explicit family of 3-XOR instances with a constant integrality gap of  $1$  v.s.  $1 - \mu$  for 3-XOR, which falls short of reaching the  $1$  v.s.  $\frac{1}{2} + \varepsilon$  gap exhibited by random instances [Gri01b, Sch08]. While standard reductions in the SoS hierarchy can improve our gap to arbitrarily close ( $1 - \varepsilon$  v.s.  $\frac{1}{2} + \varepsilon$ ), perfect completeness is lost in the process. The same issue was observed in [DFHT20]’s original explicit construction from the Ramanujan complex. They asked whether it is possible to bypass imperfect completeness by giving a direct construction with co-systolic distance at least  $\frac{1}{2} - \varepsilon$ . This remains a natural open question in our setting as well—can one directly construct a small-set boundary expander with co-systolic distance  $\frac{1}{2} - \varepsilon$ ? This would lead to a  $1$  v.s.  $\frac{1}{2} + \varepsilon$  gap for MAX- $k$ -XOR. Another natural question is whether such a bound can be transferred to 3-XOR without losing factors in the soundness. Our current reduction loses a factor in  $k$ , but we have made no attempt to optimize this step (since any constant gap is sufficient to amplify with PCP techniques if one is okay with imperfect completeness).

**Hardness Beyond XOR:** Many of the best integrality gaps known for combinatorial optimization problems (e.g. maximum independent set, chromatic number) are proved by reduction from  $k$ -CSPs [Tul09]. Unfortunately, such reductions are often *randomized*, so they do not imply explicit hard instances even when combined with our XOR construction. This raises a natural question: can we build *explicit* reductions from  $k$ -CSPs to classical combinatorial problems such as maximum independent set? Combined with our construction, this could lead to new families of hard instances for many well-studied combinatorial optimization problems. On a related note, it is worth observing that these reductions usually rely on CSPs with better integrality gaps than  $k$ -XOR. For instance, it is not hard to see that while random instances of  $k$ -XOR only exhibit a  $1$  v.s.  $1/2 + \varepsilon$  integrality gap, more constrained  $k$ -CSPs (e.g. constraints of the form  $Ax = b$  for some matrix  $A \in \mathbb{F}_2^{d \times k}$ ) can lead to much larger integrality gaps up to  $1$  v.s.  $\frac{2k}{2^k} + \varepsilon$  [Tul09]. Can we use high dimensional expanders to recover explicit  $k$ -CSPs matching these bounds?

**Small-Set HDX and Hardness of Approximation:** Small set expansion plays a fundamental role in hardness of approximation, ranging from use as a computational hardness assumption itself [RS10], to its pivotal use in the proof of the 2-2 games conjecture [KMS17, DKK<sup>+</sup>18b, DKK<sup>+</sup>18a, BKS18, KMMS18, KMS18] and recent converse use for algorithms for unique games [BBK<sup>+</sup>21, BHKL22]. This work gives the first application of *high dimensional* small-set expansion to hardness of approximation, raising the natural question: does this high dimensional variant have a broader role to play in the field as well?

## 4 Preliminaries I: SS-HDX to Hardness

We now cover the preliminary definitions required to understand our general translation of expanding chain complexes into hard instance of 3-XOR, including basics on Sum-of-Squares, chain complexes, and traditional notions of high dimensional expansion. Background required for the HDX construction itself (e.g. on left-right Cayley complexes, robust tensor codes, etc.) is postponed to Section 7.

## 4.1 Sum of Squares and Refutations

The Sum-of-Squares Semidefinite Programming Hierarchy is a powerful method for approximately solving constrained polynomial optimization problems, and is in particular the strongest known algorithmic framework for approximating CSPs. In brief, the SoS hierarchy presents a series of successively stronger SDP relaxations of a problem, where the ‘round- $t$ ’ relaxation optimizes over  $t$ -local views and runs in time  $n^{O(t)}$ . We refer the reader to [BS14, FKP19] for general information on the SoS hierarchy.

In this work, we focus in particular on the SoS relaxations of MAX- $k$ -XOR, the family of CSPs on  $n$  variables  $\{x_1, \dots, x_n\}$  and  $m$  constraints  $\{C_i\}_{i \in [m]}$  of the form:

$$x_{i_1} \oplus \dots \oplus x_{i_j} = z_i,$$

where  $z_i \in \{0, 1\}$ ,  $\{i_1, \dots, i_j\} \subset [n]$ , and  $j = j(i) \leq k$ . Let  $T_i \subset [n]$  denote the set of variables appearing in the  $i$ th constraint. Then the round- $t$  SoS SDP relaxation for MAX- $k$ -XOR can be written as:

---

**Algorithm 1:** Round- $t$  SoS Relaxation for MAX- $k$ -XOR

---

**Input:** variables  $\{v_S\}_{S \in \binom{[n]}{\leq t}}$

**Maximize:**  $\frac{1}{2} + \frac{1}{2m} \sum_{i=1}^m (-1)^{z_i} \langle v_{T_i}, v_{\emptyset} \rangle$

**Constraint to:**

1.  $\forall S_1 \oplus S_2 = S_3 \oplus S_4, |S_i| \leq t : \langle v_{S_1}, v_{S_2} \rangle = \langle v_{S_3}, v_{S_4} \rangle$
  2.  $\forall S, |S| \leq t : \|v_S\|_2 = 1$
- 

We refer to the maximum obtained by this SDP as the *value* of the round- $t$  relaxation, and say an infinite family of instances of MAX- $k$ -XOR is hard for (or cannot be refuted by)  $t$  rounds of Sum of Squares if there exists a constant  $\mu$  such that every instance is at most  $(1 - \mu)$ -satisfiable, but the round- $t$  SDP relaxation has value 1. In other words,  $t$ -rounds of the SoS hierarchy cannot distinguish between completely satisfiable and  $(1 - \mu)$ -satisfiable instances—this is often said to induce an *integrality gap* for the problem of size  $\frac{1}{1-\mu}$ .

Rather than working directly with the Sum-of-Squares SDP relaxations, we prove our hardness results through a fruitful connection with refutation complexity due to Schoenebeck [Sch08] and Tulsiani [Tul09]. More formally, following [DFHT20] we will use a proof system called  $\oplus$ -resolution where, given a system of linear equations  $\Lambda$  over  $\mathbb{F}_2$ , we may derive new equations by mod 2 summation:

$$\{\ell_1 = b_1\}, \{\ell_2 = b_2\} \implies \ell_1 \oplus \ell_2 = b_1 \oplus b_2.$$

A *refutation* in this system is a derivation that  $0 = 1$ , and in our setting corresponds to a proof that the XOR instance given by  $\Lambda$  is unsatisfiable. Schoenebeck [Sch08] and Tulsiani [Tul09] showed that any system without a short refutation has a matching SoS lower bound.

**Theorem 4.1** ([Sch08, Lemma 13] (as stated in [DFHT20])). *Let  $\Lambda$  be a system of linear equations in  $n$  variables over  $\mathbb{F}_2$ . If all refutations of  $\Lambda$  have an equation using at least  $2t$  variables, then the round- $t$  SoS Relaxation of  $\Lambda$  has value 1.*

## 4.2 Chain Complexes

While previous works constructing hard instances of CSPs rely on structure coming from graphs (e.g. [Gri01b, Sch08]) or hypergraphs [DFHT20], we take inspiration from recent work on c3-LTCs [DEL<sup>+</sup>21, LH22a] and qLDPC codes [PK21a, LZ22] and instead study a more general set of objects called *chain complexes*.

**Definition 4.2** (Chain Complex). Let  $X(0)$ ,  $X(1)$ , and  $X(2)$  be sets, and  $\partial_2 : \mathbb{F}_2^{X(2)} \rightarrow \mathbb{F}_2^{X(1)}$ ,  $\partial_1 : \mathbb{F}_2^{X(1)} \rightarrow \mathbb{F}_2^{X(0)}$  linear maps. The sequence

$$X : \mathbb{F}_2^{X(0)} \xleftarrow{\partial_1} \mathbb{F}_2^{X(1)} \xleftarrow{\partial_2} \mathbb{F}_2^{X(2)}$$

is called a (3-term) chain complex if  $\partial_1 \partial_2 = 0$ .

For the sake of intuition, let's take a moment to see why chain complexes are indeed a generalization of hypergraphs. Given an  $r$ -uniform hypergraph  $H \subseteq \binom{[n]}{r}$ , let  $X(i) \subseteq \binom{[n]}{i}$  denote any  $i$ -set contained in some  $r$ -set in  $H$ .  $H$  then induces an  $(r + 1)$ -term chain complex:<sup>17</sup>

$$X : \mathbb{F}_2^{X(0)} \xleftarrow{\partial_1} \mathbb{F}_2^{X(1)} \xleftarrow{\partial_2} \dots \xleftarrow{\partial_r} \mathbb{F}_2^{X(r)},$$

where  $\partial_i f(x)$  is given by summing  $f \pmod{2}$  over  $x$ 's 'boundary':

$$\forall f \in \mathbb{F}_2^{X(i)} : \partial_i f(x) = \sum_{y \in X(i): y \supset x} f(y). \quad (3)$$

For instance, when  $x$  is a vertex,  $\partial_2 f(x)$  averages over all edges containing  $x$ . As such,  $\partial$  is usually called the *boundary operator*, and it can be checked without too much difficulty that  $\partial_{i-1} \partial_i = 0$  (e.g. for  $r = 3$ , this follows by noting a vertex is incident to either 0 or 2 edges of any given triangle).

In fact, the boundary operators can actually always be seen to have a similar form to Equation (3), even on a generic chain complex. This follows from passing to the matrix representation as discussed in Section 1. Namely, we may view our 3-term chain complex as a pair of bipartite graphs  $B_0 = (X(0), X(1), E_1)$  and  $B_1 = (X(1), X(2), E_2)$ , whose bipartite adjacency matrices are given by the matrix representations of  $\partial_1$  and  $\partial_2$  respectively (in the standard basis). In this setting, it is easy to see that  $\partial_1$  and  $\partial_2$  are also given by mod 2 summation over neighbors on these underlying bipartite graphs:

$$\begin{aligned} \forall f \in \mathbb{F}_2^{X(1)} : \partial_1 f(x) &= \sum_{y \in X(1): (x,y) \in E_1} f(y) \pmod{2} \\ \forall f \in \mathbb{F}_2^{X(2)} : \partial_2 f(y) &= \sum_{z \in X(2): (y,z) \in E_2} f(z) \pmod{2}, \end{aligned}$$

where we have assumed for simplicity that  $\partial_1$  and  $\partial_2$  are *non-degenerate* in the sense that every row and column have at least one 1.<sup>18</sup> All complexes we study are non-degenerate, so we make this assumption throughout.

In matrix form, it is also easy to see that the transpose operators of  $\partial$ , called the *co-boundary operators* and denoted  $\delta_0 := \partial_1^T$  and  $\delta_1 := \partial_2^T$ , also form a chain complex in the opposite direction. As a result, we will usually write our chain complexes in the following form:

$$X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}.$$

We call elements of  $\mathbb{F}_2^{X(i)}$   *$i$ -chains*, and note  $\mathbb{F}_2^{X(i)}$  is often written as " $C_i$ " in the literature. We avoid this notation since it conflicts with classical notation for codes used later in the paper.

Finally, before moving on to expansion on chain complexes, we cover two further concepts that will control important parameters of our corresponding XOR instances: *maximum degree* and *explicitness*.

<sup>17</sup>Note  $X(0)$  is defined to be the empty set, and that our indexing is off by 1 from the usual notation in topology.

<sup>18</sup>In a graph, for instance, non-degeneracy corresponds to have no free-floating (degree 0) vertices.

**Definition 4.3** (Maximum Degree). *The maximum degree of a chain complex  $X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}$  is the maximum Hamming weight<sup>19</sup> across rows and columns of  $\partial_1$  and  $\partial_2$ .*

In the bipartite graph view, this is simply the maximum vertex degree across both graphs. We call an infinite family of chain complexes *bounded degree* if there exists some constant  $d \in \mathbb{N}$  such that all complexes in the family have maximum degree at most  $d$ .

Finally, in this work we will be interested in infinite families of chain complexes (and their associated XOR instances), so we need to define a notion of computational complexity over these objects. We will follow the standard notions used for expander families, and call a family of complexes *explicit* if its elements can be constructed in deterministic polynomial time (this is often called *mildly explicit*, but the difference is not particularly important in our setting).

**Definition 4.4** (Explicit Chain Complexes). *We call an infinite family of chain complexes  $\{X_i\}$  explicit if there exists a deterministic algorithm computing each  $X_i$  in time polynomial in  $|X_i(0) \cup X_i(1) \cup X_i(2)|$ .*

All complexes studied in this work will be bounded-degree, in which case this notion may equivalently be defined looking only at the size of  $X_i(0)$ . This corresponds correctly to the standard notion of complexity for the associated  $k$ -CSP family where  $|X_i(0)|$  gives the number of variables.

### 4.3 Homology and High Dimensional Expansion

High dimensional expansion is a generalization of expansion in graphs originally introduced by Linial and Meshulam [LM06] (and later independently by Gromov [Gro10]) to study the vanishing of homology in simplicial complexes. In this section we cover the basics of homology and introduce Linial and Meshulam's original notion of *(co)-boundary expansion*. These notions (or modifications thereof) will play an important role in our CSP construction.

Following standard notation, we call functions in the kernel of  $\partial_i$  *cycles*, and functions in the kernel of  $\delta_i$  *co-cycles*, denoted:

$$Z_i = \ker(\partial_i), \quad Z^i = \ker(\delta_i).$$

Since  $\partial^2 = \partial^2 = 0$ , notice that  $\text{im}(\partial_{i+1})$  are always cycles, and  $\text{im}(\delta_{i-1})$  are always co-cycles. We call functions in these classes *boundaries* and *co-boundaries* respectively, denoted:

$$B_i = \text{im}(\partial_{i+1}), \quad B^i = \text{im}(\delta_{i-1}).$$

The *homology* and *co-homology* of the chain complex correspond to (co)-cycles mod (co)-boundary:

$$H_i = Z_i/B_i, \quad H^i = Z^i/B^i,$$

where  $G/H$  denotes the quotient group. The notions of cycles and boundaries can be used to define a natural generalization of expander graphs to chain complexes called *(Co)-boundary expansion*.

**Definition 4.5** ((Co)-Boundary Expansion). *We call  $X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}$  a  $\rho$ -boundary expander if the weight of any element in  $\mathbb{F}_2^{X(1)} \setminus B_1$  is proportional to its distance from the boundary:*

$$\forall f \in \mathbb{F}_2^{X(1)} \setminus B_1 : \frac{|\partial_1 f|}{d(f, B_1)} \geq \rho,$$

where  $d(f, B_1) = \min_{b \in B_1} |f + b|$ . Similarly,  $X$  is an  $\rho$ -co-boundary expander if:

$$\forall f \in \mathbb{F}_2^{X(1)} \setminus B^1 : \frac{|\delta_1 f|}{d(f, B^1)} \geq \rho.$$

<sup>19</sup>The Hamming weight of binary vector  $v$ , denoted  $|v|$ , counts the number of entries with a 1.

Since this definition may seem un-motivated at first glance, let's again take a look at the case of a graph  $G = (V, E)$  which induces the (3-term) chain complex:

$$X : \mathbb{F}_2^\emptyset \begin{array}{c} \xrightarrow{\delta_0} \\ \xleftarrow{\partial_1} \end{array} \mathbb{F}_2^V \begin{array}{c} \xrightarrow{\delta_1} \\ \xleftarrow{\partial_2} \end{array} \mathbb{F}_2^E.$$

It is not hard to see that the co-boundary expansion of this chain is exactly Cheeger's constant:

$$h(G) := \min_{S \neq V, \emptyset} \left\{ \frac{E(S, V \setminus S)}{\min\{|S|, |V \setminus S|\}} \right\},$$

where  $E(S, V \setminus S)$  is the standard notation for the size of the edge boundary between  $S$  and the rest of the graph. This connection follows from noting that the only co-boundaries on this chain are  $V$  and  $\emptyset$ , and that  $|\delta_1 1_S|$  exactly counts the edge-boundary of  $S$ , so in particular we have:

$$\frac{|\delta_1 1_S|}{d(1_S, B^1)} = \frac{E(S, V \setminus S)}{\min\{|S|, |V \setminus S|\}}.$$

## 5 Small Set Boundary Expansion

(Co)-boundary expansion is a very strong property, and unconditional construction of bounded degree (co)-boundary expanders is still a major open question in topological high dimensional expansion. Furthermore, (co)-boundary expansion actually implies the vanishing of (co)-homology. This is an issue in and of itself in our setting, since as discussed in Section 2, our CSP construction rests crucially on the associated chain complex having non-trivial co-homology. With this in mind, we introduce a new notion of high dimensional expansion which requires boundary expansion to hold *only over small sets*.

**Definition 5.1** (Small-Set (Co)-Boundary Expansion). *We call  $X : \mathbb{F}_2^{X(0)} \begin{array}{c} \xrightarrow{\delta_0} \\ \xleftarrow{\partial_1} \end{array} \mathbb{F}_2^{X(1)} \begin{array}{c} \xrightarrow{\delta_1} \\ \xleftarrow{\partial_2} \end{array} \mathbb{F}_2^{X(2)}$  a  $(\rho_1, \rho_2)$ -small-set boundary expander if the weight of small chains in  $\mathbb{F}_2^{X(1)}$  is proportional to their distance from the boundary:*

$$\forall f \in \mathbb{F}_2^{X(1)} \setminus B_1, |f| \leq \rho_1 |X(1)| : \frac{|\partial_1 f|}{d(f, B_1)} \geq \rho_2.$$

Similarly,  $X$  is a  $(\rho_1, \rho_2)$ -small-set co-boundary expander if:

$$\forall f \in \mathbb{F}_2^{X(1)} \setminus B^1, |f| \leq \rho_1 |X(1)| : \frac{|\delta_1 f|}{d(f, B^1)} \geq \rho_2$$

We call  $X$  a  $(\rho_1, \rho_2)$ -small-set HDX if it is both a  $(\rho_1, \rho_2)$ -small-set boundary and  $(\rho_1, \rho_2)$ -small-set co-boundary expander.

Just like standard co-boundary expansion is a higher-order analog of Cheeger's constant (edge-expansion) in graphs, small-set co-boundary expansion is the natural analog of small-set expansion on graphs. As discussed in Section 3.1, various strengthened notions of (unidirectional) high-dimensional small-set expansion have been considered on both simplicial [KKL14, EK16, KM18, KM21] and chain complexes [LH22a, LH22b], but this basic generalization seems to be missing from the literature. In this work we show how small-set HDX can be transformed into explicit hard CSP instances for linear levels of Sum-of-Squares. Given the general prominence of small-set expansion throughout hardness of approximation (see e.g. [RS10, KMS18]), we expect SS-HDX may have many further applications in the area.

Before moving on, it will be useful to observe two important implications of a complex satisfying small-set (co)-boundary expansion. First, while the notion does not require the vanishing of (co)-homology like standard boundary expansion, it does still imply a strong restriction on the structure of elements in  $Z_1 \setminus B_1$ : they must be large.

**Lemma 5.2** (Small-Set (Co)-Boundary Expansion  $\rightarrow$  (Co)-Systolic Distance). *If  $X$  is a  $(\rho_1, \rho_2)$ -small-set boundary expander, then all chains  $f \in Z_1 \setminus B_1$  are large:*

$$\min_{f \in Z_1 \setminus B_1} \{|f|\} > \rho_1 |X(1)|. \quad (4)$$

*Similarly, if  $X$  is a  $(\rho_1, \rho_2)$ -small-set co-boundary expander, then all chains  $f \in Z^1 \setminus B^1$  are large:*

$$\min_{f \in Z^1 \setminus B^1} \{|f|\} > \rho_1 |X(1)|. \quad (5)$$

*Proof.* We prove the first statement only, the second follows similarly. Assume  $h_1 \in Z_1 \setminus B_1$  satisfies  $|h_1| \leq \rho_1 |X(1)|$ . Since  $h_1$  is a cycle, we have  $\partial_1 h_1 = 0$ , but then by small-set boundary expansion we have  $d(h_1, B_1) = 0$ , so  $h_1 \in B_1$  giving the desired contradiction.  $\square$

We say complexes satisfying Equation (4) have *systolic distance*  $\rho_1$ , and complexes satisfying Equation (5) have *co-systolic distance*  $\rho_1$ . As discussed in Section 2, these properties were recently crucial to the construction of good qLDPC codes [PK21a], and were also used by [DFHT20] to prove the soundness of their 3-XOR construction. Indeed it is worth noting that bounded co-systolic distance is actually enough for soundness in our construction as well, we only truly need the full power of small-set boundary expansion in one direction.

Second, we will crucially rely on a standard connection between boundary expansion and a concept known as an *isoperimetric inequality*, which relates the size of an object to the size of its boundary.<sup>20</sup> In particular, it is well known that boundary expansion is actually equivalent to an isoperimetric inequality for *minimal* chains (see e.g. [KKL16]).

**Definition 5.3** (Minimal Chains). *A function  $h \in \mathbb{F}_2^{X(1)}$  is called minimal if  $\forall b \in B_1, |h + b| \leq |h|$ .*

A similar equivalence holds for small-set boundary expansion as well, and will be crucial for the completeness of our CSP instances:  $X$  is a small-set boundary expander if and only if small, minimal chains in  $X$  satisfy an isoperimetric inequality.

**Lemma 5.4** (Small-Set (Co)-Boundary  $\leftrightarrow$  (Co)-Isoperimetric Inequality). *Let  $X$  be a  $(\rho_1, \rho_2)$ -small-set boundary expander. Then for any  $h \in \mathbb{F}_2^{X(1)}$  satisfying:*

1.  *$h$  is small:  $|h| \leq \rho_1 |X(1)|$*
2.  *$h$  is minimal:  $\forall b \in B_1 : |h + b| \geq |h|$*

*the boundary  $\partial_1 h$  must be large relative to  $h$ :*

$$|\partial_1 h| \geq \rho_2 |h|. \quad (6)$$

*Conversely if Equation (6) holds for any small minimal chain, then  $X$  is a  $(\rho_1, \rho_2)$ -small-set boundary expander.*

*Proof.* We start with the forward direction. Since  $|h| \leq \rho_1 |X(i)|$  and  $h$  is minimal, by small-set boundary expansion we have that:

$$|\partial_1 h| \geq \rho_2 d(h, B) = \rho_2 \min_{b \in B} \{|h + b|\} = \rho_2 |h|.$$

<sup>20</sup>For example the isoperimetric inequality on  $\mathbb{R}^2$  says the length (boundary) of any closed curve is at least  $2\sqrt{\pi}$  times the square root of its area.

The converse implication is similar. Let  $h \in \mathbb{F}_2^{X(1)}$  be a small chain satisfying  $|h| \leq \rho_1 |X(1)|$ , and let  $b \in B_1$  be the boundary minimizing  $|h + b|$ . Then by isoperimetry of  $h + b$ , we have:

$$|\partial_1 h| = |\partial_1(h + b)| \geq \rho_2 |h + b| = \rho_2 d(h, B_1)$$

as desired.  $\square$

We note the same result holds for co-boundary expansion by the same proof. Isoperimetry (combined with good systolic distance) will be crucial for showing completeness of our XOR instances, replacing the use of Gromov's filling inequality in [DFHT20].

## 6 From Expansion to Hardness

We now show how to translate any family of expanding, bounded-degree 3-term chain complexes with non-trivial cohomology into hard instances of 3-XOR for  $\Omega(n)$ -levels of Sum-of-Squares.

**Theorem 6.1.** *Let  $\{X_i\}$  be an explicit family of chain complexes of maximum degree  $k \in \mathbb{N}$  and  $\mu, \rho_1, \rho_2 \in (0, 1)$  constants such that:*

1.  $H^1$  is non-trivial,
2.  $X$  has  $\mu$ -co-systolic distance,
3.  $X$  is a  $(\rho_1, \rho_2)$ -small-set boundary expander.

*Then there exist constants  $\mu_1, \mu_2 \in (0, 1)$  depending only on  $k, \mu, \rho_1,$  and  $\rho_2$  and an explicit family of MAX-3-XOR instances  $\{\mathcal{I}_i\}$  on  $n_i$  variables such that:*

1. Every instance is at most  $(1 - \mu_1)$ -satisfiable,
2. No instance can be refuted by  $\mu_2 n_i$  levels of the SoS hierarchy.

*Moreover if the complex has degree lower bounded by 3,  $\{\mathcal{I}_i\}$  are instances of 3-XOR.*

Theorem 6.1 is actually proved mainly by associating an instance of MAX- $k$ -XOR to every complex  $X_i$  in the family. Moving to 3-XOR can then be done through standard NP-reduction arguments within the SoS hierarchy.<sup>21</sup> Thus the main challenge is to build hard instances of MAX- $k$ -XOR from our complexes. We'll start by overviewing our construction, which is a generalization of [DFHT20]'s 3-XOR construction from simplicial complexes to generic chain complexes.

**Construction:** It will be convenient to phrase our construction in the bipartite graph formulation discussed in Section 4. Recall that any chain complex  $X : \mathbb{F}_2^{X(0)} \xrightleftharpoons[\partial_1]{\delta_0} \mathbb{F}_2^{X(1)} \xrightleftharpoons[\partial_2]{\delta_1} \mathbb{F}_2^{X(2)}$  may be written as a pair of bipartite graphs  $B_1 = (X(0), X(1), E_1)$  and  $B_2 = (X(1), X(2), E_2)$  where  $E_1$  and  $E_2$  are uniquely determined by the matrix representations of the boundary operators. Assuming our complex has non-trivial co-homology, let  $\beta \in Z^1 \setminus B^1$ .<sup>22</sup> Our associated CSP  $\mathcal{I}_{X, \beta}$  is given by adding for every  $y \in X(1)$  the constraint:

$$C_y := \left\{ \sum_{x \in X(0): (x, y) \in E_1} x \pmod{2} = \beta(x) \right\}.$$

<sup>21</sup>Though one must be careful that the number of variables does not blow up in the reduction, as we discuss later in the section.

<sup>22</sup>Note that  $\beta$  can be found in polynomial time by standard linear algebraic techniques.

Since the choice of  $\beta \in Z^1 \setminus B^1$  will not matter, in what follows we will drop it from the notation and just write  $\mathcal{I}_X$ . We make two observations about  $\mathcal{I}_X$  before moving on. First, let's confirm  $\mathcal{I}_X$  is indeed an instance of MAX- $k$ -XOR.

**Observation 6.2.** *If  $X$  has maximum degree  $k$ , then  $\mathcal{I}_X$  is an instance of MAX- $k$ -XOR.*

*Proof.* This follows immediately from the chain complex having maximum degree  $k$ , as every  $y \in X(1)$  then has at most  $k$  neighbors in  $X(0)$  (i.e. that there are at most  $k$  elements  $x$  such that  $(x, y) \in E_1$ ).  $\square$

Second, we observe that our instances have at most a linear number of constraints.

**Observation 6.3.** *If  $X$  has maximum degree  $k$ , then  $\mathcal{I}_X$  has at most  $k|X(0)|$  constraints.*

*Proof.* Since our complex is non-degenerate and degree at most  $k$ , we have that  $|X(1)| \leq k|X(0)|$ .  $\mathcal{I}_X$  has  $|X(1)|$  constraints by construction.  $\square$

As a result, any explicit infinite family of bounded degree chain complexes with non-trivial cohomology induces an explicit infinite family of MAX- $k$ -XOR instances with linearly many constraints for some constant  $k \in \mathbb{N}$ . The main work in proving Theorem 6.1 therefore boils down to proving that the instances  $\mathcal{I}_X$  are *sound* (at most  $(1 - \mu)$ -satisfiable), and *complete* (look satisfiable to SoS).

**Theorem 6.4.** *Let  $X$  be a chain complex of maximum degree  $k$  and  $\mu, \rho_1, \rho_2 \in (0, 1)$  constants such that:*

1.  $H^1$  is non-trivial,
2.  $X$  has  $\mu$ -co-systolic distance,
3.  $X$  is a  $(\rho_1, \rho_2)$ -small-set boundary expander.

*Then  $\mathcal{I}_X$  is an instance of MAX- $k$ -CSP on  $|X(0)|$  variables satisfying:*

1. *Soundness:  $\mathcal{I}_X$  is at most  $(1 - \mu)$ -satisfiable,*
2. *Completeness:  $\mathcal{I}_X$  cannot be refuted by  $(\frac{\rho_1 \rho_2}{4k} |X(0)|)$ -levels of the SoS hierarchy.*

We'll break the proof of Theorem 6.4 into two parts, corresponding to soundness and completeness.

**Soundness:** The soundness of our construction can be proved with no further background, and is a direct generalization of arguments in [DFHT20] from simplicial complexes to general chain complexes.

*Proof of Soundness (Theorem 6.4).* Recall that our constraints are defined by some function  $\beta \in Z^1 \setminus B^1$ . Let  $f \in \mathbb{F}_2^{X(0)}$  be a potential assignment to variables in our instance. For any constraint  $y \in X(1)$ , we can check if  $f$  satisfies  $y$  by evaluating  $(\beta + \delta_0 f)(y)$ :

$$(\beta + \delta_0 f)(y) = \beta(y) + \sum_{(x,y) \in E_1} f(x).$$

In other words, the Hamming weight  $|\beta + \delta_0 f|$  exactly corresponds to the number of violated constraints in our instance. The key is now to observe that since  $\beta \in Z^1 \setminus B^1$ ,  $\beta + \delta_0 f$  also lies in  $Z^1 \setminus B^1$ . Since  $X$  has  $\mu$ -co-systolic distance, we have  $|\beta + \delta_0 f| \geq \mu|X(1)|$ , so any assignment to variables must violate at least a  $\mu$  fraction of constraints as desired.  $\square$

**Completeness:** Proving the completeness of Theorem 6.4 requires a bit more setup. As discussed in Section 4, we appeal to the general paradigm of Grigoriev [Gri01b], Schoenebeck [Sch08], and Tulsiani [Tul09] relating refutation width with Sum-of-Squares completeness. Our lower bound on the refutation width of  $\mathcal{I}_X$  can be viewed in some sense as a mix of the classical strategy of Ben-Sasson and Wigderson [BSW99] (who used traditional boundary expansion on graphs to show lower bounds against refuting Tseitin formulas) and the recent argument of [DFHT20] using Gromov’s filling inequality on the Ramanujan complex. We mostly follow the exposition given in the latter.

We will consider refutations in the  $\oplus$ -resolution proof system, in which two linear equations  $\ell_1 = b_1$  and  $\ell_2 = b_2$  can be added to derive  $\ell_1 \oplus \ell_2 = b_1 \oplus b_2$ . By Theorem 4.1, it is enough to prove that any refutation of the linear equations corresponding to  $\mathcal{I}_X$  has width at least  $\frac{\rho_1 \rho_2}{2k} |X(0)|$ , where width measures the largest number of variables appearing in any equation in the refutation. A refutation in the  $\oplus$ -resolution proof system can be modeled as a DAG where leaves correspond to linear equations (our XOR constraints), internal nodes have two incoming edges and correspond to the XOR of their parents, and the root derives the contradiction  $0 = 1$ .

To track the number of variables at each step, we follow the strategy of [DFHT20] and associate to each node  $v$  of the DAG a function  $h_v \in \mathbb{F}_2^{X(1)}$  and value  $b_v \in \{0, 1\}$  as follows. Since each leaf in the refutation corresponds to one of our XOR constraints, assign the leaf corresponding to  $s \in X(1)$  the indicator  $1_s \in \mathbb{F}_2^{X(1)}$  and value  $\beta(s) \in \mathbb{F}_2$  (where we recall  $\beta \in Z^1 \setminus B^1$  was the chain used to define our constraint values). The function and value assigned to each internal node  $v$  with parents  $v_1, v_2$  is then defined recursively to be the (mod 2) sum of its parents:

$$h_v = h_{v_1} \oplus h_{v_2}, \text{ and } b_v = \beta_{v_1} \oplus \beta_{v_2}.$$

Notice that by construction,  $\partial_1 h_v$  exactly corresponds to the variables appearing in the linear equation at node  $v$ . This means we can bound the width of the refutation by identifying some node  $v$  in the refutation whose associated function  $h_v$  has large boundary.

To this end, following [DFHT20]’s high dimensional variant of [BSW99]’s original technique we define the following potential function across nodes in our refutation:

$$\kappa(v) := \min_{b \in B_1} |h_v + b|.$$

Our goal will be to find a node in the refutation whose potential is large, but still small enough that we can apply small-set boundary expansion. Namely, if we can find  $v$  such that  $\frac{\rho_1}{2} |X(1)| \leq \kappa(v) \leq \rho_1 |X(1)|$ , then by our isoperimetric inequality for small sets (Lemma 5.4) we have:

$$|\partial h_v| = |\partial(h_v + b)| \geq \rho_2 |h_v + b| \geq \frac{\rho_1 \rho_2}{2} |X(1)| \geq \frac{\rho_1 \rho_2}{2k} |X(0)|$$

which would give the desired bound on refutation width. With this in mind, we can finally prove completeness.

*Proof of completeness (Theorem 6.4).* As discussed above, it is sufficient to prove that any refutation has width at least  $\frac{\rho_1 \rho_2}{2k} |X(0)|$ , and that this can be done by finding a node  $v$  with potential  $\frac{\rho_1}{2} |X(1)| \leq \kappa(v) \leq \rho_1 |X(1)|$ . The proof follows the classical strategy of [BSW99]. Namely it is enough to show the following three properties:

1. The root node has large potential:  $\kappa(r) > \rho_1 |X(1)|$
2. The leaves have small potential:  $\kappa(s) \leq 1$
3. The potential function is sub-additive:  $\kappa(v) \leq \kappa(v_1) + \kappa(v_2)$ .

As long as these hold, getting from the leaf potential of (at most) 1 to the root potential of  $\kappa(r) > \rho_1|X(1)|$  requires passing through some internal node  $v$  with  $\frac{\rho_1}{2}|X(1)| \leq \kappa(v) \leq \rho_1|X(1)|$  as desired.

It is left to prove the three properties, which follow from similar analysis as in [DFHT20] for the Ramanujan complex. The second and third properties are essentially immediate. Leaves are given by the indicator function of elements  $s \in X(1)$ , which are at most distance one from  $\vec{0} \in B_1$  (the all 0s function). Sub-additivity follows from the triangle inequality. For a node  $v$  with parents  $v_1$  and  $v_2$ , let  $b_1$  and  $b_2$  be boundaries minimizing  $d(h_{v_1}, B_1)$  and  $d(h_{v_2}, B_1)$ , then we have:

$$\kappa(v_1) + \kappa(v_2) = |h_{v_1} + b_1| + |h_{v_2} + b_2| \geq |h_{v_1} + b_1 + h_{v_2} + b_2| = |h_v + b_1 + b_2| \geq \kappa(v).$$

For the first property, we argue the root node  $r$  must satisfy  $h_r \in Z_1 \setminus B_1$ . If this is the case we are done by the fact that our complex has good co-systolic distance by Lemma 5.2:

$$\kappa(h_r) = \min_{b \in B_1} |h_r + b| > \rho_1|X(1)|,$$

since any  $h_r + b \in Z_1 \setminus B_1$  as well. To see that  $h_r \in Z_1 \setminus B_1$ , first note that since the root node in our refutation corresponds to the equation  $0 = 1$ , we must have  $\partial_1 h_r = 0$  and therefore  $h_r \in Z_1$ . To complete the proof we therefore only need to show  $h_r \notin B_1$ , which follows from the fact that  $b_r = 1$  for the root node. Namely, notice that for any node  $v$  we have  $b_v = \langle \beta, h_v \rangle$  by construction (since we are just summing mod 2 over the constraints), and in particular that  $\langle \beta, h_r \rangle = 1$ . On the other hand, if  $h_r \in B_1$ , then by definition there exists  $f \in \mathbb{F}_2^{X(2)}$  such that  $h_r = \partial_2 f$  and since  $\delta_1 = \partial_2^T$  we have

$$\langle \beta, h_r \rangle = \langle \beta, \partial_2 f \rangle = \langle \delta_1 \beta, f \rangle = 0$$

since  $\beta \in Z^1$ . Thus  $h_r$  is in  $Z_1$  but not  $B_1$ , which completes the proof.  $\square$

We are now one step away from proving Theorem 6.1; we just need to show how to move from a hard instance of MAX- $k$ -XOR to a hard instance of 3-XOR. Such a reduction is fairly standard within the SoS literature, but we'll include the proof for completeness. To do so, we'll need to introduce a second way to characterize completeness of an instance for  $t$  rounds of SoS through an object called a *pseudo-expectation*. Given a set of variables  $\{x_i\}_{i \in [n]}$  and  $d \in \mathbb{N}$ , let  $\text{poly}_{\mathbb{R}}(\{x_i\}, d)$  denote the set of degree at most  $d$  polynomials in  $\mathbb{R}[x_1, \dots, x_n]$ . For our purposes, it is enough to think of a degree  $2t$  pseudo-expectation as an operator  $\tilde{\mathbb{E}} : \text{poly}_{\mathbb{R}}(\{x_i\}, 2t) \rightarrow \mathbb{R}$  that 'pretends' to be an expectation in the following four ways:

1. Scaling:  $\tilde{\mathbb{E}}[1] = 1$

2. Linearity:

$$\forall a, b \in \mathbb{R}, p(x), q(x) \in \text{poly}_{\mathbb{R}}(\{x_i\}, 2t) : \tilde{\mathbb{E}}[ap(x) + bq(x)] = a\tilde{\mathbb{E}}[p(x)] + b\tilde{\mathbb{E}}[q(x)]$$

3. Positivity of Squares:

$$\forall q(x) \in \text{poly}_{\mathbb{R}}(\{x_i\}, t) : \tilde{\mathbb{E}}[q(x)^2] \geq 0$$

4. Booleanity:

$$\forall j \in [n], p(x) \in \text{poly}_{\mathbb{R}}(\{x_i\}, 2t - 2) : \tilde{\mathbb{E}}[x_j^2 p(x)] = \tilde{\mathbb{E}}[p(x)].$$

With this in mind, let  $\mathcal{I}$  be an instance of XOR on  $n$  variables  $\{x_1, \dots, x_n\}$ . It will be convenient to express constraints in  $C_i \in \mathcal{I}$  multiplicatively as:

$$C_i := \{x_{i_1} \dots x_{i_j} = b_i\}$$

where  $b_i \in \{-1, 1\}$  and assignments now range over  $\{-1, 1\}^n$ . Let  $C_i(x)$  be shorthand for the lefthand product of variables in the constraint, and  $|C_i|$  denote the degree of  $C_i(x)$ . It turns out (see e.g. [FKP19]) that completeness of  $\mathcal{I}$  against  $t$  levels of Sum-of-Squares is equivalent to the existence of a degree  $2t$  pseudo-expectation which respects every constraint  $C_i \in \mathcal{I}$  in the following strong sense:

$$\forall p(x) \in \text{poly}_{\mathbb{R}}(\{x\}, 2t - |C_i|) : \tilde{\mathbb{E}}[C_i(x)p(x)] = b_i \tilde{\mathbb{E}}[p(x)]. \quad (7)$$

With this in mind, we can finally put everything together and prove Theorem 6.1.

*Proof of Theorem 6.1.* We'll start by constructing an explicit family of hard instances of MAX- $k$ -XOR, then reduce to 3-XOR through the above machinery. By Theorem 6.4, every complex  $X_i$  in our family corresponds to an instance  $\mathcal{I}_{X_i}$  of MAX- $k$ -XOR on  $n_i = |X_i(0)|$  vertices and  $m_i \leq k|X(0)|$  constraints that is at most  $(1 - \mu)$ -satisfiable but cannot be refuted by the  $\frac{\rho_1 \rho_2}{4k} n_i$ -level SoS relaxation. Furthermore each instance  $\mathcal{I}_{X_i}$  can be constructed in  $\text{poly}(n_i)$  time. This follows immediately from the fact that  $\{X_i\}$  itself is explicit (and bounded degree), and that finding some  $\beta \in \mathbb{Z}_1 \setminus B_1$  can be done in polynomial time by basic linear algebra over dimension  $O(n_i)$  vector spaces.

It is left to argue that we can use  $\mathcal{I}_{X_i}$  to construct a corresponding instance of 3-XOR that remains hard for Sum-of-Squares. We will use the following simple approach: given a clause with more than 3 variables, split it into two clauses of about half the size whose product is the original clause. More formally, given a constraint  $C_i := \{x_{i_1} \dots x_{i_j} = b_i\}$ , we apply the transformation:

$$C_i \rightarrow \left\{ C_i^{(0)} := \{x_{i_1} \dots x_{i_{\lfloor j/2 \rfloor}} y_i = b_i\}, C_i^{(1)} := \{x_{i_{\lfloor j/2 \rfloor + 1}} \dots x_{i_j} y_i = 1\} \right\} \quad (8)$$

where  $y_i$  is a newly introduced ‘dummy’ variable. Given a generic instance of MAX- $k$ -XOR  $\mathcal{I}_k$ , let  $\Phi(\mathcal{I}_k)$  denote the CSP resulting from applying the above transformation to every constraint with more than 3 variables. We will argue that  $\Phi(\mathcal{I}_k)$  has about half as many variables per clause as the original instance, but maintains soundness and completeness up to constant factors.

**Claim 6.5.** *Let  $\mathcal{I}_k$  be an instance of MAX- $k$ -XOR for  $k \geq 4$  on  $n$  variables and  $m$  constraints such that:*

1.  $\mathcal{I}_k$  is at most  $(1 - \mu)$ -satisfiable,
2.  $\mathcal{I}_k$  cannot be refuted by  $t$  rounds of Sum-of-Squares.

*Then  $\Phi(\mathcal{I}_k)$  is an instance of MAX- $j$ -XOR for  $j = \lceil k/2 \rceil + 1$  on at most  $n + m$  variables and  $2m$  constraints satisfying:*

1.  $\mathcal{I}_k$  is at most  $(1 - \mu/2)$ -satisfiable,
2.  $\mathcal{I}_k$  cannot be refuted by  $\frac{2t}{k}$  rounds of Sum-of-Squares.

Let's first show Claim 6.5 completes the proof of our main theorem. Starting from our MAX- $k$ -XOR instance  $\mathcal{I}_{X_i}$ , Claim 6.5 shows that  $\Phi^{\lceil \log(k) \rceil}(\mathcal{I}_{X_i})$  is an instance of MAX-3-XOR on  $O_k(n_i)$  variables that is at most  $(1 - \Omega_k(\mu))$ -satisfiable but cannot be refuted by  $\Omega_k(n_i)$  rounds of Sum-of-Squares. This follows from the fact that the original (and all transformed instances) have  $m \leq O_k(n_i)$  constraints.<sup>23</sup> Finally, if the original instance had no constraints with fewer than 3 variables (which occurs if the original complex has degree lower bounded by 3),  $\Phi^{\lceil \log(k) \rceil}(\mathcal{I}_{X_i})$  is an instance of 3-XOR. With this in mind, it is left to prove the claim.

<sup>23</sup>We note that it is possible to improve the dependence on  $k$  by slightly more involved analysis, but since  $k$  is just a constant we choose to work with iterated applications of the above for simplicity of exposition.

*Proof of Claim 6.5.* The fact that  $\Phi(\mathcal{I}_k)$  is an instance of MAX- $j$ -XOR for  $j = \lceil k/2 \rceil + 1$  on at most  $n + m$  variables and at most  $2m$  constraints is immediate from construction. The main interest lies in proving soundness and completeness of the instance.

**Soundness:** Soundness of  $\Phi(\mathcal{I}_k)$  follows from observing that since  $y_i^2 = 1$ ,  $C_i = C_i^{(0)} \cdot C_i^{(1)}$ . Namely by the soundness of the original instance, any assignment of variables to  $\Phi(\mathcal{I}_k)$  must fail at least a  $\mu$  fraction of original constraints  $C_i$  (since these have no dependence on the new dummy variables). If  $C_i = C_i^{(0)} \cdot C_i^{(1)}$  is violated it must be the case that either  $C_i^{(0)}$  or  $C_i^{(1)}$  is violated, so any assignment of variables to our transformed CSP  $\Phi(\mathcal{I}_k)$  must still violate at least a  $\mu/2$  fraction of its constraints.

**Completeness:** Given a degree  $2t$  pseudo-expectation  $\tilde{\mathbb{E}}$  satisfying the constraints of  $\mathcal{I}_k$  (in the sense of Equation (7)), we must construct a new pseudo-expectation  $\tilde{\mathbb{E}}_\Phi$  on the variables of  $\Phi(\mathcal{I}_k)$  satisfying the transformed constraints. Given a polynomial  $p(x, y) \in \mathbb{R}[\{x_i\}, \{y_j\}]$ , let  $p(x, 1) \in \mathbb{R}[\{x_i\}]$  denote the result of setting each  $y$  variable to 1. The idea is to observe that each dummy variable  $y_i$  in the new instance can really be thought of as a ‘stand-in’ for the product  $b_i x_{i_1} \dots x_{i_{\lceil j/2 \rceil}} = b_i C_i^{(0)}(x, 1)$  in the sense that replacing each  $y_i$  with  $b_i C_i^{(0)}(x, 1)$  simply returns the original instance. This suggests a natural strategy for defining our new pseudo-expectation  $\tilde{\mathbb{E}}_\Phi$ : just replace  $y_i$  with  $b_i C_i^{(0)}(x, 1)$ .<sup>24</sup>

Formally, this takes a bit of work. Let  $S \subseteq [m]$  denote the set of indices on which we transformed our original instance,  $\{y_j\}_{j \in S}$  denote the newly introduced variables, and  $T : \mathbb{R}[\{x_i\}_{i \in [n]}, \{y_j\}_{j \in S}] \rightarrow \mathbb{R}[x_1, \dots, x_n]$  denote the map which independently replaces each occurrence of  $y_j$  with  $b_j C_j^{(0)}(x, 1)$  (and leaves variables in  $\{x_i\}$  unchanged). It is an elementary exercise to show that  $T$  satisfies the following useful properties:

1.  $T$  is (additively) linear:

$$T(az(x, y)) = aT(z(x, y)) \quad \text{and} \quad T(z_1(x, y) + z_2(x, y)) = T(z_1(x, y)) + T(z_2(x, y))$$

2.  $T$  is (multiplicatively) linear:

$$T(z_1(x, y)z_2(x, y)) = T(z_1(x, y))T(z_2(x, y))$$

3.  $T$  does not substantially blow up degree:

$$\text{Deg}(T(z(x, y))) \leq \lceil k/2 \rceil \text{Deg}(z(x, y)).$$

With this in mind, define the value of our new pseudo-expectation on any degree at most  $\frac{2t}{\lceil k/2 \rceil}$  polynomial  $z(x, y) \in \mathbb{R}[\{x_i\}_{i \in [n]}, \{y_j\}_{j \in S}]$  as:

$$\tilde{\mathbb{E}}_\Phi[z(x, y)] := \tilde{\mathbb{E}}[T(z(x, y))]$$

which is well-defined by the third property. It is an easy exercise to check that  $\tilde{\mathbb{E}}_\Phi$  remains a pseudo-expectation, as the linearity of  $T$  ensures scaling, linearity, positivity of squares, and booleanity are all inherited from  $\tilde{\mathbb{E}}$ . Thus it is left to check that  $\tilde{\mathbb{E}}_\Phi$  satisfies every constraint  $C_i^{(j)} \in \Phi(\mathcal{I}_k)$  in the sense of Equation (7). To see this, first observe that

$$\begin{aligned} \tilde{\mathbb{E}}_\Phi[C_i^{(j)}(x, y)z(x, y)] &= \tilde{\mathbb{E}}[T(C_i^{(j)}(x, y)z(x, y))] \\ &= \tilde{\mathbb{E}}[T(C_i^{(j)}(x, y))T(z(x, y))]. \end{aligned}$$

<sup>24</sup>We thank Sam Hopkins for suggesting this general approach.

Taking a closer look at  $T(C_i^{(j)}(x, y))$ , we have by definition that:

$$T(C_i^{(j)}(x, y)) = \begin{cases} b_i C_i(x) & \text{if } j = 1 \\ b_i (C_i^{(0)}(x, 1))^2 & \text{if } j = 0. \end{cases}$$

Breaking into case analysis, we then have for  $j = 1$ :

$$\begin{aligned} \tilde{\mathbb{E}}[T(C_i^{(1)}(x, y))T(z(x, y))] &= \tilde{\mathbb{E}}[b_i C_i(x)T(z(x, y))] \\ &= \tilde{\mathbb{E}}[T(z(x, y))] \\ &= \tilde{\mathbb{E}}_{\Phi}[z(x, y)] \end{aligned}$$

and for  $j = 0$  that:

$$\begin{aligned} \tilde{\mathbb{E}}[T(C_i^{(0)}(x, y))T(z(x, y))] &= \tilde{\mathbb{E}}[b_i (C_i^{(0)}(x, 1))^2 T(z(x, y))] \\ &= b_i \tilde{\mathbb{E}}[T(z(x, y))] \\ &= b_i \tilde{\mathbb{E}}_{\Phi}[z(x, y)] \end{aligned}$$

which match the form of the constraints given in Equation (8) as desired. □

□

## 7 Preliminaries II: Constructing SS-HDX

We now cover the tools necessary for constructing our small-set HDX, including background on basic expander graphs, left-right Cayley complexes, error correcting codes, Tanner codes, and tensor codes. We closely follow the discussion in [LZ22] who largely cover the same background material.

### 7.1 Expander Graphs

The main building block of Leverrier and Zémor's qLDPC codes are a ubiquitous class of graphs in computer science called *spectral expanders*. Let  $\mathcal{G} = (V, E)$  be an undirected  $\Delta$ -regular (multi)-graph on  $n$  vertices, and define  $\lambda(\mathcal{G}) := \max\{|\lambda_2|, |\lambda_n|\}$  where  $\Delta = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$  are the eigenvalues of the adjacency matrix of  $\mathcal{G}$ . We say  $\mathcal{G}$  is a  $\lambda$ -spectral expander if  $\lambda(\mathcal{G}) \leq \lambda$ , and call it *Ramanujan* if  $\lambda(\mathcal{G}) \leq 2\sqrt{\Delta - 1}$ , which is the optimal expansion for infinite families of fixed degree [Alo86].

We will rely on spectral expanders for two main reasons. First, as we will discuss in the following section, infinite families of these objects are well-known not only to exist, but to be explicitly constructable (see e.g. [Mor94]). Second, spectral expansion provides a useful proxy for edge-expansion in the sense that for any  $S, T \subseteq V$ , there cannot be too many edges passing between  $S$  and  $T$ . This is classically known as the *expander-mixing lemma*, and likely first appeared in [AC88]:

**Lemma 7.1** (Expander mixing lemma). *Let  $\mathcal{G}$  be a  $\Delta$ -regular graph. Then for any subset  $S, T \subset V(\mathcal{G})$  we have*

$$|E(S, T)| \leq \frac{\Delta}{|V|} |S||T| + \lambda(\mathcal{G}) \sqrt{|S||T|}.$$

When  $|S|$  and  $|T|$  are small compared with  $|V|$ , we will think of  $\lambda(\mathcal{G}) \sqrt{|S||T|}$  as the main term and  $\frac{\Delta}{|V|} |S||T|$  as the error term (we note this is the opposite of how the lemma is often applied).

It will also be important for us that the expander mixing lemma holds for *double covers* of a spectral expanders with a small modification. The double cover  $\mathcal{G}' = (V', E')$  of a graph  $\mathcal{G} = (V, E)$  has vertex set  $V' = V_0 \cup V_1$ , for  $V_0 = V \times \{0\}$  and  $V_1 = V \times \{1\}$ , and edge set  $E' = \{(v, 0), (w, 1)\} : v, w \in V, \{v, w\} \in E\}$ . The expander mixing lemma applies for double covered graphs when  $S \subset V_0, T \subset V_1$ .

**Lemma 7.2** (Expander mixing lemma for double covered graph). *Let  $\mathcal{G}$  be a  $\Delta$ -regular graph and  $\mathcal{G}'$  be its double cover. Then for any subset  $S \subset V_0(\mathcal{G}')$ ,  $T \subset V_1(\mathcal{G}')$  we have*

$$|E(S, T)| \leq \frac{\Delta}{|V(\mathcal{G})|} |S||T| + \lambda(\mathcal{G}) \sqrt{|S||T|}.$$

This can be shown easily by projecting  $S$  and  $T$  back to the original graph.

## 7.2 Left-Right Cayley Complexes

While expansion is a useful property in its own right, our arguments require higher dimensional structure. The key lies in an object called the *left-right Cayley complex* introduced in [DEL<sup>+</sup>21] to build c3-LTCs. A left-right Cayley complex is determined by a group  $G$  and two sets of generators  $A = A^{-1}$  and  $B = B^{-1}$ . The complex consists of vertices,  $A$ -edges,  $B$ -edges, and squares as follows:

- The vertices are  $V^0 = G$ .
- The  $A$ -edges are  $E_A^0$  and the  $B$ -edges are  $E_B^0$  where

$$E_A^0 = \{\{g, ag\} : g \in G, a \in A\}, E_B^0 = \{\{g, gb\} : g \in G, b \in B\}.$$

- The squares are

$$F^0 = \{\{g, ag, gb, agb\} : g \in G, a \in A, b \in B\}.$$

The main criterion for choosing  $G$ ,  $A$ , and  $B$  is to ensure the Cayley graphs  $\text{Cay}(G, A)$  and  $\text{Cay}(G, B)$  are good expanders, and in particular are Ramanujan. Besides this, for simplicity we further assume two technical conditions as in [DEL<sup>+</sup>21]: that  $|A| = |B| = \Delta$ , and the so-called *total no-conjugacy* condition

$$\forall a \in A, b \in B, g \in G, ag \neq gb.$$

The total no-conjugacy condition ensures squares are non-degenerate (contain exactly 4 distinct vertices), and that each vertex is incident to exactly  $k^2$  squares [DEL<sup>+</sup>21, Claim 3.7]. Leveraging classical results of Morgenstern [Mor94] and Lubotzky, Samuels, and Vishne [LSV05], [DEL<sup>+</sup>21] show that explicit families of left-right Cayley complexes exist for infinitely many degrees.

**Theorem 7.3** ([DEL<sup>+</sup>21, Claim 6.7]). *There exists an infinite sequence of degrees  $\Delta = q + 1$  (where  $q$  is an odd prime power) such that for each fixed  $\Delta$  there exists an explicit infinite family of left-right Cayley complexes with  $G_i = \text{PSL}_2(q^i)$  and generator sets  $A_i$  and  $B_i$  such that  $|A_i| = |B_i| = \Delta$ ,  $\text{Cay}(G_i, A_i)$  and  $\text{Cay}(G_i, B_i)$  are Ramanujan, and  $A_i, B_i$  satisfy the total no-conjugacy condition.*

As in [LZ22], we will use the *double cover* of the left-right Cayley complex, defined as:

- The vertices are  $V = V_0 \cup V_1$  where  $V_0 = G \times \{0\}$  and  $V_1 = G \times \{1\}$ .
- The  $A$ -edges are  $E_A$  and the  $B$ -edges are  $E_B$  where

$$E_A = \{\{(g, 0), (ag, 1)\} : g \in G, a \in A\}, E_B = \{\{(g, 0), (gb, 1)\} : g \in G, b \in B\}.$$

- The squares are

$$F = \{\{(g, 0), (ag, 1), (gb, 1), (agb, 0)\} : g \in G, a \in A, b \in B\}.$$

Note that every square in the original left-right Cayley complex corresponds to two squares in the double cover, and therefore that the double cover has a total of  $\frac{\Delta^2|G|}{2}$  squares. Since we will only use the double cover in our arguments, from now on the term “square” will *always refer to these double-covered squares*, not the squares in the original Cayley complex.

Following [LZ22]’s notation, we will mainly think of the double-covered complex as represented by the following graphs. First, we’ll define a graph that captures the vertices and edge-structure of the Cayley complex:  $\mathcal{G}^\cup = (V, E_A \cup E_B)$ . Second, we’ll define graphs<sup>25</sup>  $\mathcal{G}_0^\square = (V_0, E_0^\square)$  and  $\mathcal{G}_1^\square = (V_1, E_1^\square)$  capturing squares in the double cover, where

$$E_i^\square = \{(g, i), (agb, i)\} : g \in G, a \in A, b \in B\}$$

for  $i \in \{0, 1\}$ . Notice that the edges in these graphs have a one-to-one correspondence with the double-covered squares, namely that  $E_i^\square \cong F$  for  $i = 0, 1$  through the following identifications:

$$\{(g, 0), (agb, 0)\} \leftrightarrow \{(g, 0), (ag, 1), (gb, 1), (agb, 0)\}$$

and

$$\{(g, 1), (agb, 1)\} \leftrightarrow \{(g, 1), (ag, 0), (gb, 0), (agb, 1)\}.$$

These identifications will be particularly important in the proof of small-set (co)-boundary expansion as we move between the squares of our complex and their associated graph representations.

Finally, it will be important to observe that these graphs inherit the spectral properties of  $\text{Cay}(G, A)$  and  $\text{Cay}(G, B)$ . Namely that when the latter are Ramanujan,  $\mathcal{G}^\cup, \mathcal{G}_0^\square, \mathcal{G}_1^\square$  are also very good expanders.

**Lemma 7.4** ([LZ22, Lemma 4]). *If  $\text{Cay}(G, A), \text{Cay}(G, B)$  are Ramanujan graphs, then  $\lambda(\mathcal{G}_0^\square) \leq 4\Delta$ ,  $\lambda(\mathcal{G}_1^\square) \leq 4\Delta$ , and  $\mathcal{G}^\cup$  is the double cover of a  $4\sqrt{\Delta}$ -spectral expander.*

We note this is not exactly the statement given in [LZ22], but the proof is the same.

### 7.3 Error Correcting Codes

A classical  $(n, k, d)$ -error correcting (erasure) code is a method for encoding a string of  $k$  classical bits into  $n > k$  classical bits such that one can recover the original string even when up to  $d - 1$  bits of the encoded string are erased. More formally, we will consider the standard setting of *linear codes*, where the encoded space is a linear subspace  $\mathcal{C} \subset \mathbb{F}_2^n$ . Here  $n$  is the *length* of the code,  $k := \dim(\mathcal{C})$  is its *dimension*, and the minimum weight of any element (also called codeword) of  $\mathcal{C}$ ,  $d := \min_{c \in \mathcal{C}} \{|c|\}$ , is called its *distance*.<sup>26</sup> One can check that in a linear code of distance  $d$ , it is indeed possible to uniquely correct up to  $d - 1$  errors. Finally, the ratio  $r := \frac{k}{n}$  is called the *rate* of the code, and measures the overhead from the original to encoded space. We will typically be interested in families of codes that have constant rate and linear distance.

One of the main reasons to use linear codes is that there are nice linear algebraic ways of describing the objects. In particular, the linear subspace (code)  $\mathcal{C}$  is typically described either by a *parity-check* matrix, or a *generator* matrix. In particular, one can always find a parity-check matrix  $M : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{n-k}$  whose kernel is the code in question ( $\mathcal{C} := \ker M \subset \mathbb{F}_2^n$ ), and likewise a generator matrix  $M' : \mathbb{F}_2^k \rightarrow \mathbb{F}_2^n$  whose image gives the code ( $\mathcal{C} := \text{im } M' \subset \mathbb{F}_2^n$ ). When clear from context, we sometime abuse notation and write  $\mathcal{C}$  to mean the parity check matrix of  $\mathcal{C}$ .

<sup>25</sup>We note these may technically be multi-graphs as in [LZ22], but this has no effect on our arguments.

<sup>26</sup>We note that this is similar to the distance operator  $d(\cdot, \cdot)$  used to define co-boundary expansion. Indeed the distance of a code  $\mathcal{C}$  is just  $d(\emptyset, \mathcal{C})$ . We will abuse notation slightly to match standard coding theory notation and write this as  $d(\mathcal{C})$  throughout.

## 7.4 Tanner Codes

The *Tanner construction* (or *tanner code*) [Tan81] is a classical strategy in coding theory to build a linear code out of a ‘large’ regular graph and a ‘small’ *local* code that sits on the neighborhood of each vertex. Crucially, when the underlying graph is an expander, it is often the case that the Tanner code inherits desirable properties from the small code.

More formally, let  $\mathcal{G} = (V, E)$  be a  $\Delta$ -regular graph and  $E(v)$  denote the set of edges incident to any  $v \in V$ . Assume an identification of  $\mathbb{F}_2^{E(v)}$  with  $\mathbb{F}_2^\Delta$  for each  $v \in V$ , which we call the *local view* of  $v$ . Given a local code  $C_0$  with length  $\Delta$ , the Tanner code  $T(\mathcal{G}, C_0) \subset \mathbb{F}_2^E$  is given by

$$\{c \in \mathbb{F}_2^E : \forall v \in V, c|_{E(v)} \in C_0\},$$

where  $c|_{E(v)} \in \mathbb{F}_2^\Delta$  is the vector formed by the values of  $c$  on the local view of  $v$ .

It will be convenient for us to view the Tanner construction through its parity check matrix, which will make up the co-boundary operators of our chain complex. If our local code  $C_0$  has parity check matrix  $M_0$  and rate  $r_0$ , the parity check matrix of the Tanner code  $T(\mathcal{G}, C)$  is given by the composition:

$$\mathbb{F}_2^E \rightarrow \mathbb{F}_2^{V \times \Delta} \rightarrow \mathbb{F}_2^{V \times (1-r_0)\Delta}$$

where the first map copies the value on the edge to each local view of the vertices, and the second map applies  $M_0$  to each local view independently for each vertex. We will sometimes refer to this parity check matrix as the *Tanner map*.

## 7.5 Robust Tensor Codes Against Puncture

The properties of our Tanner maps are highly dependent on the local code used to instantiate them. Following [LZ22], we use a special type of local code called a *tensor code*. We closely follow the discussion of these objects given in [LZ22].

Recall that the generators of our left-right Cayley complex  $A$  and  $B$  have size  $\Delta$ . We will consider codes on  $\mathbb{F}_2^{A \times B}$  with tensor product structures. Namely, given two linear codes  $C_A \subset \mathbb{F}_2^A, C_B \subset \mathbb{F}_2^B$ , we define the *tensor code*  $C_A \otimes C_B$  to be the set of  $\Delta \times \Delta$  matrices  $M$  where each column vector  $(M_{ab})_{a \in A}$  belongs to  $C_A$  and each row vector  $(M_{ab})_{b \in B}$  belongs to  $C_B$ . We define the *dual tensor code* to be the sum  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$ , where  $C_A \otimes \mathbb{F}_2^B$  are the  $\Delta \times \Delta$  matrices whose columns belong to  $C_A$ , and  $\mathbb{F}_2^A \otimes C_B$  are the  $\Delta \times \Delta$  matrices whose rows belong to  $C_B$ . The following claims about the dimension and distance of these codes are standard and easy to verify:

1.  $\dim(C_A \otimes C_B) = \dim(C_A) \dim(C_B)$
2.  $d(C_A \otimes C_B) = d(C_A)d(C_B)$
3.  $\dim(C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B) = \Delta \dim(C_A) + \Delta \dim(C_B) - \dim(C_A) \dim(C_B)$
4.  $d(C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B) = \min(d(C_A), d(C_B))$ .

To ensure our Tanner maps have the right properties, we will actually require our local tensor codes to have a stronger property called robustness. One can think of robustness as a generalization of distance of usual linear codes to the context of tensor codes, or as we will soon see, as a sort of robust testability property.

**Definition 7.5** (Robust [LZ22, Definition 5]). *Let  $C_A \subset \mathbb{F}_2^A, C_B \subset \mathbb{F}_2^B$  be codes of length  $\Delta$  of distance  $d_A$  and  $d_B$  respectively. We say the dual tensor code  $C = C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$  is  $w$ -robust if for every codeword  $c \in C$  with Hamming weight  $|c| < w$ , there exist  $A' \subset A, B' \subset B, |A'| \leq |c|/d_B, |B'| \leq |c|/d_A$ , such that  $c_{ab} = 0$  for any  $a \notin A'$  and  $b \notin B'$ .*

Leverrier and Zémor [LZ22] prove that robust tensor codes satisfy a useful small-set robust testability property.

**Lemma 7.6** ([LZ22], Proposition 6). *Let  $C_A \subset \mathbb{F}_2^A, C_B \subset \mathbb{F}_2^B$  be codes of length  $\Delta$  of distance  $d_A$  and  $d_B$  respectively. If the dual tensor code  $C = C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$  is  $w$ -robust with  $w \leq d_A d_B / 2$ , then any word  $x$  close to both the column and row code is also close to the tensor code. More explicitly, if  $d(x, C_A \otimes \mathbb{F}_2^B) + d(x, \mathbb{F}_2^A \otimes C_B) < w$  then:*

$$d(x, C_A \otimes C_B) \leq \frac{3}{2} (d(x, C_A \otimes \mathbb{F}_2^B) + d(x, \mathbb{F}_2^A \otimes C_B)).$$

In fact, [LZ22] need a slightly stronger condition than just robustness of the code: it needs to remain robust even after the removal of a small set of rows and columns. Conceptually, this is similar to the idea of smooth codes [DSW06] where the code maintains nice properties even after the removal of a small number of variables or checks. Given a code  $C_A \subset \mathbb{F}_2^A$  and  $A' \subset A$ , let  $C_{A'} \subset \mathbb{F}_2^{A'}$  denote the *puncture code* which is the restriction of all codewords in  $C_A$  to the coordinates in  $A'$  (more precisely,  $C_{A'} = \{(c_a)_{a \in A'} : (c_a)_{a \in A} \in C_A\}$ ).

**Definition 7.7** (Robust against puncture [LZ22, Definition 7]). *Given linear codes  $C_A \subset \mathbb{F}_2^A, C_B \subset \mathbb{F}_2^B$ , we say the dual tensor code  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$  is  $w$ -robust with  $p$ -resistance to puncture if for any  $w' \leq p$  and  $A' \subset A$  and  $B' \subset B$  such that  $|A'| = |B'| = \Delta - w'$ , the dual tensor code  $C_{A'} \otimes \mathbb{F}_2^{B'} + \mathbb{F}_2^{A'} \otimes C_{B'}$  is  $w$ -robust.*

Extending prior work of [PK21a], [LZ22] show random tensor codes are robust against puncture.

**Theorem 7.8** ([LZ22, Theorem 8]). *Let  $0 < r_A < 1$  and  $0 < r_B < 1$ . Let  $0 < \varepsilon < 1/2$  and  $1/2 + \varepsilon < \gamma < 1$ . Let  $C_A$  be a random code obtained from a random uniform  $r_A \Delta \times \Delta$  generator matrix, and let  $C_B$  be a random code obtained from a random uniform  $(1 - r_B) \Delta \times \Delta$  parity-check matrix. With probability tending to 1 when  $\Delta$  goes to infinity, the dual tensor code*

$$C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$$

*is  $\Delta^{3/2-\varepsilon}$ -robust with  $\Delta^\gamma$ -resistance to puncturing.*

Because the dual of a random code is again a random code, this implies both  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$  and  $C_A^\perp \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B^\perp$  are robust against puncture with high probability.

**Corollary 7.9** ([LZ22, Theorem 17]). *Fix  $r \in (0, 1/2)$ ,  $\varepsilon \in (0, 1/2)$ ,  $\gamma \in (1/2 + \varepsilon, 1)$  and  $\delta > 0$  satisfying  $-\delta \log \delta - (1 - \delta) \log(1 - \delta) < r$ . When  $k$  is large enough, there exist codes  $C_A$  and  $C_B$  of length  $\Delta$  such that*

1.  $\dim C_A = \lfloor r\Delta \rfloor$  and  $\dim C_B = \Delta - \dim C_A$
2. The distances of  $C_A, C_B, C_A^\perp, C_B^\perp$  are all at least  $\delta\Delta$
3. Both dual tensor codes  $C_0^\perp = (C_A \otimes C_B)^\perp$  and  $C_1^\perp = (C_A^\perp \otimes C_B^\perp)^\perp$  are  $\Delta^{3/2-\varepsilon}$ -robust with  $\Delta^\gamma$ -resistance to puncturing
4.  $C_A, C_B, C_A^\perp,$  and  $C_B^\perp$  have generator matrices where every row and column have at least two ones.

We note that this is not exactly the statement of [LZ22, Theorem 17], who prove the first three conditions occur with probability going to 1 as  $\Delta$  becomes large when  $C_A$  and  $C_B$  are generated as in Theorem 7.8. The fourth item is not included in [LZ22], but also occurs under this distribution with high probability by fairly standard arguments. We give the proof in the appendix for completeness.

## 8 Constructing Small-Set HDX

We are finally ready to construct a family of 3-term chain complexes with small-set boundary and co-boundary expansion.

**Theorem 8.1.** *There exists an explicit infinite family of chain complexes  $\{X_i\}$  and constants  $d \in \mathbb{N}$  and  $\rho_1, \rho_2 \in (0, 1)$  such that each  $X_i$  satisfies:*

1.  $X_i$  has maximum degree  $d$  and minimum degree at least 3
2.  $X_i$  has non-trivial co-homology  $H^1$
3.  $X_i$  is a  $(\rho_1, \rho_2)$ -small-set HDX.

Combined with Theorem 6.1 which transforms SS-HDX into hard instances of 3-XOR, this completes the proof of our main theorem.

*Proof of Theorem 1.1.* Theorem 6.1 gives the desired explicit family of 3-XOR instances as long as it is provided an explicit family of chain complexes with bounded maximum degree, minimum degree at least 3, non-trivial co-homology, and which are  $(\rho_1, \rho_2)$ -small-set boundary expanders with  $\mu$ -co-systolic distance for some set of constants  $\mu, \rho_1, \rho_2 \in (0, 1)$ . Since any  $(\rho_1, \rho_2)$ -small-set co-boundary expander has  $\rho_1$ -co-systolic distance (Lemma 5.2), Theorem 8.1 provides an explicit family of chain complexes matching these conditions with  $\mu = \rho_1$ .  $\square$

As discussed, Theorem 8.1 is proved via Leverrier and Zémor’s [LZ22] recent construction of good qLPDC codes. They show the associated 3-term chain complex has linear systolic and co-systolic distance. Our contribution is to observe that the same construction actually satisfies the stronger small-set boundary and co-boundary expansion conditions. We note that while we only show this property for Leverrier and Zémor’s [LZ22] simplified construction, similar arguments likely hold for Panteleev and Kalachev’s [PK21a] original good qLDPC codes as well.

**Construction:** We first describe Leverrier and Zémor’s construction, which is based upon Tanner maps (parity-check matrices of Tanner codes). To start, we’ll first need to describe the underlying graphs and local codes of these maps. Recall the explicit family of left-right Cayley complexes promised by Theorem 7.3 and for any fixed complex  $X_i$  in the family let the group  $G = G_i$  and generator sets  $A = A_i, B = B_i$  be as in the theorem. The graphs underlying our Tanner maps will be the ‘square graphs’  $\mathcal{G}_0^\square$  and  $\mathcal{G}_1^\square$ , which we recall have

- Vertices  $V_i = G \times \{i\}$ ,
- Edges  $E_i^\square = \{(g, i), (agb, i)\} : g \in G, a \in A, b \in B\}$

for  $i \in \{0, 1\}$  respectively. It bears repeating that edges in these graphs are in one-to-one correspondence with squares of the double covered Cayley complex via the following identifications:

$$\{(g, i), (agb, i)\} \leftrightarrow \{(g, i), (ag, 1 - i), (gb, 1 - i), (agb, i)\}.$$

We will frequently refer to edges in  $\mathcal{G}_i^\square$  as squares due to this connection.

Since the square graphs  $\mathcal{G}_i^\square$  are  $\Delta^2$ -regular, we can define a Tanner map by combining them with any length  $\Delta^2$  local code. This role will be played by the robust dual tensor codes promised by Corollary 7.9.

Namely, letting  $C_A : \mathbb{F}_2^\Delta \rightarrow \mathbb{F}_2^{r\Delta}$  and  $C_B : \mathbb{F}_2^\Delta \rightarrow \mathbb{F}_2^{(1-r)\Delta}$  be as in Corollary 7.9 for some choice of  $r, \varepsilon$ , and  $\gamma$ , our local codes will be  $C_0^\perp$  and  $C_1^\perp$  where  $C_0 = C_A \otimes C_B$  and  $C_1 = C_A^\perp \otimes C_B^\perp$ .<sup>27</sup>

Combining these graphs and local codes gives the Tanner maps  $\mathcal{C}_0 = T(\mathcal{G}_0^\square, C_0^\perp) : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$  and  $\mathcal{C}_1 = T(\mathcal{G}_1^\square, C_1^\perp) : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ , where  $n = |F| = \Delta^2|G|/2$  is the number of squares, and  $m = r(1-r)\Delta^2|V_0| = r(1-r)\Delta^2|G|$  comes from the fact that the both dual tensor codes have dimension  $(1-r(1-r))\Delta^2$ . Associating the edges of  $\mathcal{G}_i^\square$  with squares in the discussed manner, one can check that  $\mathcal{C}_1\mathcal{C}_0^T = 0$  (see [LZ22, Section 4.1]) and therefore that these maps define a chain complex:

$$X : \mathbb{F}_2^m \xrightarrow{\delta_0 := \mathcal{C}_0^T} \mathbb{F}_2^n \xrightarrow{\delta_1 := \mathcal{C}_1} \mathbb{F}_2^m. \quad (9)$$

Moreover, this process gives an explicit family of chain complexes  $\{X_i\}$  by choosing  $G_i, A_i$ , and  $B_i$  as in the explicit family of left-right Cayley complexes promised by Theorem 7.3, and computing  $C_A, C_B$  with the desired properties by brute force search over all pairs of length  $\Delta$  codes  $C_A, C_B$  of dimensions  $r\Delta$  and  $(1-r)\Delta$  respectively.<sup>28</sup>

This completes the construction. We now move to showing that  $X$  has the three desired properties: bounded-degree, non-trivial co-homology, and small-set (co)-boundary expansion.

**$X$  has (upper) bounded-degree:** By definition  $X$  is bounded-degree if and only if the parity-check matrices of our two Tanner codes have a bounded number of ones in every row and column. By the nature of the Tanner code construction the support of any row or column is at most twice the degree of the underlying graph. Since our graphs are of degree  $\Delta^2$  (a constant with respect to the family), the resulting complex is bounded-degree as desired.

**$X$  has (lower) bounded-degree:** Recall we are promised that  $C_A, C_B, C_A^\perp$ , and  $C_B^\perp$  have generator matrices where every row and column have at least two ones. This implies that the tensor codes  $C_A \otimes C_B$  and  $C_A^\perp \otimes C_B^\perp$  can be taken to have generator matrices with at least four ones in each row and column. Since these correspond to the parity check matrices of  $C_0^\perp = (C_A \otimes C_B)^\perp$  and  $C_1^\perp = (C_A^\perp \otimes C_B^\perp)^\perp$  respectively, it can be easily checked that the parity check matrices of the associated Tanner codes  $T(\mathcal{G}_0^\square, C_0^\perp)$  and  $T(\mathcal{G}_1^\square, C_1^\perp)$  also have at least four ones in every row and column.

**$H^1$  is non-trivial:** This follows immediately from dimensionality arguments. In particular, notice that  $\dim Z^1 \geq n - m$ , whereas  $\dim B^1 \leq m$ . As a result we have  $\dim H^1 \geq n - 2m = (1/2 - 2r(1-r))\Delta^2|G|$  which is  $> 0$  whenever  $r \neq 1/2$ .

**$X$  is a small-set (co)-boundary expander:** It is left to show our complexes are small-set (co)-boundary expanders. In what follows we show the co-boundary expansion case. Since the construction is symmetric, a similar proof gives small-set boundary expansion. For convenience, we first re-formulate the problem as the following technical theorem. Note that this is the analog of [LZ22, Theorem 1] where co-systolic distance is replaced with small-set co-boundary expansion. We follow their notation when possible for consistency.

**Theorem 8.2.** *Fix  $\varepsilon \in (0, 1/2)$ ,  $\gamma \in (1/2 + \varepsilon, 1)$  and  $\delta > 0$ . For any fixed large enough  $\Delta$ , if the linear codes  $C_A$  and  $C_B$  have minimum distance at least  $\delta\Delta$  and if the dual tensor code  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B = C_1^\perp$  is  $w$ -robust with  $p$ -resistance to puncturing for  $w = \Delta^{3/2-\varepsilon/2}$  and  $p = \Delta^\gamma$ ,<sup>29</sup> then the chain complex in*

<sup>27</sup>Note we are assuming for simplicity that  $r\Delta$  and  $(1-r)\Delta$  are integer valued, but these can be replaced with  $\lceil r\Delta \rceil$  and  $\Delta - \lceil r\Delta \rceil$  without substantially affecting the proof (see [LZ22]).

<sup>28</sup>Note that since  $\Delta$  is a constant with respect to our infinite family, brute force search only requires  $O(1)$  time here.

<sup>29</sup>We note that the value of  $w$  here is slightly different than in [LZ22]. This corrects a small error in the application of robust testability (Lemma 7.6) in the original work.

Equation (9):

$$X : \mathbb{F}_2^m \xrightarrow{\delta_0 := C_0^T} \mathbb{F}_2^n \xrightarrow{\delta_1 := C_1} \mathbb{F}_2^m,$$

satisfies the following isoperimetric inequality for small, minimal chains:

$$\forall x \in \mathbb{F}_2^n \text{ s.t. } x \text{ is minimal and } |x| \leq \rho_1 n : |\delta_1 x| \geq \rho_2 |x|,$$

where  $\rho_1 = \frac{\delta}{6\Delta^{3/2+\varepsilon}}$ ,  $\rho_2 = \frac{56}{\Delta^{3-2\varepsilon}}$ .

Recall that this isoperimetric condition is equivalent to  $(\rho_1, \rho_2)$ -small-set co-boundary expansion (Lemma 5.4), so this indeed proves the desired property. The proof of Theorem 8.2 closely follows the analogous proof in [LZ22] for systolic distance. The main difference is that we must track an additional set of elements consisting of vertices in  $\mathcal{G}_0^\square$  corresponding to violated constraints. Since [LZ22] only need to consider  $x \in \mathbb{F}_2^n$  that are true codewords, this is not a relevant consideration in their result. We note that throughout we set our coefficients to match those in [LZ22] for ease of comparison.

*Proof of Theorem 8.2.* We assume  $x \neq 0$ , as the theorem holds trivially otherwise. We proceed by contradiction. Assuming  $|\delta_1 x| < \rho_2 |x|$ , we will show there exists  $y \in B_1$  such that  $|x + y| < |x|$ , contradicting minimality of  $x$ .

We first lay out some relevant notation. Thinking of  $x$  as a subset of  $E_1^\square$  (the edge set of  $\mathcal{G}_1^\square$ ), we will consider the edge-induced subgraph  $\mathcal{G}_{1,x}^\square \subset \mathcal{G}_1^\square$  and denote its vertex set by  $S \subset V_1$ . Recall that each vertex  $(g, 1) \in V_1$  has a corresponding local view made up of  $|A| \cdot |B|$  incident squares, which we'll denote by:

$$\ell(g, 1) := \{(g, 1), (ag, 0), (gb, 0), (agb, 1)\} : a \in A, b \in B\}.$$

Thinking of  $x$  now as a set of squares, let  $x|_{\ell(g,1)} \in \mathbb{F}_2^{A \times B}$  denote the restriction of  $x$  to the local view of  $(g, 1)$ , and recall that  $x$  is a co-cycle exactly when these local views correspond to codewords in  $C_1^\perp = C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$ .

Since  $x$  is arbitrary in our setting (unlike [LZ22] who only consider co-cycles) we will partition the vertices of our induced subgraph into three parts:  $S = S_v \cup S_n \cup S_e$ . First, let  $S_v \subset S$  denote the set of *violated vertices*  $(g, 1) \in V_1$  whose local views  $x|_{\ell(g,1)}$  do *not* form codewords in  $C_1^\perp$ . Following [LZ22], we split the remaining vertices in  $S \setminus S_v$  into two parts based upon their degree in the induced subgraph  $\mathcal{G}_{1,x}^\square$ : the *normal vertices*  $S_n$  with degree less than  $w_2 := \Delta^{3/2-\varepsilon}$ , and the *exceptional vertices*  $S_e$  with degree at least  $w_2$ . The intuition behind this strategy is that because  $C_1^\perp$  is  $(w > w_2)$ -robust, the codewords associated to vertices in  $S_n$  have particularly nice structure: they are zero outside of a small set of at most  $w_2/(\delta\Delta)$  rows and columns. This implies that any column (respectively row) is close to a codeword in  $C_A$  (respectively  $C_B$ ) which will eventually help us apply small-set robust testability (Lemma 7.6) to prove  $x$  is close to a co-boundary (and is therefore non-minimal).

To find such a co-boundary, we'll first need to look to the other side of the complex. Broadly speaking, the idea (which is the same as in [LZ22]) is to find a vertex  $v \in V_0$  whose local view shares many (heavy) rows and columns with local views of vertices in  $S_n$ . One can then apply robustness to see that the value of  $x$  on this local view is close to a codeword in  $C_A \otimes C_B$  which can easily be translated to the desired co-boundary.

More formally, let  $E_x \subset \mathcal{G}^\cup$  denote the set of edges incident to the squares in  $x$ ,<sup>30</sup> and call an edge *heavy* if it is incident to at least  $\delta\Delta - \Delta^{1/2-\varepsilon}/\delta$  squares in  $x$ . We will consider the set of vertices  $T \subset V_0$  which are adjacent to  $S_n \subset V_1$  through a heavy edge in the graph  $\mathcal{G}^\cup$ . Given  $v \in T$ , note that every heavy edge with an element in  $S_n$  corresponds to a row or column that is shared in their local view (and is therefore close to

<sup>30</sup>In particular for any square  $\{(g, 0), (ag, 1), (gb, 1), (agb, 0)\} \in x$ , add its four edges  $\{(g, 0), (ag, 1)\}$ ,  $\{(g, 0), (bg, 1)\}$ ,  $\{(agb, 0), (ag, 1)\}$ , and  $\{(agb, 0), (gb, 1)\}$  to  $E_x$ .

a codeword of  $C_B$  or  $C_A$  respectively). The goal is therefore to show that there exists a vertex in  $T$  that is adjacent to many elements in  $S_n$  through heavy edges, while simultaneously adjacent to few ‘bad’ vertices in  $S_e$  and  $S_v$ . This will allow us to apply robustness against puncture to find a co-boundary that reduces the weight of  $x$ . We formalize these statements below in the following two claims.

**Claim 8.3** (Modification of [LZ22, Claim 13]). *There exist  $h_1 \geq \Omega(\Delta)$ ,  $d_1 \leq O(\Delta^{1/2+\varepsilon})$ , and  $v \in T$  such that  $v$  is incident to at least  $h_1$  heavy edges and adjacent to at most  $d_1$  vertices of  $S_e \cup S_v$ .*

**Claim 8.4** (Summary of paragraph following [LZ22, Claim 13]). *For all sufficiently large<sup>31</sup>  $\Delta$ , if there exists a vertex  $v \in V_0$  incident to  $h_1 \geq \Omega(\Delta)$  heavy edges and at most  $d_1 \leq O(\Delta^{1/2+\varepsilon})$  vertices of  $S_e \cup S_v$ , then we can find a vector  $y \in B_1$  such that  $|x + y| < |x|$ .*

Together, Claim 8.3 and Claim 8.4 complete the proof of Theorem 8.2, as they promise the existence of some  $y \in B_1$  such that  $|x + y| < |x|$ , violating minimality of  $x$ . Thus it is left to prove the claims. While Claim 8.4 follows largely from arguments in [LZ22], it is helpful to present first to motivate the more technical proof of Claim 8.3.

*Proof of Claim 8.4.* Recall we are given an element  $v \in V_0$  which is incident to at least  $h_1 \geq \Omega(\Delta)$  heavy edges and adjacent to at most  $d_1 \leq O(\Delta^{1/2+\varepsilon})$  vertices in  $S_e \cup S_v$ . We consider the local view of  $x$  around  $v$  (considered as an element of  $\mathcal{G}_0^\square$ ), denoted  $x_v \in \mathbb{F}_2^{A \times B}$  here for notational simplicity. Because at most  $d_1$  vertices adjacent to  $v$  in  $\mathcal{G}^\cup$  are exceptional or violated (as considered in  $\mathcal{G}_{1,x}^\square$ ), one can find  $A' \subset A, B' \subset B$  with  $|A'| = |B'| \geq \Delta - d_1$ , such that  $A'$  and  $B'$  are indexed by either normal vertices, or vertices not in  $S$ . Furthermore, since  $d_1 \leq \Delta^\gamma$  for large enough  $\Delta$ , we also have by robustness to puncture that the restricted dual tensor code  $(C_1^\perp)' := C_{A'} \otimes \mathbb{F}_2^{B'} + \mathbb{F}_2^{A'} \otimes C_{B'}$  is  $w$ -robust.

Let  $x'_v$  be the restriction of  $x_v$  in  $A' \times B'$ . Recall each column (row) of the local view of a normal vertex is at most  $w_2/(\delta\Delta)$  away from a codeword by  $w$ -robustness. Then since each column (row) of  $x'_v$  is indexed by either a normal vertex or a vertex whose local view is all zero (i.e. not in  $S$ ), every column (respectively row) of  $x'_v$  is at most  $w_2/(\delta\Delta) = \Delta^{1/2-\varepsilon}/\delta$  away from a codeword in  $C_{A'}$  (respectively  $C_{B'}$ ). Since there are at most  $\Delta$  rows and columns, this means that  $x'_v$  is at most  $\Delta^{3/2-\varepsilon}/\delta$  away from either  $C_{A'} \otimes \mathbb{F}_2^{B'}$  or  $\mathbb{F}_2^{A'} \otimes C_{B'}$ , and moreover that:

$$d(x'_v, C_{A'} \otimes \mathbb{F}_2^{B'}) + d(x'_v, \mathbb{F}_2^{A'} \otimes C_{B'}) \leq 2\Delta^{3/2-\varepsilon}/\delta \leq w$$

for sufficiently large  $\Delta$ . Because  $(C_1^\perp)'$  is  $w$ -robust, we can apply small-set robust testability (Lemma 7.6) to infer that  $x'_v$  is close to some codeword  $c' \in C_{A'} \otimes C_{B'}$ :

$$d(x'_v, c') \leq \frac{3}{2} \left( d(x'_v, C_{A'} \otimes \mathbb{F}_2^{B'}) + d(x'_v, \mathbb{F}_2^{A'} \otimes C_{B'}) \right) \leq 3 \frac{\Delta^{3/2-\varepsilon}}{\delta}.$$

Finally, since the total number of punctured rows and columns is less than the code distance for large enough  $\Delta$ , we can extend  $c'$  uniquely to a codeword  $c \in C_A \otimes C_B$ . Taking into account the rows and columns added in this process, the distance from  $x_v$  to  $c$  then becomes at most  $d(x_v, c) \leq d(x'_v, c') + 2d_1\Delta \leq O(\Delta^{3/2+\varepsilon}) < o(\Delta^2)$  since  $\varepsilon < 1/2$ .

On the other hand, because  $v$  is incident to  $\Omega(\Delta)$  heavy edges, the weight  $|x_v| = \Theta(\Delta^2)$ . Thus for large enough  $\Delta$ , it must be the case that flipping  $c$  strictly reduces the weight of  $x$ . More precisely, set  $y$  to be  $c$  on the local view  $x_v$  and 0 elsewhere, then we have  $|x + y| < |x|$ . Since  $c \in C_A \otimes C_B = C_0$ ,  $y \in B^1$  is indeed a co-boundary which completes the proof.  $\square$

The only thing left is to show that our main technical claim actually holds, the existence of a vertex with many heavy edges that is adjacent to few violated or exceptional vertices. The proof technique is similar to that of [LZ22, Claim 13], and mostly boils down to proving that  $S_v$  and  $S_e$  are small compared to  $S_n$ .

<sup>31</sup>Here we mean in terms of  $r, \varepsilon$ , and  $\gamma$ , so  $\Delta$  remains constant with respect to the infinite family.

*Proof of Claim 8.3.* We split the proof into the following three claims. First, we claim  $T$  is non-empty.

**Claim 8.5.**  $|T| > 0$ .

With this in mind, let  $\alpha, \beta = \Theta(1)$  be constants to be set later in the proof. Following [LZ22], we claim that a reasonable fraction of  $T$  is incident to many heavy edges:

**Claim 8.6** ([LZ22, Claim 12]). *At least an  $\alpha/2$  fraction of vertices in  $T$  are incident to at least  $h_1 = \alpha\Delta$  heavy edges,*

and further that at most some smaller fraction is adjacent to greater than  $d_1$  violated and exceptional vertices:

**Claim 8.7** ([LZ22, Paragraph between Claim 4.10 and Claim 4.11]). *At most an  $\alpha/4$  fraction of vertices in  $T$  are incident to more than  $d_1 = \frac{4\beta}{\alpha}\Delta^{1/2+\varepsilon}$  vertices of  $S_e \cup S_v$ .*

Combining these claims implies at least an  $\alpha/4$  fraction of vertices satisfy the requirements of Claim 8.3. Since  $T$  is non-empty, this must apply to at least one  $v \in T$  which gives the desired result.

The key to proving all three claims lies in showing that the number of vertices in  $S_e \cup S_v$  is small compared to  $S_n$ . We will show that  $S_v$  can be upper bounded by taking  $\rho_2$  sufficiently small, and  $S_e$  can be upper bounded by the expander mixing lemma as in [LZ22].

**Lemma 8.8** (Modification of [LZ22, Claim 6]). *The number of exceptional and violated vertices is at most*

$$|S_e \cup S_v| \leq \frac{64}{\Delta^{1-2\varepsilon}}|S|. \quad (10)$$

*On the other hand, the number of normal vertices is at least*

$$|S_n| \geq \left(1 - \frac{64}{\Delta^{1-2\varepsilon}}\right)|S|. \quad (11)$$

*Proof.* The latter fact follows immediately from the former and recalling that  $S_e, S_v$ , and  $S_n$  partition  $S$ . We now show  $|S_v|$  is small. Note that by assumption we have that

$$|S_v| \leq |\delta_1 x| < \rho_2 |x|,$$

since  $\delta_1$  is the parity-check matrix of  $\mathcal{C}_1$  and every violated vertex corresponds to at least one violated constraint in  $\mathcal{C}_1$ . Because  $V_1$  has degree  $\Delta^2$  in  $\mathcal{G}_1^\square$  (i.e. each vertex sits in  $\Delta^2$  squares), we also have  $|x| \leq \Delta^2|S|/2$ . Altogether this gives

$$|S_v| < \rho_2 \Delta^2 |S| / 2 = 28|S| / \Delta^{1-2\varepsilon}$$

for our choice of  $\rho_2$ .

Now we show  $|S_e|$  is small. The degree of each non-violated vertex is at least  $\delta\Delta$  because the local view corresponds to a non-zero codeword in  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$ . This implies  $|S_n| + |S_e| \leq \frac{2|x|}{\delta\Delta}$ . Combining this with our bound on  $|S_v|$  gives

$$|S| = |S_n| + |S_e| + |S_v| \leq \left(\rho_2 + \frac{2}{\delta\Delta}\right)|x| \leq \frac{4}{\delta\Delta}|x| \quad (12)$$

where the second inequality holds for large enough  $\Delta$  (recalling that  $\rho_2 = O(\Delta^{-3+2\varepsilon})$ ). Applying the expander mixing lemma to  $E(S_e, S)$ , we then obtain

$$\begin{aligned} |E(S_e, S)| &\leq \frac{\Delta^2}{|V_1|} |S_e| |S| + 4\Delta \sqrt{|S_e| |S|} \\ &\leq \frac{4\Delta}{\delta|V_1|} |x| |S_e| + 4\Delta \sqrt{|S_e| |S|} \\ &= \frac{1}{3} \Delta^{3/2-\varepsilon} |S_e| + 4\Delta \sqrt{|S_e| |S|} \end{aligned}$$

where we have used the assumption that  $|x| \leq \frac{\delta n}{6\Delta^{3/2+\varepsilon}}$  and the fact that  $|V_1| = |G| = 2n/\Delta^2$ . On the other hand, by definition of exceptional vertices we have that  $|E(S_e, S)| \geq \Delta^{3/2-\varepsilon}|S_e|$ . Combining the inequalities we obtain  $|S_e| \leq 36|S|/\Delta^{1-2\varepsilon}$ , and plugging in our bound on  $|S_v|$  then gives  $|S_e \cup S_v| \leq 64|S|/\Delta^{1-2\varepsilon}$  as desired.  $\square$

Finally, we prove Claim 8.5, Claim 8.6, and Claim 8.7, completing the result. The latter two follow essentially as in [LZ22] (replacing  $S_e$  with  $S_e \cup S_v$ ), but we give the proofs here for completeness.

*Proof of Claim 8.5.* We wish to prove  $T$  is non-empty. First, recall that since  $x \neq 0$  by assumption,  $|S| > 0$ . By Lemma 8.8, we then have  $|S_n| > 0$  as well. We now argue that every vertex in  $S_n$  is incident to at least one heavy edge. Since  $S_n$  is non-empty, this implies  $T$  is non-empty as desired.

To see each vertex in  $S_n$  has a heavy edge, recall the local view of each normal vertex is a codeword in  $C_A \otimes \mathbb{F}_2^B + \mathbb{F}_2^A \otimes C_B$  with weight less than  $w = \Delta^{3/2-\varepsilon}$ . Because the dual tensor code is  $w$ -robust, each column (respectively row) is within  $\Delta^{1/2-\varepsilon}/\delta$  of a codeword in  $C_A$  (respectively  $C_B$ ). Since these codes all have distance at least  $\delta\Delta$ , there must be a row or column with at least  $\delta\Delta - \Delta^{1/2-\varepsilon}/\delta$  ones which exactly corresponds to a heavy edge. We note this fact also implies the total number of heavy edges is at least  $|S_n|$ , which will be useful later on.  $\square$

*Proof of Claim 8.6.* Now that we have confirmed the existence of  $T$ , we want to show it is incident to many heavy edges. To do so, we'll argue that  $T$  is small compared to the number of heavy edges.

To start, we show that  $|T| \leq \frac{64}{\delta^2\Delta}|S|$ . The proof is the same as [LZ22, Claim 11], but we give it here for completeness. First, note that by the expander mixing lemma on  $\mathcal{G}^\cup$  (which is the double cover of a  $4\sqrt{\Delta}$ -spectral expander) we have:

$$\begin{aligned} |E(S, T)| &\leq \frac{2\Delta}{|G|}|S||T| + 4\sqrt{\Delta}\sqrt{|S||T|} \\ &\leq \frac{2\Delta^{1/2-\varepsilon}}{3}|T| + 4\sqrt{\Delta}\sqrt{|S||T|} \end{aligned}$$

where as in Lemma 8.8 we have again used the fact that

$$|S| \leq \frac{4}{\delta\Delta}|x| \leq \frac{2}{3\Delta^{5/2+\varepsilon}}n = \frac{1}{3\Delta^{1/2+\varepsilon}}|G|.$$

On the other hand, since each vertex  $v \in T$  is incident to at least one heavy edge  $e$  by definition,  $v$  (and  $e$ ) are contained in at least  $\delta\Delta - \Delta^{1/2-\varepsilon}/\delta$  squares in  $x$ . Since each of these contains an additional (unique) edge incident to  $v$ , we also have the following lower bound

$$|E(S, T)| \geq (\delta\Delta - \Delta^{1/2-\varepsilon}/\delta)|T|.$$

Combining these inequalities one can check that  $|T| \leq \frac{64}{\delta^2\Delta}|S|$  for large enough  $\Delta$  as desired.

With this in hand, recall from the proof of Claim 8.5 that the total number of heavy edges in  $E_x$  is at least  $|S_n| \geq (1 - \frac{64}{\Delta^{1-2\varepsilon}})|S|$  (where the inequality is given by Lemma 8.8). Together, this implies the average number of heavy edges incident to a vertex in  $T$  is at least:

$$\frac{|S_n|}{|T|} \geq \frac{\delta^2\Delta}{64} \left(1 - \frac{64}{\Delta^{1-2\varepsilon}}\right) =: 2\alpha\Delta. \quad (13)$$

Finally given that the average degree is at least  $2\alpha\Delta$ , we want to show there is some fraction of vertices with degree  $\geq \alpha\Delta$ . This is immediate from recalling that the maximum degree of  $\mathcal{G}^\cup$  (and thus  $T$ ) is  $2\Delta$ , which implies at least an  $\alpha/2$  fraction of vertices in  $T$  are incident to at least  $\alpha\Delta$  heavy edges as desired.  $\square$

*Proof of Claim 8.7.* Finally, we want to show there are few edges between  $T$  and  $S_e \cup S_v$ . This follows from the fact that both sets are small, and the underlying graph  $\mathcal{G}^\cup$  is the double cover of a  $4\sqrt{\Delta}$ -expander on  $|G| = \frac{2n}{\Delta^2}$  vertices. In particular, combining the expander mixing lemma with our bounds from Lemma 8.8 gives:

$$\begin{aligned} E(S_e \cup S_v, T) &\leq \frac{2\Delta}{|G|} |S_e \cup S_v| |T| + 4\sqrt{\Delta} \sqrt{|T| |S_e \cup S_v|} \\ &\leq \frac{128\Delta^{2\varepsilon}}{|G|} |S| |T| + 32\Delta^\varepsilon \sqrt{|T| |S|}. \end{aligned}$$

Recall that  $|S| \leq \frac{1}{3\Delta^{1/2+\varepsilon}} |G|$ . Further, since each normal vertex is adjacent to  $T$  and the degree of  $T$  is at most  $2\Delta$ , we have  $(1 - \frac{64}{\Delta^{1-2\varepsilon}}) |S| \leq |S_n| \leq 2\Delta |T|$ , and thus for large enough  $\Delta$  that  $|S| \leq 4\Delta |T|$ . Altogether we therefore have:

$$E(S_e \cup S_v, T) \leq \frac{128}{3\Delta^{1/2-\varepsilon}} |T| + 64\Delta^{1/2+\varepsilon} |T| \leq \beta \Delta^{1/2+\varepsilon} |T|$$

where  $\beta = 64 + \frac{128}{3\Delta}$ . As a result, at most an  $\alpha/4$  fraction of vertices in  $T$  are incident to more than  $d_1 = \frac{4\beta}{\alpha} \Delta^{1/2+\varepsilon}$  vertices of  $S_e \cup S_v$  as desired, which completes the proof of Claim 8.3 and Theorem 8.2 in turn.  $\square$

$\square$   
 $\square$   
 $\square$

Putting everything together, we now prove the existence of an explicit family of SS-HDX.

*Proof of Theorem 8.1.* Fix any  $r \in (0, 1/2)$ ,  $\varepsilon \in (0, 1/2)$ ,  $\gamma \in (1/2 + \varepsilon, 1)$ , and  $\delta \in (0, 1)$  satisfying  $-\delta \log \delta - (1 - \delta) \log(1 - \delta) < r$ , and let  $\Delta = \Delta(r, \varepsilon/2, \gamma, \delta) \in \mathbb{N}$  be sufficiently large that the guarantees of Corollary 7.9 and Theorem 8.2 are met. Brute forcing over pairs of length  $\Delta$  codes  $C_A, C_B$  of dimensions  $r\Delta$  and  $(1 - r)\Delta$  respectively, Corollary 7.9 promises we can find in  $O_\Delta(1)$  time codes  $C_A, C_B$  such that:

1.  $\dim C_A = \lfloor r\Delta \rfloor$  and  $\dim C_B = \Delta - \dim C_A$ ,
2. The distances of  $C_A, C_B, C_A^\perp, C_B^\perp$  are all at least  $\delta\Delta$ ,
3. Both dual tensor codes  $C_0^\perp = (C_A \otimes C_B)^\perp$  and  $C_1^\perp = (C_A^\perp \otimes C_B^\perp)^\perp$  are  $\Delta^{3/2-\varepsilon/2}$ -robust with  $\Delta^\gamma$ -resistance to puncturing.
4.  $C_A, C_B, C_A^\perp$ , and  $C_B^\perp$  have generator matrices where every row and column have at least two ones.

Following the construction and the discussion earlier this section, the Tanner maps resulting from these codes and the explicit left-right Cayley complexes of [DEL<sup>+</sup>21] give an explicit family of chain complexes with degree between 3 and  $2\Delta^2$  and non-trivial co-homology. Furthermore each individual complex in the family satisfies the requirements of Theorem 8.2 in both directions, so by symmetry the complexes are  $(\rho_1, \rho_2)$ -small-set HDX for  $\rho_1 = \frac{\delta}{6\Delta^{3/2+\varepsilon}}$  and  $\rho_2 = \frac{56}{\Delta^{3-2\varepsilon}}$ . This concludes the proof of Theorem 8.1.  $\square$

**Remark 8.9.** We note that the proof of Theorem 8.2 actually gives a stronger guarantee than small-set (co)-boundary expansion. In particular, because the boundary  $y$  that reduces the weight of  $x$  is supported on a local view of a single vertex, the result actually gives an isoperimetric inequality for the broader class of small, locally minimal functions:

$$\forall x \in \mathbb{F}_2^n \text{ s.t. } x \text{ is locally minimal and } |x| \leq \rho_1 n : |\delta_1 x| \geq \rho_2 |x|,$$

where  $x$  is locally minimal if  $|x| \leq |x + \delta_0(e_v)|$  for all basis vectors  $e_v \in \mathbb{F}_2^m$ . As discussed in Section 3.1, this stronger isoperimetric inequality has seen prior use in the topological HDX literature [KKL14, EK16, KM18, KM21] as well as in recent work on c3-LTCs [LH22a] and qLDPC codes [LH22b].

## Acknowledgements

The authors thank Noah Fleming, Sam Hopkins, and Russell Impagliazzo for helpful discussion on reductions within the Sum-of-Squares hierarchy, Amy Kanne for helpful discussions on qLDPC codes and [LZ22], and Tali Kaufman for many fruitful discussions on high dimensional expansion. The authors also thank Sam Hopkins, Shachar Lovett, and Anthony Ostuni for helpful comments on an earlier version of the manuscript.

## References

- [AAT05] Mikhail Alekhovich, Sanjeev Arora, and Iannis Tourlakis, *Towards strong nonapproximability results in the lovász-schrijver hierarchy*, Proceedings of the thirty-seventh annual ACM symposium on theory of computing, 2005, pp. 294–303.
- [AC88] Noga Alon and Fan RK Chung, *Explicit construction of linear sized tolerant networks*, Discrete Mathematics **72** (1988), no. 1-3, 15–19.
- [AC02] Noga Alon and Michael Capalbo, *Explicit unique-neighbor expanders*, The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings., IEEE, 2002, pp. 73–79.
- [AJT19] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani, *Approximating constraint satisfaction problems on high-dimensional expanders*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 180–201.
- [AL20] Vedat Levi Alev and Lap Chi Lau, *Improved analysis of higher order random walks and applications*, arXiv preprint arXiv:2001.02827 (2020).
- [Alo86] Noga Alon, *Eigenvalues and expanders*, Combinatorica **6** (1986), no. 2, 83–96.
- [ALO20] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, arXiv preprint arXiv:2001.00303 (2020).
- [ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant, *Log-concave polynomials ii: high-dimensional walks and an fpras for counting bases of a matroid*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 1–12.
- [BBK<sup>+</sup>21] Mitali Bafna, Boaz Barak, Pravesh K. Kothari, Tselil Schramm, and David Steurer, *Playing unique games on certified small-set expanders*, STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021 (Samir Khuller and Virginia Vassilevska Williams, eds.), ACM, 2021, pp. 1629–1642.
- [BCK15] Boaz Barak, Siu On Chan, and Pravesh K Kothari, *Sum of squares lower bounds from pairwise independence*, Proceedings of the forty-seventh annual ACM symposium on Theory of computing, 2015, pp. 97–106.
- [BE21] Nikolas P Breuckmann and Jens Niklas Eberhardt, *Quantum low-density parity-check codes*, PRX Quantum **2** (2021), no. 4, 040101.

- [BGIP01] Sam Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi, *Linear gaps between degrees for the polynomial calculus modulo distinct primes*, Journal of Computer and System Sciences **62** (2001), no. 2, 267–289.
- [BGMT12] Siavosh Benabbas, Konstantinos Georgiou, Avner Magen, and Madhur Tulsiani, *Sdp gaps from pairwise independence*, Theory of Computing **8** (2012), no. 1, 269–289.
- [BHKL21] Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett, *Hypercontractivity on high dimensional expanders*, arXiv preprint arXiv:2111.09444 (2021).
- [BHKL22] ———, *High dimensional expanders: Eigenstripping, pseudorandomness, and unique games*, Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2022, pp. 1069–1128.
- [BIK<sup>+</sup>96] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák, *Lower bounds on hilbert’s nullstellensatz and propositional proofs*, Proceedings of the London Mathematical Society **3** (1996), no. 1, 1–26.
- [BKS18] Boaz Barak, Pravesh K Kothari, and David Steurer, *Small-set expansion in shortcode graph and the 2-to-2 conjecture*, arXiv preprint arXiv:1804.08662 (2018).
- [BS14] Boaz Barak and David Steurer, *Sum-of-squares proofs and the quest toward optimal algorithms*, arXiv preprint arXiv:1404.5236 (2014).
- [BSW99] Eli Ben-Sasson and Avi Wigderson, *Short proofs are narrow—resolution made simple*, Proceedings of the thirty-first annual ACM symposium on Theory of computing, 1999, pp. 517–526.
- [CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo, *Using the groebner basis algorithm to find proofs of unsatisfiability*, Proceedings of the twenty-eighth annual ACM symposium on Theory of computing, 1996, pp. 174–183.
- [Cha16] Siu On Chan, *Approximation resistance from pairwise-independent subgroups*, Journal of the ACM (JACM) **63** (2016), no. 3, 1–32.
- [CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev, *Integrality gaps for sherali-adams relaxations*, Proceedings of the forty-first annual ACM symposium on Theory of computing, 2009, pp. 283–292.
- [CS96] A Robert Calderbank and Peter W Shor, *Good quantum error-correcting codes exist*, Physical Review A **54** (1996), no. 2, 1098.
- [DD19] Yotam Dikstein and Irit Dinur, *Agreement testing theorems on layered set systems*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 1495–1524.
- [DDFH18] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha, *Boolean function analysis on high-dimensional expanders*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [DEL<sup>+</sup>21] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes, *Locally testable codes with constant rate, distance, and locality*, arXiv preprint arXiv:2111.04808 (2021).

- [DFHT20] Irit Dinur, Yuval Filmus, Prahladh Harsha, and Madhur Tulsiani, *Explicit sos lower bounds from high-dimensional expanders*, arXiv preprint arXiv:2009.05218 (2020).
- [DK17] Irit Dinur and Tali Kaufman, *High dimensional expanders imply agreement expanders*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 974–985.
- [DKK<sup>+</sup>18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra, *On non-optimally expanding sets in grassmann graphs*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 940–951.
- [DKK<sup>+</sup>18b] ———, *Towards a proof of the 2-to-1 games conjecture?*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, 2018, pp. 376–389.
- [DSW06] Irit Dinur, Madhu Sudan, and Avi Wigderson, *Robust local testability of tensor products of ldpc codes*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Springer, 2006, pp. 304–315.
- [EH17] Lior Eldar and Aram W Harrow, *Local hamiltonians whose ground states are hard to approximate*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 427–438.
- [EK16] Shai Evra and Tali Kaufman, *Bounded degree cosystolic expanders of every dimension*, Proceedings of the forty-eighth annual ACM symposium on Theory of Computing, 2016, pp. 36–48.
- [EKZ20] Shai Evra, Tali Kaufman, and Gilles Zémor, *Decodable quantum ldpc codes beyond the square root distance barrier using high dimensional expanders*, 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2020, pp. 218–227.
- [FKP19] Noah Fleming, Pravesh Kothari, and Toniann Pitassi, *Semialgebraic proofs and efficient algorithm design*, Foundations and Trends in Theoretical Computer Science, 2019.
- [GLL21] Tom Gur, Noam Lifshitz, and Siqi Liu, *Hypercontractivity on high dimensional expanders*, arXiv preprint arXiv:2111.09375 (2021).
- [GMPT10] Konstantinos Georgiou, Avner Magen, Toniann Pitassi, and Iannis Tourlakis, *Integrality gaps of  $2-o(1)$  for vertex cover sdps in the lovász–schrijver hierarchy*, SIAM Journal on Computing **39** (2010), no. 8, 3553–3570.
- [Gri98] Dima Grigoriev, *Tseitin’s tautologies and lower bounds for nullstellensatz proofs*, Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No. 98CB36280), IEEE, 1998, pp. 648–652.
- [Gri01a] ———, *Complexity of positivstellensatz proofs for the knapsack*, computational complexity **10** (2001), no. 2, 139–154.
- [Gri01b] ———, *Linear lower bound on degrees of positivstellensatz calculus proofs for the parity*, Theoretical Computer Science **259** (2001), no. 1-2, 613–622.
- [Gro83] Mikhael Gromov, *Filling riemannian manifolds*, Journal of Differential Geometry **18** (1983), no. 1, 1–147.

- [Gro10] Mikhail Gromov, *Singularities, expanders and topology of maps. part 2: From combinatorics to topology via algebraic isoperimetry*, Geometric and Functional Analysis **20** (2010), no. 2, 416–526.
- [HHO21] Matthew B Hastings, Jeongwan Haah, and Ryan O’Donnell, *Fiber bundle codes: breaking the  $n/2$  polylog( $n$ ) barrier for quantum ldpc codes*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1276–1288.
- [JMO<sup>+</sup>21] Fernando Granha Jeronimo, Tushant Mittal, Ryan O’Donnell, Pedro Paredes, and Madhur Tulsiani, *Explicit abelian lifts and quantum ldpc codes*, arXiv preprint arXiv:2112.01647 (2021).
- [JST21] Fernando Granha Jeronimo, Shashank Srivastava, and Madhur Tulsiani, *Near-linear time decoding of ta-shma’s codes via splittable regularity*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1527–1536.
- [Kho02] Subhash Khot, *On the power of unique 2-prover 1-round games*, Proceedings of the thirty-fourth annual ACM symposium on Theory of computing, 2002, pp. 767–775.
- [KKL14] Tali Kaufman, David Kazhdan, and Alexander Lubotzky, *Ramanujan complexes and bounded degree topological expanders*, 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, IEEE, 2014, pp. 484–493.
- [KKL16] ———, *Isoperimetric inequalities for ramanujan complexes and topological expanders*, Geometric and Functional Analysis **26** (2016), no. 1, 250–287.
- [KM18] Tali Kaufman and David Mass, *Cosystolic expanders over any abelian group.*, Electron. Colloquium Comput. Complex., vol. 25, 2018, p. 134.
- [KM21] ———, *Unique-neighbor-like expansion and group-independent cosystolic expansion*, 32nd International Symposium on Algorithms and Computation (ISAAC 2021), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [KMMS18] Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra, *Small set expansion in the johnson graph.*, Electronic Colloquium on Computational Complexity (ECCC), vol. 25, 2018, p. 78.
- [KMOW17] Pravesh K Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer, *Sum of squares lower bounds for refuting any csp*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 132–145.
- [KMS17] Subhash Khot, Dor Minzer, and Muli Safra, *On independent sets, 2-to-2 games, and grassmann graphs*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 576–589.
- [KMS18] ———, *Pseudorandom sets in grassmann graph have near-perfect expansion*, 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2018, pp. 592–601.
- [KO20] Tali Kaufman and Izhar Oppenheim, *High order random walks: Beyond spectral gap*, Combinatorica (2020), 1–37.
- [KT21] Tali Kaufman and Ran J. Tessler, *New cosystolic expanders from tensors imply explicit quantum LDPC codes with  $\Omega(\sqrt{n}) \log^k(n)$  distance*, STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021, 2021, pp. 1317–1329.

- [LH22a] Ting-Chun Lin and Min-Hsiu Hsieh, *c3-local testable codes from lossless expanders*, arXiv preprint arXiv:2201.11369 (2022).
- [LH22b] ———, *Good quantum ldpc codes with linear time decoder from lossless expanders*, arXiv preprint arXiv:2203.03581 (2022).
- [LM06] Nathan Linial\* and Roy Meshulam\*, *Homological connectivity of random 2-complexes*, *Combinatorica* **26** (2006), no. 4, 475–487.
- [LSV05] Alexander Lubotzky, Beth Samuels, and Uzi Vishne, *Explicit constructions of ramanujan complexes of type ad*, *European Journal of Combinatorics* **26** (2005), no. 6, 965–993.
- [LZ22] Anthony Leverrier and Gilles Zémor, *Quantum tanner codes*, arXiv preprint arXiv:2202.13641 (2022).
- [Mor94] Moshe Morgenstern, *Existence and explicit constructions of  $q+1$  regular ramanujan graphs for every prime power  $q$* , *Journal of Combinatorial Theory, Series B* **62** (1994), no. 1, 44–62.
- [MS09] Claire Mathieu and Alistair Sinclair, *Sherali-adams relaxations of the matching polytope*, *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 293–302.
- [PK21a] Pavel Panteleev and Gleb Kalachev, *Asymptotically good quantum and locally testable classical ldpc codes*, arXiv preprint arXiv:2111.03654 (2021).
- [PK21b] ———, *Quantum ldpc codes with almost linear minimum distance*, *IEEE Transactions on Information Theory* **68** (2021), no. 1, 213–229.
- [Pra22] Kevin Pratt, Personal communication, March 2022.
- [RS10] Prasad Raghavendra and David Steurer, *Graph expansion and the unique games conjecture*, *Proceedings of the forty-second ACM symposium on Theory of computing*, 2010, pp. 755–764.
- [RST12] Prasad Raghavendra, David Steurer, and Madhur Tulsiani, *Reductions between expansion problems*, *2012 IEEE 27th Conference on Computational Complexity*, IEEE, 2012, pp. 64–73.
- [Sch08] Grant Schoenebeck, *Linear level lasserre lower bounds for certain  $k$ -csp*s, *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, IEEE, 2008, pp. 593–602.
- [Ste96] Andrew M Steane, *Error correcting codes in quantum theory*, *Physical Review Letters* **77** (1996), no. 5, 793.
- [STT07] Grant Schoenebeck, Luca Trevisan, and Madhur Tulsiani, *Tight integrality gaps for lovász-schrijver lp relaxations of vertex cover and max cut*, *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, 2007, pp. 302–310.
- [Tan81] R Tanner, *A recursive approach to low complexity codes*, *IEEE Transactions on information theory* **27** (1981), no. 5, 533–547.
- [Tul09] Madhur Tulsiani, *Csp gaps and reductions in the lasserre hierarchy*, *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 303–312.
- [Tul21] ———, *On sos algorithms and lower bounds for constraint satisfaction problems over high-dimensional expanders.*, SoS+TCS Reading Group, 2021.

## A Existence of Good Base Codes

In this section we prove Corollary 7.9, the existence of base codes  $C_A$  and  $C_B$  with the properties needed for our SS-HDX construction in Theorem 8.1. We restate the result here for convenience.

**Corollary A.1.** *Fix  $r \in (0, 1/2)$ ,  $\varepsilon \in (0, 1/2)$ ,  $\gamma \in (1/2 + \varepsilon, 1)$  and  $\delta > 0$  satisfying  $-\delta \log \delta - (1 - \delta) \log(1 - \delta) < r$ . When  $k$  is large enough, there exist codes  $C_A$  and  $C_B$  of length  $\Delta$  such that*

1.  $\dim C_A = \lfloor r\Delta \rfloor$  and  $\dim C_B = \Delta - \dim C_A$
2. The distances of  $C_A, C_B, C_A^\perp, C_B^\perp$  are all at least  $\delta\Delta$
3. Both dual tensor codes  $C_0^\perp = (C_A \otimes C_B)^\perp$  and  $C_1^\perp = (C_A^\perp \otimes C_B^\perp)^\perp$  are  $\Delta^{3/2-\varepsilon}$ -robust with  $\Delta^\gamma$ -resistance to puncturing
4.  $C_A, C_B, C_A^\perp$ , and  $C_B^\perp$  have generator matrices where every row and column have at least two ones.

*Proof.* We assume for notational simplicity that  $r\Delta$  and  $(1 - r)\Delta$  are integral (the proof is essentially the same without this assumption). We will argue that all four properties are satisfied with probability going to one (as  $\Delta$  becomes large) under some distribution for the generation of  $C_A$  and  $C_B$ . By a union bound, a pair satisfying all properties must then exist for large enough  $\Delta$ .

Consider the distribution over codes  $C_A$  and  $C_B$  given by generating  $C_A$  by a uniformly random  $r\Delta \times \Delta$  generator matrix, and  $C_B^\perp$  from an independent uniformly random  $r\Delta \times \Delta$  generator matrix. Leverrier and Zémor [LZ22] prove that the first three conditions occur with probability going to one under this distribution (see [LZ22, Theorem 17]), so we need only show the last condition holds.

This follows easily from a few basic observations. Let  $r_0 \in (0, 1)$  be any constant. First, observe that conditioned on being full rank, a uniformly random  $r_0\Delta \times \Delta$  generator matrix corresponds to a uniformly random subspace of dimension  $r_0\Delta$ , and furthermore that such a matrix is full rank with probability going to 1 as  $\Delta$  grows large. Second, note that by a Chernoff and union bound, the probability this random generator matrix has any row or column with less than two ones also quickly goes to zero. This implies that for any fixed  $r_0$ , as  $\Delta$  grows large the probability that a random subspace of dimension  $r_0\Delta$  has a generator matrix satisfying condition 4 goes to 1.

Since  $C_A$  and  $C_B^\perp$  are generated by uniformly random  $r\Delta \times \Delta$  generator matrices, they clearly satisfy condition 4 with high probability. The trick is then simply to notice that (conditioned on full rank),  $C_A^\perp$  and  $C_B$  are uniformly random subspaces of dimension  $(1 - r)\Delta$ , and therefore also satisfy condition 4 with probability going to one by the above observation.  $\square$